

NASA/SP-20240014019

Version 2.0

November 2024




NASA RISK MANAGEMENT HANDBOOK

Version 2.0

PART 1

GUIDANCE FOR AN INTEGRATED FRAMEWORK
OF RISK LEADERSHIP AND MANAGEMENT
ACROSS THE NASA ORGANIZATION:
CONCEPTS AND PRINCIPLES



NASA SP-20240014019
Version 2

NASA SP-20240014019 Version 2 supersedes NASA SP-2011-3422 Version 1 dated November 2011.

Cover image: Brown dwarf identified by James Webb Space Telescope in star cluster IC 348, about 1000 light-years away

Comments, questions, and suggestions regarding this document can be sent to:

Dr. Mary Coan Skow
Agency Risk Management Officer
NASA Headquarters, Office of Safety and Mission Assurance (OSMA)
mary.coan@nasa.gov

National Aeronautics and Space Administration
NASA Headquarters Washington, D.C. 20546
November 2024

ACKNOWLEDGMENTS

The authors express their gratitude to NASA's Leadership for their support and encouragement in developing this document, the second edition of the NASA Risk Management Handbook. Building upon the work that resulted in the first edition of this handbook, the development effort leading to this two-volume handbook was conducted in stages and was supported through reviews and discussions by the Agency Risk Management Working Group (ARMWG) and the Agency Risk Management Officer (ARMO) Team and by the Agency reviewers of the pre-publication version of the handbook. The authors also acknowledge the contribution of Michael Yau of ASCA Inc. to the development of this handbook.

AUTHORS:

Homayoon Dezfali, NASA Headquarters

Sergio Guarro, ASCA Inc.

Chris Everett, Idaho National Laboratory

Allan Benjamin, Quality Assurance & Risk Management Services, Inc.

Mary Coan Skow, NASA Headquarters

REVIEWERS:

Reviewers who provided comments on the drafts leading up to this version are

Tahani Amer, NASA Headquarters

Wilma Anton, NASA Johnson Space Center

Dan Blackwood, NASA Goddard Space Flight Center

Willie Blanco, NASA Goddard Space Flight Center

Alfredo Colón, NASA Headquarters

Frank Fried, NASA Goddard Space Flight Center

Scott Graham, NASA Glenn Research Center

Christine Greenwalt, NASA Glenn Research Center

Danielle Griffin, NASA Glenn Research Center

Gene Griffith, NASA Langley Research Center

Marjorie Haskell, NASA Headquarters

Linda Hastings, NASA Glenn Research Center

Don Helton, NASA Headquarters

Steven Hirshorn, NASA Headquarters

Vicky Hwa, NASA Headquarters
Maggie Jones, NASA Headquarters
Prince Kalia, NASA Goddard Space Flight Center
Michele King, NASA Headquarters
Terry Lambert, NASA Glenn Research Center
Marion Meissner, NASA Headquarters
Eric Miller, NASA Armstrong Flight Research Center
Wendy Morgenstern, NASA Headquarters
John Orme, NASA Headquarters
Ariel Pavlick , NASA Headquarters
David Payne, NASA Space Communications and Navigation
Kelli Peterson, NASA Glenn Research Center
Robin Ripley, Goddard Space Flight Center
Jeffrey Sheehy, NASA Headquarters
Virginia Stouffer, NASA Headquarters
Madelyn Suttle, NASA Marshall Space Flight Center
Sharon Thomas, NASA Johnson Space Center
Tim Trenkle, NASA Goddard Space Flight Center
Mike Viens, NASA Goddard Space Flight Center
Zion Young, NASA Ames Research Center

Contents

ACKNOWLEDGMENTS	ii
LIST OF FIGURES.....	viii
LIST OF TABLES	x
ACRONYMS AND ABBREVIATIONS	xi
PREFACE	xiv
1 INTRODUCTION	1
1.1 PURPOSE, FOCUS, AND SCOPE	1
1.2 APPLICABLE POLICY DOCUMENTS	1
1.3 STRUCTURE AND UTILIZATION	1
1.3.1 <i>Handbook Organization</i>	1
1.3.2 <i>Use Cases and Utilization</i>	4
1.4 WHAT IS NEW IN VERSION 2.0	4
1.5 REFERENCES FOR CHAPTER 1	8
2 OBJECTIVES-DRIVEN RISK MANAGEMENT FRAMEWORK FOUNDATIONS AND KEY ELEMENTS	9
2.1 RISK AND OPPORTUNITY IN THE NASA CONTEXT	9
2.1.1 <i>Definition of Risk in the NASA Enterprise Context</i>	9
2.1.2 <i>The Objectives that Define Risk</i>	9
2.1.3 <i>Flow-down of Strategic Objectives to the Execution Level</i>	10
2.1.4 <i>Allocation and Execution of Program/Project and Institutional Objectives</i>	11
2.1.5 <i>The Performance Measures that Define Objectives Achievement and Related Risk Levels</i>	13
2.1.6 <i>Types of Risk</i>	13
2.1.7 <i>Opportunity in the NASA Context</i>	16
2.2 PRINCIPLES AND ELEMENTS OF THE NASA ODRM FRAMEWORK	17
2.2.1 <i>Policy Basis of the NASA Risk Management Framework</i>	17
2.2.2 <i>Technical and Cross-organizational Integration of Risk Perspective</i>	18
2.2.3 <i>Risk Leadership</i>	20
2.2.4 <i>Rigor in Risk Assessment and Decision Processes</i>	28
2.2.5 <i>Pros and Cons of Simplification versus Rigor in Graded Approach to Risk Assessment and Management</i>	30
2.2.6 <i>Risk Management Integration Across Technical, Organizational, and Life-cycle Boundaries</i>	33
2.3 REFERENCES FOR CHAPTER 2	42
3 RISK MODELS, ANALYSIS, AND DECISION CONCEPTS	44
3.1 FOUNDATIONAL RISK CONCEPTS	44
3.2 CHARACTERIZATION OF RISK	45
3.2.1 <i>Aggregate Risk to an Objective</i>	45
3.2.2 <i>Individual Risk Scenarios</i>	46
3.2.3 <i>Forms and Metrics for Individual Risk Scenario Contributions to Aggregate Risk</i>	53
3.2.4 <i>Unknown and/or Underappreciated Risk</i>	56
3.2.5 <i>Further Observations on Accidental vs. Adversarial Scenarios</i>	58
3.2.6 <i>The Organization-Specific Risk Model</i>	60
3.3 ORGANIZATIONAL RISK POSTURE AND TOLERANCE	61

3.3.1	<i>Definition and Application of Risk Tolerance Levels</i>	61
3.3.2	<i>Risk Classification Based on Risk Tolerance Levels</i>	65
3.3.3	<i>Use of Risk Tolerance Level Thresholds in Decisions for Risk Acceptance</i>	69
3.3.4	<i>Risk Tolerance Levels for Individual Risk Scenarios</i>	69
3.3.5	<i>Risk Classification of Individual Risk Scenarios</i>	73
3.3.6	<i>Display of Individual Risk Scenarios in Traditional Risk Matrix Format</i>	74
3.3.7	<i>Stepwise Recap and Example of Risk Leadership and Objectives-Driven Risk Management Application</i>	77
3.4	REFERENCES FOR CHAPTER 3	77
4	RISK-INFORMED DECISION MAKING	79
4.1	RANGE AND OBJECTIVES OF RIDM APPLICATION	79
4.1.1	<i>Adaptive and Graded Approach to RIDM Application</i>	80
4.2	OVERVIEW OF THE RIDM PROCESS	82
4.2.1	<i>Part 1, Identification of Alternatives</i>	84
4.2.2	<i>Part 2, Analysis of Alternatives</i>	86
4.2.3	<i>Part 3, Risk-Informed Alternative Selection</i>	87
4.2.4	<i>Performance Markers in Activity Planning vs. Activity Execution Stages</i>	89
4.3	ACCOUNTING FOR U/U RISKS	90
4.4	THE RIDM PROCESS STEPS	91
4.4.1	<i>Steps in Part 1, Identification of Alternatives</i>	92
4.4.2	<i>Steps in Part 2, Analysis of Alternatives</i>	93
4.4.3	<i>Steps in Part 3, Risk-Informed Alternative Selection</i>	94
4.5	DETAILS OF RIDM STEP 1 (PART 1), IDENTIFY OBJECTIVES AND PERFORMANCE MEASURES	96
4.5.1	<i>Constructing an Objectives Hierarchy</i>	97
4.5.2	<i>Fundamental vs. Means Objectives</i>	99
4.5.3	<i>Performance Measures</i>	100
4.5.4	<i>Risk Minimization Is Not a Performance Objective</i>	104
4.5.5	<i>Example Performance Measures</i>	104
4.6	DETAILS OF RIDM STEP 2 (PART 1), IDENTIFY DECISION ALTERNATIVES	104
4.6.1	<i>Compile an Initial Set of Alternatives</i>	104
4.6.2	<i>Identify Viable Decision Alternatives by Use of a Trade Tree or Matrix</i>	105
4.7	DETAILS OF STEP 3 (PART 2), CONDUCT INTEGRATED RISK ANALYSIS OF EACH ALTERNATIVE	106
4.7.1	<i>Set the Analytical Framework</i>	108
4.7.2	<i>Choose the Analysis Methodologies Using a Graded Approach</i>	108
4.7.3	<i>Conduct the Risk Analysis</i>	110
4.8	DETAILS OF STEP 4 (PART 2), DEVELOP THE TECHNICAL BASIS FOR DELIBERATION	120
4.9	DETAILS OF STEP 5 (PART 3), DELIBERATE	122
4.9.1	<i>Convene a Deliberation Forum</i>	123
4.9.2	<i>Develop Performance Targets and Risk Tolerances for Individual Performance Measures</i>	123
4.9.3	<i>Sequentially Establish Risk-Normalized Performance Targets for All Performance Measures</i>	126
4.9.4	<i>Pare Down the Contending Alternatives</i>	130
4.9.5	<i>Communicating the Contending Alternatives to the Decision Maker</i>	135
4.10	DETAILS OF STEP 6 (PART 3), SELECT AN ALTERNATIVE AND ACCEPT THE ASSOCIATED RISK	137
4.10.1	<i>Select a Decision Alternative</i>	137
4.10.2	<i>Finalize the Performance Targets and Assist the Decision Authority's Deliberation on Requirements</i>	137
4.10.3	<i>Accept the Risks of Selected Alternative and Document Decision Rationale</i>	138
4.11	GRADED AND SPECIAL FOCUS RIDM APPLICATIONS	138

4.11.1	<i>RIDM Specialization by Type</i>	139
4.11.2	<i>RIDM Graded Approach by Activity Class</i>	141
4.12	REFERENCES FOR CHAPTER 4	149
5	CONTINUOUS RISK MANAGEMENT	151
5.1	INITIALIZATION OF CRM	152
5.1.1	<i>Development of the Risk Management Plan</i>	152
5.1.2	<i>Inputs to CRM from the Activity-Planning RIDM Process</i>	153
5.1.3	<i>Inputs to CRM from Local Organizational Management</i>	155
5.1.4	<i>Establishing the Performance Measures To Be Considered, Performance Markers, and Associated Risk Tolerance levels</i>	155
5.1.5	<i>Risk Burn-Down Schedules</i>	157
5.1.6	<i>The Risk Database</i>	158
5.2	CRM STEP 1: IDENTIFY	159
5.2.1	<i>Identify Individual Risk Scenarios, Opportunities, and Leading Indicators</i>	159
5.2.2	<i>Develop Risk and Opportunity Statements</i>	162
5.2.3	<i>Validate the Risk and Opportunity Statements</i>	165
5.2.4	<i>Develop Risk and Opportunity Narratives</i>	166
5.2.5	<i>Categorize the Risk or Opportunity using Risk, Opportunity, and Leading Indicator Taxonomies</i>	167
5.3	CRM STEP 2: ANALYZE	167
5.3.1	<i>Implementing a Graded Approach to Analysis</i>	168
5.3.2	<i>Develop a Risk Scenario Diagram</i>	171
5.3.3	<i>Analysis of Individual Risk Scenarios</i>	172
5.3.4	<i>Analyze Opportunities</i>	188
5.4	CRM STEP 3: PLAN	192
5.4.1	<i>Generate Risk Response Options</i>	192
5.4.2	<i>Generate One or More Risk Response Alternatives</i>	196
5.4.3	<i>Perform Risk Analysis of Mitigation Alternatives</i>	197
5.4.4	<i>If Needed, Deliberate and Select a Mitigation Alternative</i>	199
5.4.5	<i>Implement the Risk Response</i>	200
5.5	CRM STEP 4: TRACK	201
5.6	CRM STEP 5: CONTROL	204
5.7	COMMUNICATE AND DOCUMENT	205
5.7.1	<i>Communication of Risk Information and Deliberations</i>	205
5.7.2	<i>Documentation of Risk Information and Deliberations</i>	207
5.8	REFERENCES FOR CHAPTER 5	208
6	ORGANIZATIONAL AND MANAGERIAL ASPECTS OF OBJECTIVES-DRIVEN RISK MANAGEMENT	209
6.1	PRINCIPLES OF ORGANIZATIONALLY INTEGRATED RISK MANAGEMENT	209
6.2	RISK MANAGEMENT INTEGRATION ACROSS LIFE-CYCLE STAGES	210
6.2.1	<i>Application of RIDM and CRM Across Life-cycle Stages</i>	210
6.2.2	<i>Evaluation of the Risk Management Effort at Life-cycle Reviews</i>	213
6.3	RECOGNITION AND HANDLING OF CROSS-CUTTING RISK	214
6.4	INTEGRATION ACROSS DOMAIN AND ORGANIZATION BOUNDARIES	218
6.5	RISK MANAGEMENT EXECUTION PLANNING	220
6.5.1	<i>Contents of the RMP</i>	221
6.5.2	<i>Cross-Organizational Integration of RM Plans</i>	226
6.6	REFERENCES FOR CHAPTER 6	227

7	DEFINITIONS	228
APPENDIX A	ROADMAP FOR RISK MANAGEMENT HANDBOOK UTILIZATION	234
APPENDIX B	EXAMPLE OF LEADING INDICATORS.....	239
APPENDIX C	FURTHER TECHNICAL CONSIDERATIONS ON RISK EVALUATION AND DECISION MAKING	242
APPENDIX D	AGGREGATE RISK EFFECTS OF CUMULATIVE-CONSEQUENCE INDIVIDUAL RISK SCENARIOS	256
APPENDIX E	STEPWISE RECAP AND EXAMPLE OF RISK LEADERSHIP AND OBJECTIVES-DRIVEN RISK MANAGEMENT APPLICATION.....	260
APPENDIX F	CONTENT GUIDE FOR THE TECHNICAL BASIS FOR DELIBERATION	274
APPENDIX G	CONTENT GUIDE FOR THE RISK-INFORMED SELECTION REPORT	276
APPENDIX H	ESTIMATE THE PERFORMANCE MEASURE MARGINS NEEDED TO ACCOMMODATE IMPLIED RISK SCENARIOS (THE U/U RISK).....	278
APPENDIX I	DEVELOP RISK, OPPORTUNITY, AND LEADING INDICATOR TAXONOMIES.....	282

LIST OF FIGURES

Figure 1-1.	Intersection of NPR 8000.4 with Program Project and Domain-Specific Directives and Requirements [2]	2
Figure 2-1.	NASA 2022 Strategic Plan Themes, Strategic Goals, and Strategic Objectives	10
Figure 2-2.	Flow-down Definition of Program/Project and Institutional Objectives	10
Figure 2-3.	Assignment and Allocation of Objectives Execution Responsibilities	12
Figure 2-4.	Flow of Risk Leadership Application within the Organizational Hierarchy	22
Figure 2-5.	Definition of Risk Posture and Tolerance Levels According to Risk Leadership Principles	22
Figure 2-6.	Risk Leadership and Management Integrated Implementation and Application Cycle	25
Figure 2-7.	High-Level View of the Interfaces of RIDM and CRM with Organizational Management.....	34
Figure 2-8.	Coordinated Use of RIDM and CRM in the Risk Management of an Activity	36
Figure 2-9.	Top-Level View of the Flow of Risk-Relevant Information to and From a Root Entity	42
Figure 3-1.	Top-Level Anatomy of Risk.....	46
Figure 3-2.	Risk Operationally Characterized as a Set of Risk Triplets.....	47
Figure 3-3.	Schematic of a Risk Scenario Diagram (RSD).....	50
Figure 3-4.	RSD Representation of “Standard” vs. “Hostile Agent” Risk Scenarios	51
Figure 3-5.	Anatomy of Risk: Known Risk and U/U Risk	57
Figure 3-6.	Robotic Spaceflight Project Outcome as a Function of Cost and Design Complexity	59
Figure 3-7.	The Risk Model (of Objective B).....	60
Figure 3-8.	Aggregation of Risk within the Integrated Organizational Risk Model	62
Figure 3-9.	Satisfaction of Risk Tolerance Levels Relative to a Performance Measure PDF	64
Figure 3-10.	Satisfaction of Risk Tolerance Levels Relative to a Performance Measure CDF	64
Figure 3-11.	The Relationship between PDFs and CDFs	65
Figure 3-12.	Risk-Bar Classification for Case of Increasing PM Direction of Goodness.....	67
Figure 3-13.	Satisfaction of Risk Tolerance Levels Relative to a Performance Measure CCDF	68
Figure 3-14.	Risk-Bar Classification for Case of Decreasing PM Direction of Goodness	69
Figure 3-15.	CCDF for an Individual Risk Scenario in Continuous and Approximate Discrete Form.....	72
Figure 3-16.	Individual Risk Scenario CCDF Relative to the IRTLs	74
Figure 3-17.	Risk Matrix for Case of Two Performance Markers	75
Figure 3-18.	Risk Matrix for Case of a Single Performance Marker	77
Figure 4-1.	The RIDM Process	83
Figure 4-2.	Functional Roles and Information Flow in RIDM Deliberations	83
Figure 4-3.	Uncertainty of Performance Outcomes Due to Uncertainty of Determining Conditions across the Enterprise, Program/Project, and Institutional Activity Domains	87
Figure 4-4.	Risk Posture Expressed by Risk-Informed Performance Markers	89
Figure 4-5.	RIDM Process Steps.....	92
Figure 4-6.	RIDM Process Flowchart: Part 1, Identification of Alternatives.....	93
Figure 4-7.	RIDM Process Part 2, Risk Analysis of Alternatives	94
Figure 4-8.	RIDM Process Part 3, Risk-Informed Alternative Selection.....	95
Figure 4-9.	Notional Objectives Hierarchy	98
Figure 4-10.	Fundamental vs. Means Objectives [8]	100
Figure 4-11.	Types of Performance Measures	102
Figure 4-12.	The Relationship between Performance Objectives and Performance Measures.....	103
Figure 4-13.	Example Launch Vehicle Trade Tree from ESAS	107
Figure 4-14.	Risk Analysis Framework (Alternative Specific).....	109
Figure 4-15.	Risk Analysis Using a Monte Carlo Sampling Procedure.....	111
Figure 4-16.	Uncertain Performance Parameters Leading to Performance Measure Histograms	114
Figure 4-17.	Application of Risk Margin to the Known Risk to Account for U/U Risk.....	115
Figure 4-18.	Robustness and Uncertainty	116
Figure 4-19.	Notional Depiction of Decision Sensitivity to Input Parameters.....	118

Figure 4-20.	Analysis Level Matrix	120
Figure 4-21.	Notional Imposed Constraints Risk Matrix	121
Figure 4-22.	Notional Band Aid Chart for Performance Measure X	121
Figure 4-23.	Comparison of Uncertainty Distributions.....	122
Figure 4-24.	Evaluation of Risk of Alternatives with Respect to a Risk Tolerance Level for Known Risk	124
Figure 4-25.	Consideration of Performance Constraints in Risk-Normalized Evaluation of Alternatives.....	125
Figure 4-26.	Establishing Risk-Normalized Performance Targets.....	127
Figure 4-27.	RPTs and Risk Tolerances for Three Alternatives	131
Figure 4-28.	An Example Uncertainty Consideration: The Potential for High Performance.....	133
Figure 4-29.	Notional Performance Target Chart	136
Figure 4-30.	Notional Risk List for Alternative X	136
Figure 5-1.	The CRM Process.....	151
Figure 5-2.	Decreasing Risk over Time for an Activity that is “On Track” to Being Within the Established Risk Posture	157
Figure 5-3.	Risk Burn-Down Schedule for a Hypothetical Organizational Objective	158
Figure 5-4.	Analysis Methodology Guidance Chart	169
Figure 5-5.	Schematic of a Risk Scenario Diagram (RSD).....	172
Figure 5-6.	Analysis of an Individual Risk Scenario, Neglecting Epistemic Uncertainty	175
Figure 5-7.	Quantitative and Qualitative Representations of Epistemic Uncertainty	177
Figure 5-8.	Analysis of an Individual Risk Scenario, Including Epistemic Uncertainty	178
Figure 5-9.	Confidence that the Risk of Not Meeting a PM X Performance Marker Due to Individual Risk Scenario q is Within the Individual Risk Scenario RTL	179
Figure 5-10.	Analysis of Aggregate Risk, Including Epistemic Uncertainty	181
Figure 5-11.	Aggregate Risk Spider Chart for an Activity	184
Figure 5-12.	Top Risks Spider Chart for an Objective.....	184
Figure 5-13.	Aggregate Risk Spider Chart of the Effects of Seizing Opportunity Y	185
Figure 5-14.	Heuristic Approach to Risk Analysis	189
Figure 5-15.	Example Illustration of the Use of an Opportunity Matrix.....	191
Figure 5-16.	The “Mitigate” Risk Response Disposition.....	194
Figure 5-17.	Relationship between Risk Response Options and Risk Response Alternatives	197
Figure 5-18.	Invocation of RIDM within the CRM Plan step	199
Figure 5-19.	Requirement Risk Tracking Chart.....	203
Figure 6-1.	Life-cycle Integration of RIDM and CRM Processes	212
Figure 6-2.	Recognition and Classification of Cross-Cutting Risks	216
Figure 6-3.	Schematic of desirable coordination of RM plans.....	227

LIST OF TABLES

Table 1-I.	Risk Management Roles and Functions by Activity Stage.....	5
Table 2-I.	Programmatic Requirements Hierarchy	14
Table 2-II.	Risk Management Approach Advantage & Disadvantage Factors.....	31
Table 2-III.	Notional LCR Success Criteria Topics.....	39
Table 3-I.	Classification of Risk Based on Satisfaction of Tolerance Levels	66
Table 3-II.	Classification of Individual Risk Scenarios.....	73
Table 4-I.	A Constructed Scale for Stakeholder Support (Adapted from [3]).....	101
Table 4-II.	Performance Measures Examples for Planetary Spacecraft and Launch Vehicles.....	105
Table 4-III.	Key Aspects of Modeling and Simulation Credibility Assessment Levels	119
Table 4-IV.	Specialization of RIDM Steps by Activity Type	142
Table 4-V.	RIDM Graded Approach by Activity Class	145
Table 5-I.	Example Net Benefit Ranking for Tactical Opportunities.....	191
Table 5-II.	Example Likelihood Ranking for Tactical Opportunities	192

ACRONYMS AND ABBREVIATIONS

A/P	Activity/Project
A/P-O	Activity/Project Objective
AC	Added Cost
ACS	Attitude Control System
AHP	Analytic Hierarchy Process
AI	Artificial Intelligence
AoA	Analysis of Alternatives
AR	Aggregate Risk
ARMO	Agency Risk Management Officer
ARTL	Aggregate Risk Tolerance Level
ASAP	Aerospace Safety Advisory Panel
ASARP	As Safe as Reasonably Practicable
ASO	Administrator Staff Offices
BC	Baseline Cost
CAS	Creditability Assessment Scale
CCDF	Complementary Cumulative Distribution Function
CCIRS	Cumulative-Consequence Individual Risk Scenario
CD	Center Director
CDF	Cumulative Distribution Function
CDR	Critical Design Review
CFD	Computational Fluid Dynamics
ConOps	Concept of Operations
CORCP	Cross-Organizational Risk Communication Protocols
CORHP	Cross-Organizational Risk Handling Protocols
CPMO	Chief Program Management Officer
CRAFT	Comparative Risk Assessment Framework and Tools
CRM	Continuous Risk Management
CRO	Chief Risk Officer
DE	Departure Event
DRM	Design Reference Mission
EEO	Equal Employment Opportunity
EOM	End of Mission
ERM	Enterprise Risk Management
ESAS	Exploration Systems Architecture Study
ESD	Event Sequence Diagram
FMEA	Failure Modes and Effects Analysis
FTA	Fault Tree Analysis
FTE	Full-Time Equivalents
GAO	Government Accountability Office
HAZOP	Hazard and Operability Studies
HCIRS	High-Consequence Individual Risk Scenario
HQ	Headquarters
ICR	Individual Cost Risk
ICRP	Independent Comprehensive Review Panel

IRS	Individual Risk Scenario
IRTL	Individual Risk Scenario Tolerance Level
IT	Information Technology
ITAR	International Traffic in Arms Regulations
JCL	Joint Confidence Level
KAR	Known Aggregate Risk
KACR	Known Aggregate Cost Risk
KDP	Key Decision Point
KPM	Key Performance Measure
LCR	Life-cycle Review
LEO	Low Earth Orbit
LOC	Loss of Crew
LOM	Loss of Mission
M&S	Modeling and Simulation
MAUT	Multi-Attribute Utility Theory
MCR	Mission Concept Review
MD	Mission Directorate
MDAAs	Mission Directorate Associated Administrators
MDR	Mission Definition Review
ML	Machine Learning
MSEO	Mission Support Enterprise Offices
MSO	Mission Support Office
MTO	Mass-to-Orbit
N/A	Not Applicable
NAPA	National Academy of Public Administration
NASA	National Aeronautics and Space Administration
NIST	National Institute of Standards and Technology
NPD	NASA Policy Directive
NPR	NASA Procedural Requirements
NPS	NASA Policy Statement
ODRM	Objectives-Driven Risk Management
ODRMF	Objectives-Driven Risk Management Framework
ORR	Operational Readiness Review
OSMA	Office of Safety and Mission Assurance
PC	Performance Constraint
PC	Project Cost
pdf	Probability Density Function
PDR	Preliminary Design Review
PIR	Program Implementation Review
P(LOC)	Probability of Loss of Crew
P(LOM)	Probability of Loss of Mission
PM	Performance Measure
PMG-M	Performance Goal Margin
PMK	Performance Marker
PMK-G	Performance Goal
PMK-R	Performance Requirement

PMR-M	Performance Requirement Margin
PMX	Performance Measure X
POC	Point-of-Contact
PT	Performance Target
R&D	Research and Development
RIDM	Risk-Informed Decision Making
RISR	Risk-Informed Selection Report
RL	Risk Leadership
RM	Risk Management
RMF	Risk Management Framework
RMP	Risk Management Plan
RMTT	Risk Management Tiger Team
ROM	Rough Order-of-Magnitude
RPT	Risk-Normalized Performance Target
RRW	Risk Reduction Worth
RSD	Risk Scenario Diagram
RTL	Risk Tolerance Level
S&MA	Safety and Mission Assurance
SDR	System Definition Review
SEMP	Systems Engineering Management Plan
SF	System Failure
SIR	System Integration Review
SMA	Safety and Mission Assurance
SME	Subject Matter Expert
SP	Special Publication
SRB	Standing Review Board
SRR	System Requirements Review
SV	Shortfall Value
TAR	Total Aggregate Risk
Tas	Technical Authorities
TBD	To Be Determined
TBfD	Technical Basis for Deliberation
TCR	Total Cost Risk
TRL	Technology Readiness Level
U/U	Unknown and/or Underappreciated
UIR	U/U Implied Risk
UUR	U/U Risk

PREFACE

Since the initial publication of the NASA Risk Management Handbook Version 1.0 (NASA/SP-2011-3422) in 2011, risk management as a discipline at the National Aeronautics and Space Administration (NASA) has continued to evolve toward a vision of a fully integrated, holistic discipline that effectively manages all sources of internal and external risk that threaten Agency objectives. This vision is codified in NPR 8000.4C, Agency Risk Management Procedural Requirements, and elaborated on in the two volumes of this revision (Version 2.0).

In 2008, NPR 8000.4A introduced Risk-Informed Decision Making (RIDM) as a complement to the Continuous Risk Management (CRM) process to ensure that direction-setting decisions were informed by an understanding of the risks associated with each alternative. Although NPR 8000.4A applied to all Agency activities, the focus of the original handbook was on programs and projects, with a promise that future versions would expand to include enterprise and institutional domains of risk management. We hope you find this revision responsive to that promise.

Under the recently created Agency Risk Management Officer (ARMO), NASA's risk management framework aligns with the flow-down of NASA's strategic goals into the activities of programs, projects, institutional organizations, and international and commercial *Providers*; addresses risks to all categories of objectives (e.g., safety and mission success, cybersecurity, physical security, technical performance, cost, schedule, reputational, compliance); supports the judicious pursuit of meaningful opportunities; identifies and communicates systemic, cross-cutting risks that affect two or more activities and/or organizations; and enables identification and communication of top risks to NASA's governing boards and forums.

An organizing principle of the framework is the philosophy of 'Risk Leadership' articulated in the NASA Governance and Strategic Management Handbook (NPD 1000.0C). Risk leadership entails decision-making, risk-acceptance, and accountability of Management Authorities within a risk posture that establishes the limits of acceptable risk to top-level objectives. This includes the authority to allocate portions of the risk posture to subordinate and/or supporting organizational entities in tandem with the flow down of objectives and requirements to those entities.

Understanding an activity's top-level objectives and focusing on adherence to the risk posture associated with those objectives are the essence of 'objectives-driven' risk management. It gives NASA organizations and partners the flexibility to engage in innovative and transformational solutions so long as the risks are understood, documented, communicated, and consistent with the established risk posture. Such flexibility is becoming increasingly important as NASA expands acquisition strategies to include commercial space systems and services, builds ever more complex in-space networks and infrastructure such as Gateway, develops or adopts new technologies such as additive manufacturing and in-situ resource utilization, and transitions systems engineering practices from document-centric to model-based.

This NASA Risk Management Handbook (Version 2.0) is presented in two volumes. Part 1 focuses on fundamental risk management concepts and principles, including the overall structure of NASA's integrated risk management framework and the risk management and decision processes that are to be implemented within the framework. Part 1 also focuses on the risk assessment

techniques that should be utilized in support of such processes and the management and organizational interfaces that should enable effective integration of risk management activities within the Agency. Part 2 includes examples of core risk management activities that practitioners may exercise when performing their risk management responsibilities.

This handbook is a living document within a dynamic NASA risk management community of practice whose mission embodies world-class excellence in the development, adoption, execution, and continuous improvement of risk management concepts, policies, procedures, and practices. The discipline of risk management will continue to evolve in concert with the Agency's governance, acquisition, and environments.

Finally, this handbook provides guidance for risk management practices and is not intended to be a directive. The principles contained herein should be adapted to the situation at hand.

Homayoon Dezfuli, Ph.D.

Project Manager, NASA Headquarters

Mary Coan Skow, Ph.D.

Agency Risk Manager Officer, NASA Headquarters

November 2024

1 Introduction

1.1 Purpose, Focus, and Scope

The purpose of this handbook is to provide an in-depth reference for the practice of risk management in NASA, updating the guidance offered in its original version, NASA/SP-2011-3422 (November 2011) [1], and closely aligning the updated guidance with the current NASA Procedural Requirements for Agency Risk Management, NPR 8000.4 [2], and the parent NASA Policy Directive for NASA Governance and Strategic Management, NPD 1000.0 (January 2020) [3].

NPD 1000.0 introduces with emphasis the concept of “Risk Leadership,” making it a fundamental tenet and pillar of the risk management culture that it advocates for the Agency. NPR 8000.4 applies this concept and establishes Risk Management (RM) requirements for the Agency as an integrated enterprise, as well as the RM requirements for portfolio elements within the enterprise. Such elements include the various programs and projects that contribute to the Agency’s objectives and the various institutional activities carried out by entities that contribute to mission support.

The present version of the handbook also emphasizes the integration of risk management processes across activity and project life cycles and their coordination and interaction with day-to-day programmatic and organizational functions. Areas of application of risk assessment and management that were not covered with specific guidance in the preceding version are addressed in this version with in-depth examples.

1.2 Applicable Policy Documents

The requirements of NPR 8000.4 that this handbook supplements with implementation guidance are aligned with NASA policy expressed by, besides NPD 1000.0, the directives of NPD 1000.3 [4], NPD 1000.5 [5], NPD 7120.4 [6], NPD 8700.1 [7], NPD 8900.5 [8], NPD 2810.1 [9], and other Agency domain-specific policy documents, as shown in Figure 1-1.

1.3 Structure and Utilization

The handbook is structured into two parts, whose chapters are in turn organized in a sequential order intended to facilitate a gradual and progressive introduction of the reader to risk management principles and practices, as further explained in the following introductory sections.

1.3.1 [Handbook Organization](#)

Part 1 of the handbook is dedicated to the introduction of the basic foundations of the NASA integrated risk management framework, the related fundamental risk concepts, the description of the risk management and decision processes that are to be implemented within the framework, the discussion of the risk assessment techniques that should be utilized in support of such processes, and the management and organizational interactions and interfaces that should be enabled to implement an effective integration of risk management activities within the Agency.

Part 2 provides examples of application of the processes and techniques introduced in Part 1, in the context of both programmatic (i.e., project and/or mission related) and institutional activities.




Figure 1-1. *Intersection of NPR 8000.4 with Program Project and Domain-Specific Directives and Requirements [2]*

1.3.1.1 Utilization of Parts 1 and 2

The handbook Part 1 presents a progression of material that is intended to reflect the order of priority by which a team or individual responsible for organizing and executing risk management processes and tasks within a NASA organization carrying out a programmatic or institutional activity may generally want to approach and become familiar with the guidance contained therein.

Although such a hypothetical individual or team will usually be engaged in risk management execution within a relatively self-contained type of activity or project, it is important that the associated execution objectives be viewed in the broader context of Agency strategic and organizational objectives, and that the potential relations and interactions of risk across the horizontal and hierarchical borders of the NASA organizational structure be fully considered. Accordingly, after the introductory information presented in Chapter 1, the first and basic elements of risk management instruction and guidance introduced in Chapter 2 are focused on the presentation and explanation of the foundational principles and elements of the current NASA risk

management framework, as set up and translated into operational requirements by NPR 8000.4. Prominent among these are:

- The application of risk leadership,
- The consideration of all risks and their aggregate effect in relation to activity objectives,
- The integration and planning of risk management implementation across activity life-cycle stages,
- The integration of risk communication and risk-related decision-making as a fundamental aspect of standard program/project and activity management.

With those principles properly identified and established as overarching risk management goals, the reader is then led to examine and consider in Chapter 3 the risk concepts that constitute the technical background and building blocks necessary to support, with the necessary rigor, effective and efficient execution of the risk management processes introduced in the following chapters.

Chapter 4 and 5 are respectively dedicated to the introduction of the Risk-Informed Decision Making (RIDM) and Continuous Risk Management (CRM) processes and to their explanation, step by step in terms of their implementation and execution, with back reference to the risk assessment and risk evaluation techniques and provisions previously discussed in Chapter 3 and the related appendices. The two chapters also discuss the types of interfacing and coordinated integration of RIDM and CRM that may occur to pursue specific risk management objectives that are relevant under specific activity/project conditions, e.g., selection and optimization of risk control strategies of special importance and relevance, or re-baselining of an entire activity or project, or major portion thereof, in terms of requirements and/or associated risk posture.

Chapters 2 through 5 seek to cover the risk management foundational principles, fundamental risk concepts, basic assessment techniques, and operational processes. With the benefit of this information and guidance the reader can proceed to the selection of the techniques and processes best suited for specific forms of risk management execution at different stages of an activity life cycle. To complement and complete the risk management organization and tailoring guidance process Chapter 6 provides the reader with a discussion of the organizational aspects of risk management execution. This covers the range from the necessary formal planning and tailoring of processes and techniques and its documentation in a Risk Management Plan (RMP), to the various forms of recommended risk-related interactions and communication, across and up-and-down the organizational and management structure of the Agency, which may also need to be formally recognized and documented in the RMPs defined by the organizations interfacing at various levels of the structure.

The organization and utilization of Part 2 reflect the same sequential order of the subjects presented in Part 1. Thus, after the Chapter 1 introduction, Chapter 2 discusses examples of flow-down and identification of objectives that provide the focus towards which the risk management for an activity or project are to remain oriented, as well as examples of selection of risk management processes and tasks that suit the context and stage of a specific type of activity or project. In their ensemble, these materials illustrate and exemplify the concepts discussed in Chapters 2 and 3 of Part 1.

Chapters 3 and 4 of Part 2 provide, respectively, practical examples of the application of RIDM and CRM processes in both program/project and institutional activity contexts. Therefore Part 2 Chapter 3 provides examples for the RIDM concepts and processes discussed in Part 1 Chapter 4, whereas Part 2 Chapter 4 has a corresponding role with regard to the CRM concepts and processes discussed in Part 1 Chapter 5.

1.3.2 Use Cases and Utilization

While the description of the handbook structure given in the preceding section is also an outline of its order and mode of utilization by a generic reader, the topics discussed throughout the document undoubtedly have different degrees of relevance for readers with different roles and responsibilities in the organization and execution of risk management tasks. This section provides a general overview intended to facilitate the order of priority by which readers with different roles may approach the principal topics discussed in the handbook.

Table 1-I presents a depiction of the common types of roles and objectives that an individual or team with risk management organizational or executional responsibility may have in the two principal top-level stages of a project or activity – i.e., *Definition/Planning* vs. *Execution*. This identification of risk management roles is then used with the detailed roadmap presented in Appendix A to identify the topics covered in this handbook which may have different degrees of relevance and priority in relation to the risk management role and stage of application identified in Table 1-I.

A few clarifications are in order with respect to the use of the two tables (Table 1-I and the detailed roadmap in Appendix A). First, the roles defined in Table 1-I represent general definitions as represented in the table itself: they should not be construed as having any direct pre-defined correspondence with the personnel roles officially assigned within the NASA organizational and programmatic hierarchies. Second, the terms “Definition/Planning” and “Execution” are used to identify the two major activity stages that are relevant in relation to the type of risk management processes that are to be executed within them: in the context of a formally structured project life cycle and the corresponding systems engineering definition of project phases, Definition/Planning corresponds to the Pre-Phase A and Phase A portions of the project life cycle, and Execution corresponds to the remainder of the life cycle.

1.4 What is New in Version 2.0

A majority of the general principles and fundamental concepts described in the original version of this handbook [1] remain valid in the current NASA application contexts. However, new principles have been articulated over the last 10 years, from activities such as the recent Agency-level Risk Management Tiger Team assessment [10], which among other things emphasized the importance of risk management in procurement and acquisition decision making [11]. Such principles include:

- Emphasis on Objectives-Driven Risk Management (ODRM) and Risk Leadership (RL).

This theme is centered around the complementary concepts of risk as the potential for shortfalls with respect to the achievement of established objectives (aka ODRM), and of the need for the establishment, by overseeing decision-authorities and in consultation with affected stakeholders, of limits of acceptable risk to an organization’s or activity’s top-

level objectives (aka *risk tolerances*), within which activity managers have risk management authority and accountability (aka RL). The risk tolerances on an organization's or activity's top-level objectives are collectively referred to as its *risk posture*. ODRM and RL focus risk management on the objectives that define success for an organization or activity; establish a risk posture for those objectives; and give activity managers leadership authority to manage the risk as they see fit, as long as it remains within the risk posture. Moreover, ODRM and RL flow down limits of acceptable risk to supporting organizations in tandem with the flowdown of derived and allocated objectives, thereby establishing those organizations' top-level objectives and associated risk postures within which they, too, exercise risk leadership. ODRM and RM principles can be found in Chapters 2 and 3, as well as in the application examples presented in Part 2.

Table 1-I. Risk Management Roles and Functions by Activity Stage

Project / Activity Stage ---->	Definition and Planning	Execution
USER TYPE	RM ROLE & FUNCTIONS	
Deliberation & Decision Board / Council	<i>Providing High Level Directions and Making Decisions</i>	<i>Providing High Level Directions and Making Decisions</i>
High-Level Leader / Manager	<i>Providing Top-Level RM Direction / Focus</i>	<i>Providing Top-Level RM Direction / Focus</i>
Review & Recommendation Board / Council	<i>Providing Oversight</i>	<i>Reviewing and Providing Oversight</i>
Program / Project / Activity Manager	<i>Organizing and Directing RM Tasks</i>	<i>Organizing / Directing RM Tasks</i>
	<i>Making and Accounting for Risk-Relevant Decisions</i>	<i>Making and Accounting for Risk-Relevant Decisions</i>
Program / Project / Activity RM Specialist	<i>Executing RM Processes and Tasks</i>	<i>Executing RM Processes and Tasks</i>
Program / Project / Activity Risk Analyst	<i>Executing Specific Risk Analysis Tasks</i>	<i>Executing Specific Risk Analysis Tasks</i>

- New Focus on the Assessment, Communication, and Management of Aggregate Risk.

The definition of risk in NPR 8000.4 places the focus of risk assessment, communication and management processes on the *total* risk to an organization's or activity's objectives. Such risk is referred to in the NPR and in this handbook as *aggregate risk*, to indicate that

it includes the full ensemble of *individual risk scenarios* (which in the general RM literature are commonly referred to as “individual risks”) that can result in shortfalls. The focus on aggregate risk stands in contrast to the exclusive focus on individual risk scenarios that has historically characterized NASA risk management. Aggregate risk is introduced in Chapter 2, with further detailed discussion, including the different types of “risk aggregation” that may be encountered, in Chapter 3.

- Rigorous Integration of RIDM and CRM Processes throughout Activity Life Cycles.

The present handbook emphasis and discusses how RIDM and CRM can be fully integrated, via the application of the necessary procedural and analytical rigor, throughout an activity/project life cycle. This theme is introduced in Chapter 2 and examined in relation to its various contexts of operational application in Chapters 4 and 5, and in the corresponding examples provided in Part 2.

- RM Integration across Agency Organizational Structures.

While it remains a central principle of effective risk management that its execution is to be conducted by organizational units within their sphere of responsibility and capability, this update dedicates considerable attention to its cross-organizational interface aspects. The backbone of these interfaces is the flow-down of objectives from the NASA Strategic Goals through the NASA organizational structure and into programs, projects, and institutions. Acceptable levels of risk flow down in tandem with the objectives, and the actual risks that each organizational unit faces are communicated back up along the same pathways to the higher-level organizations that they serve. Additionally, RM integration includes those circumstances that call for horizontal communication of risk information and for the elevation of risk management decision-making. These issues are addressed in Chapter 6.

Besides the principles identified above, the topics identified below are also new important additions to the scope of the handbook:

- Alternative Means of Risk Communication and Display.

Visual displays that constitute alternatives to the use of traditional 5 x 5 risk matrices are discussed in Chapters 3 through 5, contextually to the nature of the underlying risk information to be communicated. Two underlying and fundamental principles are established as premises of these discussions. The first is that the rating of individual risk scenarios cannot any longer be based on across-the-board definitions, as routinely done in the past by pinning them to standardized 5 x 5 risk matrices, but is to be determined and communicated in relation to objective-specific risk tolerances that are explicitly tied to the degree to which they threaten an organization’s or activity’s top-level objectives. Therefore, there are as many sets of risk rating criteria as there are top-level objectives, and each set is anchored to the limits of acceptable risk to the associated objective. The second principle is that a clear distinction is to be made and maintained between the analysis of risk, whether aggregate or individual, and the acceptability of that risk per the risk rating criteria. The former involves the objective determination of the likelihoods of shortfalls with respect to the top-level objectives, whereas the latter involves the subjective

assessment of the tolerability of those likelihoods. In general, risk analysis must be completed prior to, and separately from, risk rating. As discussed in Chapter 3 and following chapters, once individual risk scenarios and their aggregate risk effects on activity/project objectives have been analyzed and rated, the communication of their risk acceptability or non-acceptability may be summarized in simple stoplight form, or in an objective-specific 3 x 3 or 5 x 5 form. Whatever the form of communication and display selected, it needs to be anchored to the risk posture established for the organization or activity of concern, with respect to its specific objectives.

- Risk Management Execution in the Enterprise and Institutional Domains.

While the original version of the RM Handbook concentrated on the relevance of risk management application in the Program/Project domain, this version also covers contexts and means of risk management application in the Enterprise and Institutional domains that are relevant for NASA. These subjects are addressed in Chapters 4 and 5, and detailed application examples are provided in Part 2.

- Risk, Benefit, and Opportunity Management.

Balancing risk and benefits, and exploiting opportunities within an organization's deliberately adopted risk posture is treated in this handbook as being an integral part of risk management. As an informing principle this is primarily discussed in Chapter 2 and expanded on in the context of CRM in Chapter 5.

- Consideration of Physical Security and Cybersecurity Risks.

- Physical Security protects NASA's physical assets and resources (facilities, people, equipment, etc.) from environmental threats and physical harm. This includes both active measures like security personnel and passive measures like barriers and surveillance. NPR 1620.3B Physical Security Requirements for NASA Facilities and Property, responsible office: OPS provides more details.
- Cybersecurity protects NASA's information and data in the various states that it exists (during processing, in transit and in storage) from the risk of inaccessibility, corruption, or exposure. NPR 2810.1F Security of Information and Information Systems, responsible office: OCIO provides more details.
- Both have become increasingly central areas of concern for organizations, assets, and missions, especially in the last decade. Risks within these domains often involve intentional attacks, which require special treatment in terms of modeling and analysis. While the themes of cybersecurity and physical security risk are addressed throughout the handbook, the nature of the tailoring of risk assessment and management provisions necessary to account for the special nature of intentionally driven scenarios is discussed primarily in dedicated sections of Chapter 3.

- Graded Approach in RIDM and CRM Application Processes.

This present handbook version provides a much more detailed accounting of how a graded approach may be used in defining the scope, completeness, and mix of quantitative vs.

qualitative considerations in the RIDM and CRM processes and analyses used in the execution of risk management activities. For a general introduction to this important subject the reader is urged to first review the discussion given in Chapter 2, Section 2.2.5, where the subject is addressed from the perspective of the pros and cons of adopting simplified vs. more rigorous processes of risk assessment and management. The concept of graded approach is then discussed in greater detail in Chapters 4 and 5, with specific reference to the aspects of its implementation within the RIDM and CRM processes. Examples of application of the approach are provided in Part 2.

- Consideration of Unknown and/or Underappreciated Risk and Associated Leading Indicators.

Besides the familiar theme of managing known risks, the present handbook also accounts for the potential magnitude of unknown and/or underappreciated (U/U) risks, using historical experience as a guide. This subject is primarily discussed in Chapter 3.

1.5 References for Chapter 1

1. NASA Special Publication, NASA/SP-2011-3422, NASA Risk Management Handbook, Version 1.0. November 2011.
2. NASA Procedural Requirements, NPR 8000.4C, Agency Risk Management Procedural Requirements. April 2022.
3. NASA Policy Directive, NPD 1000.0C, NASA Governance and Strategic Management Handbook. January 2020.
4. NASA Policy Directive, NPD 1000.3E, The NASA Organization w/Changes 106-108. April 2015.
5. NASA Policy Directive, NPD 1000.5C, Policy for NASA Acquisition - Updated w/Change 2. July 2020.
6. NASA Policy Directive, NPD 7120.4E, NASA Engineering and Program/Project Management Policy. June 2017.
7. NASA Policy Directive, NPD 8700.1F, NASA Policy for Safety and Mission Success. July 2022.
8. NASA Policy Directive, NPD 8900.5B, NASA Health and Medical Policy for Human Space Exploration (Revalidated 3/28/17). December 2011.
9. NASA Policy Directive, NPD 2810.1F, Security of Information and Information Systems. January 2022.
10. NASA Internal Report, Risk Management Tiger Team Report. September 07, 2023.
11. NASA Policy Statement, NPS 1001.105, NASA Chief Acquisition Officer's Intent, June 18, 2024.

2 Objectives-Driven Risk Management Framework Foundations and Key Elements

This chapter is dedicated to an examination of the foundational principles and perspectives on risk that define and inform the entire NASA Objectives-Driven Risk Management framework. The understanding and application of these principles at all levels of the NASA organization and activities is the key for a successful and fully integrated implementation of the framework and of the processes that within it are applied in the wide range of all its application contexts.

2.1 Risk and Opportunity in the NASA Context

This section considers the definitions, and the related concepts, that provide the fundamental perspective from which risk, and its specular opposite, opportunity, should be viewed and considered from within the NASA enterprise and organization. When considering and utilizing the more detailed elements of guidance and information contained in the remainder of this handbook, users are invited to do so without ever losing sight of the connection between those practical and operational guidelines and the informing principles and perspectives discussed in the present chapter.

2.1.1 Definition of Risk in the NASA Enterprise Context

NPR 8000.4, *Agency Risk Management Procedural Requirements* [1], defines risk as “the potential for shortfalls with respect to achieving explicitly established and stated objectives.” This definition, which is also the definition assumed for the entire set of contexts addressed by this handbook, is tailored so as to focus attention on the “explicitly established and stated objectives” that are being pursued by the organizational entity in question, be it a mission, a project, a program, a NASA Center, a mission support organization, or the Agency as a whole. These are the objectives that define success for the entity: if the entity meets its objectives, then it is successful; if it fails to meet its objectives then it is not successful, or at least not fully successful. Defining risk in terms of the objectives that define an entity’s success ensures that entity’s risk management activities, like all the entity’s activities in general, remain focused on success. Section 2.1.2 below discusses how the objectives used to define risk are established.

2.1.2 The Objectives that Define Risk

As an Agency, NASA is governed by objectives, beginning with the strategic goals identified by NASA in its periodically updated Strategic Plan [2] and the strategic objectives that are derived from them, as illustrated in Figure 2-1 (reproduced from [2]). These strategic objectives are then allocated to NASA’s directorates and staff offices, and into its Centers, programs, projects, and mission support organizations. At each level of allocation, a given organizational entity’s top-level objectives toward which the entity operates define success for that entity, and the success of every entity is contingent upon the success of the subordinate entities into which it allocates sub-objectives, as well as the success of equal-level entities with which it may be cooperating for the achievement of some of its objectives. Thus, NASA is structured as a coordinated conglomerate of semi-autonomous organizational entities, each of which is connected to other entities through a network of allocated and shared objectives, as further discussed below.

Program/project, and Institutional domains), without referring to any specific assignment of the corresponding activities to specific entities in the NASA organizational structure. Also, it should be noted that the figure does not attempt to identify how the objectives that are generically identified in it, by type and domain association, may correspond to the goal and objectives identified in the current NASA Strategic Plan. A brief discussion of such a correspondence, however, is provided in Part 2 of this handbook.

The key points presented and illustrated by Figure 2-2 are as follows:

- The Enterprise Domain defines strategic and cross-agency objectives, of both “technical” and “institutional” nature, which are usually developed into execution-level sub-objectives to be pursued via projects and activities defined within the Program/Project and Institutional Domains.
- Technical objectives defined at the top Enterprise Domain level are eventually articulated into two broad categories of Program/Project Domain sub-objectives, i.e.:
 - a. a portfolio of mission objectives, and
 - b. a set of technology Research and Development (R&D) objectives intended for the development and improvement of technical capabilities needed for the execution of future operational missions.
- Institutional Domain objectives can, in their derivation and definition, be viewed as being in either of two categories:
 - a. objectives that proceed from cross-agency institutional development and support needs identified at the Enterprise Domain level;
 - b. objectives that have their source in the specific support needs of program and projects defined within the Program/Project Domain.
- As mentioned earlier, the figure does not explicitly identify the types of organizational units to which different types of execution objectives are assigned; however, its breakdown and flow-down of objectives on the Institutional Domain side does imply that both the sources and the executions of the two types of institutional objectives it identifies may reside at substantially different levels of the organization. This is further discussed below in Section 2.1.4.

2.1.4 Allocation and Execution of Program/Project and Institutional Objectives

Figure 2-3 presents a depiction, necessarily simplified yet representative in its essential traits, of the basic layers in the NASA organizational structure where technical and institutional objectives are defined, and where responsibilities are assigned for execution of corresponding activities and projects.

In alignment with the present NASA policies and directives, the figure indicates that:

- the higher levels of the NASA organizational structure hold the primary responsibility and functions for establishing agency objectives and the overall principles to be applied in their pursuit, including Risk Leadership principles;

- the NASA Directorate organizations are responsible for setting up and managing the programs, projects, and institutional activities by means of which agency objectives, decomposed and articulated into an executable portfolio of missions, projects, and institutional tasks, can be pursued;
- the NASA Centers and Facilities are ultimately assigned the responsibility and functions of providing executable projects and activities with the necessary management, technical, and support personnel, as well as the needed physical / technical infrastructure (offices, labs, assembly and test facilities, Information Technology (IT) equipment and services, etc.).




Figure 2-3. Assignment and Allocation of Objectives Execution Responsibilities

The figure should not be interpreted as a formal depiction of a specific organizational set-up, as any such specific organizational structure may be subject to adjustments or even significant changes over time. Rather, it is intended to depict in broad outline, and in the form presently recognizable along and across the principal layers of the NASA organizational structure, the allocation of the main types of responsibilities concerning the definition and execution of agency technical and institutional objectives, without implying any corresponding and rigidly defined subdivision of roles and responsibilities in all relevant practical cases. Consistently with this perspective, a double path for assignment of institutional objective execution responsibilities is shown in the figure, with one side indicating responsibility flowing to Centers and Facilities directly from the Enterprise executive level (i.e., the NASA Administrator Offices and Councils), and the other addressed to the former via elaboration and definition carried out by the Mission Support Directorate. Also to be recognized is that, although the primary role of personnel residing

at a Directorate level is for the most part that of managing and controlling, rather than executing, project and institutional tasks, there may be even at that level individuals who actually carry out relevant and significant portions of the objective-execution tasks.

As a further illustration of the flowdown of objectives and the meaning of the term, “top-level objectives,” as used in this handbook, Table 2-I, reproduced from NPR 7120.5, shows the flow down of objectives from Agency strategic planning through Agency, directorate, program, and project requirement levels to the systems that will be developed and used to achieve Agency goals. At each level of the hierarchy the objectives that flow down to an organizational unit are the top-level objectives for that unit. Thus, for example, the top-level objectives for a program are the “Detailed requirements levied on a program to implement the PCA and high-level programmatic requirements allocated from the program to its projects”, whereas the top-level objectives for a system are the “Detailed requirements allocated from the project to the next lower level of the project.”

2.1.5 The Performance Measures that Define Objectives Achievement and Related Risk Levels

Upon the development of a set of top-level objectives to be levied on an organizational unit, one or more performance measures are identified for each objective as the metric(s) by which its degree of fulfillment can be characterized and/or quantified. In most cases the appropriate performance measure(s) to use is/are self-evident from the definition of the objective itself. In other cases, the choice may not be as clear, and work must be done in order to assure that the objective is not only assessable, but that the performance measure used to characterize it or quantify it is adequately representative of the objective to begin with. Sections 4.5.3, 4.5.5, and 5.1.4 will provide targeted guidance on defining performance measures that enable the gauging of the degree to which an objective is being successfully accomplished. As will be explained and discussed in depth in Chapter 3, the full quantitative expression of the risk to an objective involves an estimation of the probability that the performance at issue, as characterized in terms of the performance measure(s) related to the objective, is deficient with respect to the target value(s) that define(s) achievement of the objective.

2.1.6 Types of Risk

The nature of an organizational entity’s risks, and therefore of the risk management processes brought to bear on them, are a function of the corresponding objectives. Risks to objectives allocated to NASA’s Agency-level organizations are *enterprise risks*; risks to objectives allocated to NASA’s infrastructure and mission support organizations are *institutional risks*, and risks to objectives allocated to NASA’s programs and projects are *program and project risks*. Given that NASA projects often involve the execution of specific space missions, *mission risks* generally may be considered as a type of project risks.

Another way of distinguishing categories of risk is according to the nature of the objectives being pursued. For example, NPR 8000.4 defines the “mission execution domains” of *safety*, *mission success (technical)*, *physical security*, *cybersecurity*, *cost*, and *schedule*. Correspondingly, the risk associated with a particular objective can be categorized in terms of the category of the objective. Thus, the potential shortfalls with respect to a safety objective constitute a *safety risk*, and so on.

Table 2-I. Programmatic Requirements Hierarchy

Requirements Level	Content	Governing Document	Approver	Originator
NASA Strategic Goals	Agency goals, objectives, and strategic direction	NPD 1001.0, NASA Strategic Plan; and Strategic Planning Guidance	NASA Administrator	OCFO
Mission Directorate Requirements	High-level requirements levied on a program to carry out strategic and architectural direction, including programmatic direction for initiating specific projects	Program Commitment Agreement (PCA)	NASA AA	MDAA
Program Requirements	Detailed requirements levied on a program to implement the PCA and high-level programmatic requirements allocated from the program to its projects	Program Plan	MDAA	Program Manager
Project Requirements	Detailed requirements levied on a project to implement the Program Plan and flow down programmatic requirements allocated from the program to the project	Project Plan	Program Manager	Project Manager
System Requirements	Detailed requirements allocated from the project to the next lower level of the project	System Requirements Documentation	Project Manager	Responsible System Lead

MDAA = Mission Directorate Associate Administrator; NASA AA = NASA Associate Administrator

A more general distinction between types of risk, one which applies in fact regardless of any other distinction, concerns whether one is considering *aggregate risk* or *individual risk scenarios* within such an aggregate. The corresponding definitions and an in-depth discussion of what in practice this distinction entails are provided in Chapter 3. As an introduction, it suffices here to be said that *aggregate risk* refers to the full ensemble of potential situations and conditions that may lead to the non-fulfillment of a specific organizational and/or activity objective, whereas an *individual risk scenario* is the specific definition of a single scenario that may compromise, or in combination with other individual risk scenarios contribute to compromising, one or more objectives.

Common Types of Risk

Organizational Risk Types

Organizational entities have been traditionally classified according to their role within the NASA organizational structure and one top-level, corresponding way of classifying risk is according to the type of organization it affects, e.g.:

- *Enterprise (Strategic, Agency) risk*
- *Institutional risk*
- *Program/Project risk*

Activity-Domain-Specific Risk Types

Enterprise-level organizations are responsible for different types of objectives that, at lower levels, are assigned to organizations that operate in program/project or institutional domains. As each of these organizations has specific types of objectives and corresponding activities for which it is responsible, risks can also be further classified in terms of the more specific types of activities that they affect. Herein and in the remainder of the handbook objectives and corresponding risks that fall under the jurisdiction of enterprise organizations are referred to as being part of the “enterprise activity domain,” whereas those under the jurisdiction of program, project, or mission organizations are referred to as being within the “program/project activity domain,” and objectives under the jurisdiction of institutional organizations as being part of the “institutional activity domain.” Common sub-categories to which risks identified within each of three above domains are traditionally and commonly associated with are listed in the table below:

Common Risk Types Typically Assigned to Each Activity Domain		
Enterprise	Program/Project	Institutional
• Strategic risk • Operations risk • Compliance risk • Acquisition risk • Fraud risk • Reputational Risk • Etc.	• Safety risk • Technical risk • Security risk • Cost risk • Schedule risk • Etc.	• Staffing risk • Training risk • Maintenance risk • Supply chain risk • Facility safety risk • Facility availability risk • Etc.

Caveats

The bucketing of risks into activity domains and sub-types should not be taken too literally. Individual risks are scenarios in which a chain of events depicts a progression from causes to consequences, and the classification attributes listed above may refer to any element of that cause-effect chain. Thus, a given risk may be “technical” in its root-causes and “cost/schedule” in its consequences. It is also common for most risk types to span two and sometimes all three of the top-level activity domains. For example, cybersecurity risks can threaten enterprise-level (Agency-wide) networks, program/project networks, or institutional (e.g., Center-wide) networks. Technical risks, cost risks, and acquisition risks can occur at any level. The cross-cutting nature of many risks over various activity domains is a topic that will be addressed later in several places in this handbook, particularly within the context of the implementation of risk management when risks affect multiple activity domains.

2.1.7 Opportunity in the NASA Context

As part of the development and maturing of the concept of Risk Leadership and the changes in organizational culture that have ensued from it, risk management presently seeks to explicitly take into consideration not only the negative potentialities traditionally associated with risk but also the positive possibilities associated with potential benefits and opportunity. As will be further discussed later in this chapter, this means that the application of risk leadership is intrinsically tied to the definition of a balanced risk posture, which informs decisions relative to risk also with the consideration of the benefits and opportunities offered by the selection of activity objectives and associated courses of action.

The concept of opportunity is contextualized for NASA application by the corresponding definition given in [3], which is emphasized in the “blue box” shown below.

Opportunity

“The possibility of an existing goal, objective, or desired outcome being met more efficaciously, or a new goal, objective, or desired outcome becoming feasible.” [3]

In the NASA context, opportunity may typically have two different aspects and manifestations.

The first possible aspect applies to an opening for modifying an organizational unit’s *fundamental objectives* – i.e., the true top-level technical and programmatic goals of an activity or ensemble of activities – to better align them with the objectives of the organizational unit from which they were allocated. From the Agency point of view, the ability to consider and adopt this type of organizational redirection is an expression of overall agility as an enterprise. For example, the emergence of new technologies can open up possibilities that justify redirecting existing efforts towards new objectives that more effectively advance strategic goals. This type of opportunity pertains to promoting accomplishment of the Agency’s mission through strategic re-planning, rather than merely by reducing the risks in its existing network of execution objectives.

The latter type of opportunity is associated with actions that have a more tactical than strategic connotation, and that as such usually have an impact in a more constrained context, by reducing the risk to specific project or activity execution objectives, and/or by inducing the generation of additional and collateral benefits from the execution of certain project or activity tasks, via the modification of task means objectives.

The term *means objectives* refers to the means by which the fundamental objectives of the organization are achieved. Such means generally concern the operational accomplishment of the programs, projects, enterprise initiatives, institutional initiatives, tasks, and activities that comprise an organization’s performance plan. For example, for an organization that has begun execution of a project, an emerging opportunity to share a research and development task with a partner organization with specialized expertise in that area might result in both a reduction of the risk of failing in that task and in the benefit of better practical results from it. The event that leads to the possibility of a partnership (e.g., the partnering organization expressing a willingness to participate) is an opportunity because it offers a lower risk path to the achievement of the first

organization's objectives and overall results potentially improved significantly above the minimum level of satisfactory performance.

Opportunities (like risk responses) typically have a timeframe associated with them, i.e., a “window of opportunity,” during which the opportunity must be either seized or forfeited. Correspondingly, effective opportunity management requires an organization to be both proactive in its identification of opportunities, which typically involves active engagement with, and awareness of, the socio-technical environment in which the organization operates, and agile in its ability to assess the potential values of identified opportunities and adjust its execution of tasks to pursue the opportunities it deems fruitful. This is one reason why an organization must be agile. Significant gains in advancement or progress may also involve investment in the creation of opportunities, e.g., by putting resources into basic or applied research, with the expectation that on the whole these efforts will bear fruit and speed the rate of progress toward long-term goals. In the words of Francis Bacon in 1612 [4]: “A wise man will make more opportunities than he finds.”

2.2 Principles and Elements of the NASA ODRM Framework

Having established the basic definitions of risk and opportunity that are operationally meaningful in the context of NASA programs, projects, and activities, it is important to also examine the foundational principles that inform the NASA ODRM framework (ODRMF) and the key components that support its implementation, in accord and alignment with those principles. Regardless of the organizational level and the corresponding importance and breadth of the involved activities, the former should remain constant points of reference, and the latter should be utilized consistently with their definitions and guidelines in any execution of risk management processes and tasks.

This chapter discusses the following themes as informing principles and foundational elements of the NASA ODRMF implementation:

- Cross-organizational and integrated perspective on risk
- Application of risk leadership
- Rigor in execution of risk decision processes and supporting analyses
- Integration of risk management execution and communication across organizational boundaries and activity life cycles

The discussion concerning the above ODRMF themes is preceded in the following by a summary review of the primary requirements and guidance, established in NASA directive and mandate documents as well as other official documentation, that have a clear direct or indirect bearing on the related risk management subjects.

2.2.1 Policy Basis of the NASA Risk Management Framework

Fundamental policy elements that support the NASA ODRMF are expressed at the higher level in the NASA directives NPD 1000.0 [5], NPD 1000.5 [6], and NPD 8700.1 [7]. These are then fully articulated and expanded into the risk management requirements of NPR 8000.4 [1]. Direct reference to the ODRMF foundational themes and principles discussed in this chapter can therefore be found, albeit at different levels of elaboration, in all of the above documents, as well as in

recently distributed policy statements, such as NPS 1001.105, which contains a strong endorsement of the underlying principles [8]. To provide the necessary policy perspective for the subjects being addressed, this will be identified and highlighted, as applicable, in the discussion that follows.

2.2.2 Technical and Cross-organizational Integration of Risk Perspective

The integration of risk assessment and management activities, both in an organizational sense throughout the Agency structure, and in a technical sense across the boundaries of risk domain definitions (e.g., technical or safety vs. cost or schedule) is a fundamental prerequisite for the identification of risks of a cross-cutting nature and the application of resources for their handling and control, at the appropriate level and with the appropriate breadth and range.

Direct and indirect references to the need for cross organizational integration of risk related processes and activities can be found in both NASA official policy documents (NPDs and NPRs) and in other organizational and policy-related documentation. NPR 8000.4 provides a definition of cross-cutting risk (see blue-box below) and identifies multiple provisions intended to address risks of such a nature and encourage the cross-organizational and vertical integration of management practices pertaining specifically to the identification and handling of risk within the Agency. Consistently with the policy stance set forth by the NPR, the recently published final report [9] of a Risk Management Tiger Team (RMTT) appointed at top agency level recognizes the importance of risk management cross-integration and of the identification and handling of cross-cutting risks. The report identifies these themes as being central to an effective application of the NASA risk management, to the point that to remedy deficiencies noted by the RMTT in the way they are addressed in the risk management execution it recommends the appointment of an Agency top-level Chief Risk Officer (CRO), having the primary function of “risk integrator” for the Agency, with responsibility for identification of “Top Enterprise Risks,” the development of cross-organizational Enterprise Risk Management (ERM) processes, communications, and methods, and the promotion of Risk Leadership initiatives. The key RMTT recommendations were adopted by the Agency and further endorsed in the Policy Statement NPS 1001.105 [8] issued by the NASA Deputy Administrator. This document refers to the newly instituted position of Agency Risk Management Officer (ARMO)¹ and to the associated functions and responsibility in the following manner:

- *“Strengthening Objectives-Driven Risk Management and Risk Leadership. The ARMO will work in concert with Mission Directorates, Centers, and Offices to increase awareness, understanding, and consideration of risks to their key objectives to support risk-informed decision-making in all phases of the acquisition life cycle. This collaboration includes the application of risk leadership principles to define and communicate the risk postures that express the amount of risk NASA is willing to accept to achieve its acquisition objectives.*
- *Strengthening Risk Oversight of our Acquisitions. Under the guidance of the Associate Administrator, the Assistant Administrator for Procurement/Deputy CAO, Chief Program*

¹ ARMO is the officer position that the RMTT had recommended instituting, using however the denomination of Chief Risk Officer (CRO)

Management Officer (CPMO), and ARMO will work in concert with the Office of the Executive Secretariat to strengthen our existing governance councils and policies. This includes setting clear expectations for acquisition strategy meetings that risk-informed data and evidence be utilized to support recommendations. Decisional outcomes will continue to be rigorously documented and tracked, aligning objectives and associated risk postures prior to execution of acquisition activity and development of the procurement strategy.”

To the end of promoting risk management integration across the Agency, the Deputy Administrator charges via the NPS *“our Associate Administrator, Mission Directorate Associate Administrators, Center Directors, and other Officials-in-Charge and senior leaders to: 1) promote a culture that fosters and values Risk Leadership; and 2) encourage objectives-driven risk management activities that enable productive discussions through open multidisciplinary discussions of concerns, risks, and issues with the potential to impact NASA’s success.”*

As mentioned above, a truly integrated perspective on risk has both technical and organizational implications with regard to how it is to be assessed and managed. The technical dimension lies in the fact that a full understanding and appreciation of risk requires not only the identification of its individual sources and effects but also a serious and rigorous consideration and assessment of its aggregated impact. NPR 8000.4 provides an operational definition of aggregate risk (see again blue box below) and a set of specific requirements for its appropriate assessment and handling.

Aggregate Risk

“The cumulative risk associated with a given goal, objective, or performance measure, accounting for all significant risk contributors. For example, the total probability of loss of mission is an aggregate risk quantified as the probability of the union of all scenarios leading to loss of mission”. [1]

Cross-cutting Risk

“A risk that is generally applicable to multiple mission execution efforts, with attributes and impacts found in multiple levels of the organization or in multiple organizations within the same level.”

In the organizational sense, an integrated perspective on risk requires the placement of risk management practices within the broader integrated structuring and organization of program, project, and institutional activity management. In its definition of Program Management Integration, NPD 1000.0 identifies the formulation of Risk Management Plans (RMPs) by programs and projects as one of its key ingredients. In NPR 8000.4 the importance of cross-organizational integration of risk assessment and management activities is underlined by the following statement:

“1.2.1.3 Managing risk requires the coordination of risk identification, assessment, decision and communication activities. All levels of NASA executives and managers, at all levels of the organizational hierarchy, are responsible to enable such a coordination.”

NPR 8000.4 also addresses other specific aspects of risk management integration and coordination requirements, such as those concerning the relationship between the procuring and executing sides – i.e., the “Acquirer” and the “Provider” – of a project or activity, and the coordination of the

execution of risk management processes across organizational boundaries, when this is necessary for an effective management of cross-cutting risks. The details of these subjects are discussed within the related contexts throughout this handbook.

2.2.3 Risk Leadership

Risk leadership can be defined as being the application, by the leaders, managers, and execution staff of an activity or project, of a clear and consistently shared identification and communication of the activity priority objectives, and of the associated risk posture to be applied in the pursuit of such objectives and of the associated targeted benefits.

Risk leadership is therefore expressed and enabled via the identification and communication by higher-level organizational leaders of the strategic and institutional objectives of the concerned entity or institution, and of accompanying guidance on how to apply an appropriate posture in risk related decision making, and balance between rigor and speed. A complementary and enabling element of Risk Leadership is the feedback of appropriately prioritized risk information and Risk Management results from the execution levels to the higher managers and leaders of the organization and enterprise.

NPD 1000.0 dedicates an entire section to the introduction and discussion of the principle of risk leadership, and NPR 8000.4 makes multiple references to it in relation to its interrelations with other key risk management principles, and especially in regard to its role for promoting coordination and integration of risk posture and risk related decision processes across the management structures and hierarchies of projects and activities.

When viewed and interpreted from the above perspective, the application of risk management emerges from and is informed by, a specific Risk Leadership context established by the Agency's organizational and institutional managers and shared with the entire workforce that is charged with the execution of risk management and related activities. A definition of this context provides criteria through which risk-informed judgments and decisions can be formulated in modes compatible and consistent with the top programmatic and institutional objectives of the NASA enterprise and organization. In complementary fashion within the same context, personnel with risk management execution responsibilities are also responsible for applying criteria and modes of risk communication that are appropriate for providing priority feedback information on risk activities and results in a timely fashion, both upward and across the organizational lines of communication. This makes it so that higher management decisions are appropriately informed, and that risk leadership application guidance can be adjusted or redefined, as it might become necessary to make risk-related decision making sufficiently agile and adaptable in the ever-changing environments of program/project and institutional activity executions.

2.2.3.1 *Application of Risk Leadership Principles in the Objectives-Driven Risk Management Framework*

The application of risk leadership principles within the NASA ODRMF involves the execution of certain specific operational and communication processes that are not defined solely within the confines of risk management, and that in fact in some measure transcend the latter as more general and necessary premises for its implementation and execution.

The following implementation steps can be considered as being essential to a successful application of risk leadership principles by all personnel, from high level managers to technical staff members, involved in the definition and execution of risk management and risk-related decision functions:

- Definition and Communication of Organizational Objectives

A clear definition and statement of an organization's objectives (both the fundamental and means objectives), values and priorities, is the foundation upon which an effective risk management system can be built. This can be understood from the very definition of what constitutes risk in the context of an organization such as NASA, i.e. "risk to the achievement of objectives" (see NASA's definition of risk in Section 2.1 of this handbook and in NPR 8000.4C). It follows from that definition that risk cannot be described, let alone assessed and managed, if the potentially impacted objectives are not clearly identified, communicated, and understood throughout an enterprise and organization. It also follows that the responsibility for the clear setting, declaration and communication of organizational objectives is a management responsibility that constitutes an essential component of risk leadership.

- Definition and Communication of Risk Posture

The management of risk involves the execution of decision processes at all levels of an organization. If risk management decisions are to be elaborated and made with any degree of consistency and coherence, it is necessary that the risk leadership exerted from the higher levels of the enterprise include clear guidance on the risk posture that the organization leaders have agreed upon and wish to see applied at all levels of the organization. This should remain valid regardless of whether the decision subject is risk to objectives established at the enterprise, institutional, or program and project level. To achieve consistency and coherence in risk posture across a large and complex organization is no easy feat. In this regard, risk leadership requires that qualitative criteria, elaborated and agreed upon at the top of the organization, be clearly and effectively communicated to lower-level managers, so that the latter fully understand how they are expected to make decisions on the types and levels of risk that can be accepted, or not, in the pursuit of organizational objectives. A visual illustration of the relation between the identification and definition of organizational objectives and the definition of a corresponding risk posture is provided in the upper portion of Figure 2-4.

- Translation of Risk Posture Directives into Risk Tolerance Criteria for Objectives

Risk posture can be defined by high level leaders in global qualitative terms, but, as illustrated in the lower portion of Figure 2-4, it is the responsibility of lower-level managers and technical risk experts to translate such guidance into risk tolerance criteria relative to the performance measures by which the fulfillment of a specific project / activity objectives is evaluated. This translation of risk posture guidance into operational risk tolerance criteria, and eventually into the definition of quantitative risk tolerance measures applicable in the pursuit of an activity objectives, may in practice require iteration and feedback steps between the higher and lower levels of the organization hierarchy, but is a critical step to avoid ambiguity in the definition and misinterpretations in the application of risk leadership

and posture guidance. Figure 2-5 conceptually illustrates this key risk leadership application principle.




Figure 2-4. Flow of Risk Leadership Application within the Organizational Hierarchy




Figure 2-5. Definition of Risk Posture and Tolerance Levels According to Risk Leadership Principles

- Empowerment of Decision Makers and Control of Risk Decision Processes
The next key principle to be applied in Risk Leadership is a controlled empowerment of risk management decision making at all organizational management levels. “Controlled empowerment” is not a contradiction in terms: on the contrary, it means that decisional power is delegated and allocated within clearly assigned limits of acceptable risk, and that a corresponding continuous monitoring of how risk management decision-making is executed must be also carried out. Monitoring and control is necessary not to second guess every risk related decision made by lower-level managers, but to verify that the goals of consistency and coherence in risk posture established at the top are reflected by the actually executed decision processes, and, in cases where issues and deviations might be found in this regard, to apply the directives and corrections that may be necessary.
- Well Established Communication Lines and Protocols for Risk Information and Outcomes
The last and fundamental principle of effective Risk Leadership implementation is the enablement and encouragement of proactive communication of pertinent risk information and RM application results transversally across organizational units that are potentially affected and vertically, especially in the direction of providing feedback to higher level managers and leaders on the practical effect of the application of risk posture and risk tolerance criteria in the day-to-day execution of project / activity RM tasks and processes. Appropriate information prioritizing protocols are necessary to avoid excessive and needless misuse of communication lines, however an effective communication and feedback mechanism is crucial to prevent cross-cutting risks from being overlooked or underestimated and to provide risk leaders with the information they may need to adjust or recalibrate their risk posture perspective vis-à-vis the actual outcomes of the related RM activities.

Figure 2-6 illustrates the key aspects of a full cycle of integrated risk management that implements and applies the risk leadership principles introduced and discussed in this section. The figure identifies in this respect: a) the RIDM-supported top-down definition of activity objectives and associated risk posture and tolerances; b) the CRM-supported life-cycle risk identification, risk assessment, and risk control activities; and c) the bottom-up feedback and reporting of the outcomes of these activities to higher organizational levels. The figure also identifies in the pyramid at its center the principal underlying elements of project and activity management with which the risk management and leadership cycle actions are interconnected, and therefore are to be coordinated. These elements are identified starting at the top with the management agents and entities that provide direction and definition, proceeding down onto the project/mission execution activities and support tasks within which the risk leadership directives are to be applied.

Risk Posture, Risk Tolerance, and Appetite for Seeking Opportunity

“Risk Posture” is the attitudinal framework that expresses an organization’s valuation of the potential gains and losses that may result from its activities and operations. The adopted Risk Posture guides the organization in its Decision-Making, enabling it to strike a balance between its appetite for seeking opportunities and its level of tolerance for the concomitant risks. An organization’s Risk Posture defines the character of the organization in terms of its eagerness to accept challenges in the pursuit of potential gain, vs. its desire to avoid risks in order to minimize the potential for loss.

Risk Posture may be initially defined and communicated by the leadership of an organization in qualitative terms, however it is eventually articulated in more specific quantitative terms via the definition of Risk Tolerance Levels (RTLs), i.e., risk thresholds that express the limits of probability of shortfall or loss the organization is willing to accept in pursuit of achieving particular objectives (see definition and discussion in Section 3.3.1). A full set of RTLs comprises the operational definition of Risk Posture that an organization seeks to express and apply across the range of organizational objectives and performance dimensions. As such, RTLs may be applied to objectives whose measure of success is ultimately observable (e.g., the objective of obtaining X years of useful data), as well as to objectives whose measure of success can only be estimated probabilistically but not directly observed (e.g., the objective of maintaining a probability of loss of crew less than Y). Risk Posture is ultimately defined in operational terms by the balance and apportionment of the RTLs across all the objectives of a given activity, and expresses the organization’s view of how the likelihood of gains or shortfalls in the pursuit of any one objective balances against the corresponding likelihoods of gains or shortfalls in the pursuit of the other objectives.

More detailed perspectives on Risk Posture and the other elements of risk terminology introduced here, including the probabilistic interpretation of these terms, are presented and discussed in Chapter 3.

Application of Risk Leadership Integrated with Agency Management Processes at All Levels

- Top-down definition and communication of risk posture and risk-tolerance in key mission / activity dimensions
- Bottom-up feedback on activity execution, application of risk-tolerances and effectiveness of associated risk controls



Figure 2-6. Risk Leadership and Management Integrated Implementation and Application Cycle

2.2.3.2 Desirable Risk Management Outcomes Enabled by Risk Leadership

Beyond its earlier-provided definition, the meaning and significance of risk leadership in relation to project and activity management and its risk management aspects can be further identified by considering the principal effects that its application is intended to produce in the risk-related management and decision processes that risk management executes and supports. The following are key desired outcomes that can be identified from such a perspective:

- Balancing of risk against benefits and opportunities
- Reduction of “process burden”
- Balance between risk assessment rigor and decision velocity

In the pursuit of these outcomes, effective risk leadership additionally assumes the responsibility of providing confidence in the robustness of the risk-related decisions that are made. In general, confidence is obtained by assuring that the generation of data and analyses used to support each decision has exercised the level of rigor needed to provide the desired confidence. The objectives of reducing process burden and increasing decision velocity mandate that the level of applied rigor is both necessary and sufficient for the specific decisions to be made. In other words, sufficient rigor is to be expected but excessive rigor is to be avoided. The definition of “decision robustness” provided in the blue box below is consistent with this description of necessary and sufficient rigor.

Decision Robustness

A robust decision is one that is based on sufficient technical evidence and characterization of uncertainties to determine that: a) the selected decision alternative best reflects decision-maker preferences and values, consistently with the informing state of knowledge at the time of the decision, and b) can be deemed to be insensitive to credible modeling perturbations and realistically foreseeable new information.

With regard to the balancing of risks and opportunities, there often is a tendency for letting the related decision processes be influenced by psychological factors that are not always in the interest of applying a truly balanced judgment. Studies on factors affecting decision processes show that, when people are confronted with two choices where the balance between opportunity for success and risk of loss is neutral or even moderately favorable to the opportunity, a majority will tend to choose the path with lower risk. This aversion is related to the Ellsberg paradox [10], which concerns people's choices between situations that exhibit different levels of certainty (i.e., most people have “ambiguity aversion”). Use of risk leadership and risk management in a structured approach helps counter risk aversion and ambiguity aversion by ensuring that strategic decisions are made more objectively.

Ultimately, the decision makers have the responsibility to define their risk posture, rather than simply accept a risk-averse or ambiguity-averse stance. When this responsibility is accepted and applied, the ensuing implementation of risk management will not be perceived by the managers and staff members accountable for the practical execution of project tasks as being a cause of “process burden.” Thus, it is important that all involved understand that risk management

execution is to be kept consistent from the top down with the overall objectives of the projects and missions for which the associated activities are applied, and with all other program and mission execution constraints. In this perspective, the application of risk leadership must inform and infuse risk management, by means of its upfront identification, communication, and clarification of mission objectives, risk posture, and efficient risk processes. If this is achieved, risk management can then be applied throughout the levels and ramifications of program, project, or institutional executions, in a way that is conducive to the desired outcomes of balancing risk and opportunity, containing process burden, and improving decision velocity and effectiveness.

Direct or indirect reference to the above desired outcomes can be found both in official NASA documents and in the statements of top-level managers. NPD 1000.0C makes explicit reference to applying a “proper risk posture,” and it is not difficult to view that as being in direct correlation with the above desired outcome of balancing risks and opportunities. The directive also explicitly identifies “decision velocity” as a risk leadership goal.

The definition and consistent assumption of a proper risk posture, as advocated by NPD 1000.0C, is not just as an enabling condition of the application of risk leadership, but should be pursued as one of its primary objectives. As an indispensable complement of its definition and adoption, effective communication of risk information, properly contextualized by providing the perspective of the objectives and opportunities with which it is correlated, needs to flow from the top down within the organization, and also transversally and in feedback mode from the bottom up, as suggested by the cross-cutting nature of the risk involved and/or by the magnitude of risk mitigation resources that may be needed. An open and unimpeded flow of risk information is a key component of the implementation of risk leadership and of the process of effectively risk-informing program and mission decision processes. A reduction of decision process burden and enhancement of decision velocity may not be a guaranteed effect of fostering a balanced risk posture and risk-informing decisions with properly contextualized risk information, but it is certainly facilitated by the pursuit of these principles.

In concluding the introduction of the above concepts, it is useful to consider the words of a notable Harvard Business School study and publication [11] detailing the characteristics and attributes of what it considers case-study examples of successful implementation and execution of risk management philosophies and processes. In its closing statements, the study identifies the management leadership principles that it found to be common denominators of the successful case studies examined, and which, together, it considers to be key for the effective application of risk management in all functional and programmatic areas of a large enterprise:

“In fact, coming to an agreement about the company’s belief system, about its objectives, values and priorities, is the first—and in some ways the most important—of the three parts of developing an effective risk management system. The second is to formulate the firm’s risk posture about how much and what kind of risk can be tolerated. And third is the continuous monitoring and benchmarking of a firm’s risk-taking behavior against its risk posture. The firm’s risk posture should clarify what risks can be accepted and left unattended, and what risks need immediate attention and action. It starts with the company leadership team reaffirming its mission and values.” [Note: In the quote, in keeping with the terminology in the blue

box of Section 2.2.3.1 defining risk posture, we have substituted the term ‘risk posture’ for the term ‘risk appetite,’ which was used in the original source.]

It can be easily recognized that a substantial alignment and consistency exists between the principles stated in the above quote and the NASA NPD 1000.0 risk leadership concept and objectives—i.e., more specifically, in underscoring the importance of a shared risk posture and of effective risk-informed decision-making. It is thus worthwhile to expand on the concepts quoted above from the paper, into a discussion of what risk leadership philosophy and principles may be applicable across the NASA enterprise, and how their application can provide much needed context, at all organizational levels, for an effective implementation and execution of risk management.

2.2.4 Rigor in Risk Assessment and Decision Processes

The application of well-balanced rigor is a principle to which all risk management implementations must adhere if the applicable policies and guidelines are to be translated effectively into day-to-day application practices. That this theme deserves to be upheld as a risk management guiding principle is underscored by the RMTT report [9]. In key observations and recommendations related to the theme, the report in fact finds that:

- *“NASA’s RM Policy is sufficient, but implementation improvements and additional consistency are needed across the agency.”*
- *“Additional rigor is required in pre-formulation.”*
- *“Risk Training is under-utilized, ...”*

Applying rigor in an effective and efficient manner across risk management processes has multiple challenges. The first is in the historical roots of risk management practices, whereby a majority of past application processes were not necessarily linked or associated with supporting formal techniques of risk assessment and decision making. The lack of specific training of personnel charged with the execution of risk management processes aggravated the problem, in that those techniques, even when proposed by guidance documents like the earlier version of this handbook, were frequently not well understood, or were even perceived as an unnecessary burden, rather than as an aid for successful risk management executions. An example of this issue is the past underutilization of RIDM as an activity pre-formulation instrument, even though it was originally formulated specifically to support the risk-informed decisions by which the course of projects or institutional activities is selected before their actual execution is initiated. It is easy to see the connection between this under-utilization and the above Ref. [9] call for additional rigor in pre-formulation.

Ref. [9] also recommends *“Ensuring consistent knowledge and execution of the Risk Framework.”* This recommendation can be connected to the above recommendation for greater rigor if one interprets it as a call for a consistently balanced implementation of the NASA risk framework, by which the depth of utilization of formal technical risk assessments and decision-support instruments is commensurate to the importance of the risk related decisions to be made, and consistent with the risk posture deliberately adopted and applied in the affected activities. In this respect the application of well-balanced risk management rigor can also be connected to the

concept of a “graded approach” in the utilization of risk management processes and techniques – an approach which is predicated in this handbook (see Sections 4.11 and 5.3.1) as the means by which rigor is pursued not for its own sake, but in a way commensurate to what is at stake and to the nature of the decisions that are supported by the relative assessment and deliberation processes.

The application of rigor as a principle of effective risk management is therefore one that applies not just to the technical analyses supporting the latter, but even more importantly to the decision processes within or directly connected to it. The context for this is provided by the policy set forth in NPD 1000.5, where it states:

“It is NASA policy to ... incorporate a risk-informed decision-making process that includes the identification, analysis, and management of programmatic, institutional, technical, cost, schedule, environmental, safety, management, industry, and external policy risks that might jeopardize the successful execution of the Agency's acquisition strategies. The process shall include the philosophy of Risk Leadership, as described in NPD 1000.0.”

In the context of decision-making, rigor concerns the clear identification and documentation of the factors used as input to the decision process, of the basis of deliberation, and of the rationales that support the selected courses of action that constitute the decision output. Balanced rigor means that the above can be articulated and tailored at different levels of complexity depending on what is at stake and on the corresponding risk posture, but also that the essential elements of deliberation and documentation should always be present and traceable in any risk-related decision process of consequence. This is especially necessary in decisions concerning risk acceptance. For these, NPR 8000.4 prescribes that explicit decision criteria be established in an organization's Risk Management Plan (RMP). Per the NPR, in this respect the RMP:

“...Reflects the overall programmatic risk posture by documenting risk acceptance criteria/thresholds and elevation protocols (the specific conditions under which a risk management decision is elevated through management to the next higher level).”

Beyond the above, and as the discussion in this section also advocates, the NPR predicates a graded and balanced application of rigor in the section dedicated to the requirements for the RIDM process [1]:

“The manager shall ensure that key decisions, including risk acceptance decisions, are informed by Analysis of Alternatives carried out by applying the RIDM process with a level of rigor that is commensurate with the significance and the complexity of the decisions.”

The discussion of the principle of rigor in risk management can be concluded by summarizing and emphasizing its key corollaries:

- Rigor applies to both the definition and documentation of risk-related decision processes and to the selection and utilization of the analytical tools that support them.
- Rigor is to be applied in a balanced fashion that takes into account the nature of the risks being considered and the importance of the related decisions.

- Risk leadership sets the risk posture that in any given organization provides the context for the balanced and graded application of rigor, via the definition of criteria for opportunity-seeking and risk-acceptance.
- Balanced application of rigor also involves, besides the appropriate utilization of risk-related decision criteria and supporting assessment tools, the careful definition, selection, and tailoring of technical and safety standards, consistent with the adopted institutional or programmatic risk postures.

2.2.5 Pros and Cons of Simplification versus Rigor in Graded Approach to Risk Assessment and Management

One of the practical principles set forth in this handbook is that in the application of risk management processes, and of the risk analysis techniques that are utilized to support them, a **graded approach** should in general be followed. This principle is discussed and exemplified in several passages. More specifically, Chapter 6 provides guidance on how it should be formalized and documented for any given activity or project context in the **risk management plan** (RMP) that constitutes the declared roadmap for implementation of risk management in compliance with the requirements of NPR 8000.4 [1]. The RMP should therefore be the first source of guidance which is accessible within a project or activity for defining the degree of tailoring and adaptation of recommended risk management and analysis standards and practices. The goal of any such graded approach definition and declaration is to achieve an appropriate balance between the goals of effectiveness and utility of the risk management processes and the inevitable constraints programmatically existing on the resource practically available for their implementation and execution.

To illustrate what the application of the graded approach principle may entail in terms of the selection of processes and analytical techniques of greater or lesser level of sophistication and accuracy, certain parts of this handbook present both examples of more accurate and complex applications of fully quantitative probabilistic risk analytical processes and techniques and of simplified semi-quantitative types of analyses or assessments. This introductory discussion of the graded approach subject, on the other hand, intends to provide an initial but quite necessary overall perspective on the pros and cons that are associated with the selection of one type of approach to risk management and analysis instead of another.

Between the two theoretical ends of the spectrum of possible approaches, i.e., a one hundred percent fully quantitative framework based on rigorous logic-probabilistic models of risk at one end, and a purely qualitative framework assigning risks to qualitatively defined levels of classification at the other end, different degrees of “mix and match” of rigorous and simplified processes and techniques is theoretically possible. Consistently with the varying degree of risk management and assessment techniques “grading” that may be chosen and applied, Table 2-II lists a set of relevant associated pros and cons factors in relative comparative rather than absolute terms. The contents of the table are self-explanatory, and it is left to the reader to judge the relevance of the factors considered therein in the specific context of the activity or project to which the risk management process is to be applied.

Table 2-II. Risk Management Approach Advantage & Disadvantage Factors

		ORIENTATION OF RISK MANAGEMENT & ASSESSMENT APPROACH	
		<i>More simplified and qualitative</i>	<i>More rigorous and quantitative</i>
ADVANTAGE & DISADVANTAGE FACTORS	<i>Resources for Execution</i>	<i>Less</i>	<i>More</i>
	<i>Ease of set-up & Execution</i>	<i>Easier</i>	<i>More challenging</i>
	<i>Interpretation & Communication of Results</i>	<i>Easier</i>	<i>More challenging</i>
	<i>Objectivity of Results</i>	<i>Higher incidence of subjective bias</i>	<i>Greater objectivity</i>
	<i>Validity of Results</i>	<i>Higher incidence of assessment errors</i>	<i>More realistic assessment</i>
	<i>Level of Depth of Results</i>	<i>Less depth & granularity</i>	<i>Greater depth and granularity</i>
	<i>Inherent Uncertainty in Results</i>	<i>More inherent uncertainty in assessment models</i>	<i>Less inherent uncertainty in assessment models</i>

Besides the pros and cons of what can be expected in certain key characteristics of a selected type of risk management and assessment approach, it is also useful to bring to the reader's attention some additional factors that should be given due weight in the formulation of the type of risk management approach to be chosen. Ref. [12] appropriately suggests that certain intrinsic characteristics of risk-related decision concerning an activity or project should be drivers for the selection of a formal RIDM process as an activity planning and pre-execution risk management and decision support process of choice. These characteristics are identified in [12] as follows:

- High Stakes — High stakes are involved in the decision, such as significant costs, significant potential safety impacts, or the importance of meeting the objectives.
- Complexity — The actual ramifications of alternatives are difficult to understand without detailed analysis.
- Uncertainty — Uncertainty in key inputs creates substantial uncertainty in the outcome of the decision alternatives and points to risks that may need to be managed.
- Multiple Attributes — Greater numbers of attributes cause a greater need for formal analysis.
- Diversity of Stakeholders — Extra attention is warranted to clarify objectives and formulate performance measures when the set of stakeholders reflects a diversity of values, preferences, and perspectives.

When such characteristics are prominent attributes of the conditions concerning an activity or project, they constitute a strong motivation for the activity/project decision makers and leaders to adopt a ***low or very-low risk tolerance posture***. The adoption of more rigorous and formal

definitions and applications of risk management is a choice that should be made whenever it is decided that an activity or project is to be set-up and executed with a posture of low or very-low risk tolerance.

Before closing the discussion on risk management rigor and graded approach choices, two more considerations need to be made concerning the characteristics of the output that can be expected as a result of the type of approach that is selected. The first consideration concerns a common misconception that exists when the terms “qualitative” and “quantitative” are used to refer to analytical approaches, and more specifically to the approaches that may be adopted when the assessment of risk and related parameter is concerned. The misconception is in the fact that the key difference between a “qualitative” and a “quantitative” approach is often perceived as being definable in terms of the degree of accuracy and precision that the available information allows the analysis to achieve. I.e., the choice between the adoption of a “qualitative” versus “quantitative” approach is assumed to have to be driven by how precise and accurate the result can be, in view of the type of information which is available for feeding either into either type of assessment. According to this interpretation of what constitutes the key attribute of a qualitative or quantitative approach, the use of quantitative forms of assessment is not worth adopting when the information that is available for input is “soft,” as such analyses will produce results that are deceptively “precise” (i.e., very detailed and granular) at face value, but actually inaccurate and invalid (because of the “garbage in, garbage out” effect).

The above perspective needs to be corrected, as the real issue is not in the inherent precision or uncertainty of a qualitative or quantitative assessment, but in whether the framework and parameters of the assessment are set up in objectively defined terms. A qualitative definition of assessment scales, if not anchored by objectively defined bounds or intervals, is often subjectively biased by the perspective of the individuals or groups who provide the definitions of the qualitative scales. For example, a project cost defined as being “high” by an assessor who is operating within the context of a low budget project would instead be seen as being “low” or even “insignificant” when judged by enterprise-level entity.

A better perspective from which to judge the merits of qualitative or quantitative risk-parameter assessment criteria is reached by recognizing that a definition in quantitative terms is needed to provide objectivity, not precision. The definition and use of quantitative measurement scales that can be understood independently of the context of the assessment and/or communication of the results is a means of injecting objectivity in the assessment of concern. It does not mean that levels of precision the existing information cannot support are artificially created and forced into the assessment process. Any “soft” or uncertain nature of the assessment input information can easily be recognized and properly accounted for by providing definition and quantification of uncertainty ranges, however wide, that affect the estimation of relevant assessment quantities and parameters.

A last important consideration to be made in regard to the choice of an appropriate type of risk management graded-approach concerns the intended utilization of the risk related information. Risk management is intrinsically a tool supporting the decision-making processes that activity and project managers apply day-in and day-out in the execution of their jobs. Thus, the quality of the risk management output should match the importance and potential impact of the decisions for

which it is used. Both in the context of the RIDM processes utilized in pre-execution activity and project planning, and in the context of CRM processes whereby risk control measures need to be identified and selected, decision makers frequently face the same type of “cost vs. benefit” questions. That is, between alternative project path, design solution, and/or risk control measure A, B, C, etc., to decide and select from, the key question to determine an answer for is which is the most appropriate and favorable in terms of risk control, system effectiveness, and resource expenditure. The assessment processes that may provide the answer to the question, including those involving specific types of risk assessment and management processes and techniques, should be chosen according to whether the level of depth and resolution in their outputs is sufficient to clearly distinguish between the advantages and disadvantages of the alternatives being decided upon. Overly simplified approaches that limit the depth of their results to an approximate “binning” of risk levels may not be provide enough resolution in decision processes where the ultimate impact may involve the ultimate success of an important mission, or the expenditure of program resources in the order of tens or hundreds of millions of dollars.

2.2.6 Risk Management Integration Across Technical, Organizational, and Life-cycle Boundaries

The integration and coordination of risk management across all agency activities is the last, but certainly not least important of the NASA ODRMF focal elements discussed in this chapter. As a theme, it is closely related to the theme of integration of risk perspective discussed earlier in Section 2.2.1, however the aspects of integration to be considered within it are those that more specifically concern the implementation of operational risk management provisions and processes across:

- a) The different major stages of a project or activity life cycle,
- b) The areas of responsibility assigned to organizational entities and programs, and
- c) The boundaries of technical disciplines and domains.

Historically, the implementation of risk management at NASA was initially focused on the identification and handling of individual risk scenarios of a technical, safety, cost or schedule relevance within specific programs or projects, once their execution was underway. This was entrusted to the implementation of the Continuous Risk Management (CRM) process by programs and projects. Subsequently, the Risk-Informed Decision Making (RIDM) process was developed for the primary purpose of addressing risk from an Analysis of Alternatives (AoA) perspective in the concept selection and pre-execution stages of projects and activities. More than a decade after its introduction, RIDM has to date seen limited application in the contexts for which it was initially conceived. This is indirectly noted in the previously cited RMTT report [9] where it acknowledges that NPR 8000.4 – which prominently identifies RIDM as the risk management tool to be used to support risk-informed pre-execution decisions concerning the selection of mission design concepts and development strategies – represents a “reasonable framework for Risk Management,” while at the same time identifying “adding rigor to early-formulation” as one of the top-priority recommendations to “enhance the Risk Management system's effectiveness to address key issues.”

This handbook emphasizes that the two primary instruments of risk management implementation within the NASA ODRMF, i.e. RIDM and CRM need to be applied according to the stage and nature of the risk management issues that are to be dealt with.

2.2.6.1 Integration of RIDM and CRM with Management and Acquisition Activities across Development Stages

Although they are structured according to, and thus also defined by, their own process and analytical steps, RIDM and CRM are intended as complementary components that support an integrated application of RM from the onset to the conclusion of any major NASA mission, project, initiative, or activity. This includes acquisition activities where NASA chooses to delegate the actual development of a system, or even its operational utilization in execution of a mission, to a commercial contractor. In fact, as Figure 2-7 conceptually illustrates, the interlacing and interaction of RIDM, CRM and organizational management can be expected take place from beginning to end of any acquisition life cycle where a *primary Acquirer* and a *primary Provider* organization play the major roles².




Figure 2-7. High-Level View of the Interfaces of RIDM and CRM with Organizational Management

It is further noted that, although the formal instruments of *Acquirer-to-Provider* delegation (i.e., MOUs, formal contracts, etc.) change from the cases where both the *Acquirer* and *Provider* organizations are NASA entities to the cases where the *Acquirer* is a NASA entity and the *Provider* is a commercial contractor, the overall and top-level framework of management activities, including risk management ones, is to be selected and defined by the *Acquirer* organization before any formalization occurs. That is, it remains in all cases the *Acquirer's* privilege and responsibility to decide the degree of delegation it wants to see applied in the implementation of a system

² The term “primary” is used here to indicate the top-level of Acquirer / provider interfacing and interaction, regardless of any further delegation and subcontracting that may be put into effect and occur at lower levels.

acquisition and the execution of an associated mission. This concerns also and includes the risk management activities that may be implemented via the RIDM and CRM processes in the course of a system acquisition and mission execution life cycle. Even in the case of a “turn-key” purchase of a service or system, where the *Provider* is granted the largest degree of delegation of decision authority and responsibility, the NASA *Acquirer* entity may use RIDM in the planning stages of the turn-key acquisition to risk inform the initial selection of the *Provider* entity to which such authority and responsibility will be delegated, and to identify the principal programmatic and technical risks associated with each of the possible selections. Furthermore, even in such a form of “hands off” acquisition, the *Acquirer* may want to contractually require as a condition of the system or service purchase that the *Provider* address the risks initially identified by the *Acquirer* via a process that can be mapped to a typical CRM set of management steps and that produces CRM-like documentation for the *Acquirer* to review, and possibly use for decisions to be made at any predefined acquisition key decision points (KDPs).

In Figure 2-7, the flow of activities for organizational management presented in the center box is categorized according to the definition of three conceptual stages of activity development: a) concept / execution strategy selection, b) implementation budgeting / planning, and c) activity execution / product utilization. Consistently with the general considerations relative to degree of delegation of decision authority between *Acquirer* and *Provider*, the figure indicates that the former is generally taking responsibility for activity and acquisition strategy selection decisions, while the latter is usually delegated at least a primary portion of the activity and mission execution decisions. Responsibility for execution planning decisions is shown to be assigned either to the former or latter type of process player, depending on the type of acquisition strategy initially selected by the *Acquirer*. In the case of programs and projects defined and executed according to the established canons of systems engineering, more formally defined life-cycle stages that characterize program/project timelines and progress from one “key decision point” (KDP) to another are operationally and customarily referred to with specific terminology, however those stages can still be conceptually mapped to the above simplified definitions.

A more detailed depiction of how an integrated flow of risk management application is to be executed via a selection of RIDM and CRM coordinated processes is provided by Figure 2-8. The figure indicates the sequence and coordination of RIDM AoA and traditional CRM steps, under the common and less common conditions that can be encountered in the course of an activity or project definition and execution life cycle. The corresponding modes of RIDM and/or CRM execution are defined and explained below.




Figure 2-8. Coordinated Use of RIDM and CRM in the Risk Management of an Activity

- In the pre-execution stages of any activity, and as part of risk-informing the associated AoA processes, *Activity-Planning RIDM* identifies a set of top-level performance risks for each of the candidate alternative mission, system, or institutional strategy and plan solutions that are being considered in the AoA. The identification of these risk profiles for each of the alternatives being considered enables the selection and planning of the solution that will be actually implemented and executed to be carried out with an appropriate consideration of the balance between risk and benefits that the deliberating body considers optimal. Given the importance of this element of risk information in the AoA process, it is the responsibility of the management authority that assembles and controls the deliberating body or council that carries out the activity selection and planning deliberations to make sure that an appropriate cadre of technical support experts is also assembled to carry out the Activity-Planning RIDM steps and tasks (as per the detailed steps of the process discussed in Chapter 4) and provide the resulting information to the deliberating body. For projects and missions that are organized and evolved according to the formal stages of a standard NASA systems engineering life cycle, the Activity-Planning RIDM should be concluded before the Mission Concept Review (MCR) and its results should be made part of the Formulation Authorization Document for the project and/or mission. For other types of activities an equivalent timetable and formalization of resulting information should be identified and implemented. Once a specific solution among the alternative candidates is identified and selected as a result of the selection and planning process, the set of risks identified by RIDM for that solution constitutes the initial profile communicated to the activity management and to the CRM process. CRM therefore receives these individual risk scenario contributors identified by RIDM as an initial set to be further assessed, handled, and tracked as the project or activity progresses from the planning to the implementation and execution stages.
- In the implementation and execution stages, the *Activity-Execution CRM* process is applied for the identification, control and tracking of individual risk scenarios, with additional provisions for the assessment and control of unknown and/or underappreciated (U/U) risks, and of aggregate risk (as fully discussed in Chapter 5). Opportunity management, i.e., the identification and evaluation of any conditions or openings in technical or operational conditions that may bring additional benefits from the execution of a project or activity, should also be part of a full application of the CRM process. Within the CRM execution framework, the weighing of plausible risk mitigation alternatives against one another may require rigorous analysis in the form of a RIDM-based AoA that considers risk reduction worth against other factors such as cost of implementation or introduction of other forms of risk, so that an optimal risk mitigation solution can be reliably selected. That is, under such conditions CRM necessitates the coordinated application of an *Activity-Execution RIDM* AoA for the selection of optimal risk mitigation solutions, among the alternatives that may be theoretically possible. In this context, an “optimal risk mitigation” solution may be defined and identified on the basis of criteria that are tailored for the context of the risk being considered. A typical, and relatively simple optimization criterion is one where risk reduction for the performance measure of concern is compared with the cost of implementation and application of the controls being

considered, and with any “collateral” risk profile impacts on other performance measures that may be impacted by the control application.

- A major type of condition may occasionally occur, again in the implementation and execution stages, when the control and handling of relevant risks cannot be accomplished without serious interventions on the design and/or operational requirements for a given system. If it is recognized within the execution of the corresponding CRM “Plan” step that significant changes from, or additions to, the initial project / activity objectives or requirements may be needed in order to keep execution risk at acceptable levels, to the extent of involving a reconsideration of design and/or operational alternatives for the activity, then a reactivation of the AoA function of RIDM may be accordingly resorted to, in order to provide a new risk-informed basis for the rebaselining of the project or activity objectives or requirements and associated risk posture. Such a reactivation is referred to as *Activity-Rebaseline RIDM*. Given the major implications and impacts that activity rebaselining would generally have, it can be executed only upon deliberation and approval by the project or activity accountable managers. As in the case of activity-planning applications, the results produced by RIDM in this rebaselining mode become input to the CRM processes for the rebaselined project or activity.
- A different type of rebaseline trigger-condition may be encountered, again in the activity execution stage and within the opportunity management side of CRM, when an opportunity of major potential benefit – e.g., the availability of a previously not fully mature technology that would greatly reduce the cost and schedule of a project execution – is identified as exploitable via the redefinition of an activity objectives and execution plan. Under such a condition the AoA function of *Activity-Rebaseline RIDM* may be invoked, with accountable managers’ approval, to risk-inform the decision of whether, and by what practical execution path, such a major opportunity should be pursued.

2.2.6.2 Adherence to the Risk Posture throughout the Program/Project Life Cycle

NASA programs and projects are initiated and implemented to accomplish objectives that flow down from the Agency’s Strategic Plan. Consistent with the philosophy of risk leadership, risk postures flow down in tandem with the flow-down of objectives, and it is the purpose of each program/project RMP to document the approach that the program/project will take to ensure that the risk to the achievement of its objectives is within the risk posture established for it [1].

Programs/projects are partitioned into life-cycle phases, each of which has one or more associated life-cycle reviews (LCRs) that provide a periodic assessment the program/project technical and programmatic status and health at key points in the life cycle. For spaceflight programs and projects, NPR 7120.5 requires the participation of an independent Standing Review Board (SRB) in the conduct of the System Requirements Review (SRR), System Definition Review (SDR)/ Mission Definition Review (MDR), Preliminary Design Review (PDR), Critical Design Review (CDR), System Integration Review (SIR), Operational Readiness Review (ORR), and Program Implementation Review (PIR).

At LCRs, programs/projects are assessed against six criteria, one of which is the adequacy of the risk management approach [13]. These criteria are specialized to the specific objectives of each

life-cycle phase and LCR, resulting in LCR-specific sets of LCR success criteria that are used as the basis for program/project evaluation. Guidance on the development of LCR success criteria is presented in [14], with the expectation that the criteria will be customized appropriately to the particular program/project being reviewed.

With respect to the adequacy of the risk management approach, the LCR success criteria should be developed in enough detail to provide a sound basis for determining whether or not the risk to the achievement of the program's/project's top-level objectives is within the established risk posture, given the state of the program/project at the time of the review. The specific criteria should cover all aspects of the program/project upon which adherence to the risk posture depends. Early in the life cycle, before top-level program/project requirements have been baselined, the criteria are expected to focus heavily on the adequacy of RIDM during concept studies. During formulation, the criteria are expected to be more of a mix of RIDM-related and CRM-related criteria, focusing both on the program's/project's exposure to risk as a result of decision-making, as well as on the management of the risks that are taken on board as a result of decision-making. During implementation, the criteria are expected to focus more heavily on CRM and the identification and management of emerging sources of risk.

Table 2-III illustrates the kinds of topics that might be addressed by the LCR success criteria. The items in the table do not represent LCR success criteria themselves, nor is the table partitioned into specific LCRs. Rather, the listed items are meant to illustrate the kinds of topics that could reasonably be expected to be addressed by LCR success criteria in the corresponding portions of the life cycle.

Table 2-III. Notional LCR Success Criteria Topics

LCR Success Criteria Topics Addressing the Adequacy of the Risk Management Approach (Notional)	
Pre-Formulation	• Concept studies included a broad range of ideas • Selection of the preferred concept(s) was risk-informed • No "showstopper" sources of risk were identified
Formulation	• A feasible risk posture has been established • The risk management effort is adequately resourced • System definition decisions were risk informed • Major sources of risk have been analyzed and mitigation strategies have been defined • The risk posture has been flowed down to subordinate organizations in tandem with the flow-down of objectives
Implementation	• A credible case was made that the program/project is within its established risk posture • The program/project is operating within defined limits and in accordance with programmatic commitments • Effective measures are in place to identify and manage emerging sources of risk

The success criteria for a given LCR should be developed prior to the execution of the associated life-cycle phase. NPR 7120.5 requires the baselining of all LCR success criteria at System Requirements Review (SRR). The RMP should align with the LCR success criteria, in that the risk management activities conducted in each phase should be directed towards satisfying the LCR success criteria of that phase (and possibly those of future phases, for longer-term activities). Moreover, the set of risk management activities should include the production of the evidence that will be used at each LCR to substantiate that the success criteria have been met. Like the success criteria themselves, the evidence that will be used to substantiate their accomplishment should be specified prior to the execution of the phase, to prevent “backfitting” a case based on unspecified evidence.³

2.2.6.3 Risk Management Coordination in Alignment with Organizational Objectives

NASA risk management policy guidance provided by NPR 8000.4C defines risk in terms of the severity and likelihood by which the declared objectives of an activity may be missed, due to any possible sets of conditions and causes. It is therefore a necessary precondition for an integrated and coordinated execution of risk management across organizational boundaries to have an overall understanding of the relation between Agency high-level strategic objectives – i.e., the objectives that are pursued within the “Enterprise Domain” – and the more specific objectives, belonging to the “Project/Program Domain” and/or “Institutional Domain” spheres of activity, into which the Enterprise Domain objectives are decomposed and allocated.

The definition of Agency organizational units, and of the hierarchical and functional framework within which they are to operate, is inevitably subject to changes in time. Nevertheless, a sound basis for integrated and effective risk management plans and processes can be identified and applied only if a clear vision and perspective of the logic flow of agency objectives, from the higher Enterprise Domain level to the lower Program/Project Domain and Institutional Domain execution levels can be established.

A conceptual illustration of the interrelation and flow-down of objectives from the strategic enterprise domain level and across the program/project and Institutional domains was provided earlier in Section 2.1.3 by Figure 2-2. It remains the responsibility of the activity execution entities and/or organizational units to which the realization of objectives is allocated – see in this regard the illustration provided in Section 2.1.4 by Figure 2-3 – to identify what lines and protocols of risk communication, elevation and handling coordination should be followed for an integrated execution of risk management, aligned with the logic and hierarchical interconnections of organizational objectives across activity execution and risk domains. The identification of such lines and protocols of risk communication, elevation and cooperative handling should be part of risk management plans formulated at all primary organizational levels and for all major activities, regardless of the domain within which they are executed.

³ This is analogous to the systems engineering practice of defining a requirement’s verification protocol in tandem with requirement definition, rather than waiting until verification itself when there might be a temptation to craft a protocol that yields a positive result regardless of what was actually produced.

2.2.6.4 Risk Feedback and Communication Concerning Risk

An important and often overlooked aspect of risk management integration is the non-attributional communication and feedback concerning risk and its handling, as recognized by the previously cited RMTT report [9] in its recommendation for “Encouragement of additional Risk Dialogue.”

All of the risk types identified earlier in Section 2.1.6 can flow down or propagate upward and across the organization and are of potential concern to all organizational entities within it. Correspondingly, the risk management function within the organization must also cut across the various organizational entities, making vertical and cross-organizational communication essential to maintaining full awareness and consistency in the execution of risk management activities and for the effective implementation of risk leadership objectives.

For risk management in the service of risk leadership, communications across the organization should occur not only upward and downward, but also horizontally, as depicted below in Figure 2-9. The figure shows a hypothetical microcosm of several entities, designated as *superordinate*, *parallel*, and *subordinate*, surrounding a *root* entity. In the interaction between the management of these organizational entities and the corresponding management of interrelated or cross-cutting risks, the entities allocate their objectives, constraints, requirements, and risk postures downward to their subordinate organizational entities. These then become inputs to the identification, evaluation, and management of risks and opportunities at the subordinate level. The information generated there through each entity’s risk management activities is communicated both upward to entities that are superordinate to them and laterally to entities that are parallel with them and potentially affected by common and cross-cutting risks.




Figure 2-9. Top-Level View of the Flow of Risk-Relevant Information to and From a Root Entity

2.3 References for Chapter 2

1. NASA Procedural Requirements, NPR 8000.4C, Agency Risk Management Procedural Requirements. April 2022.
2. NASA Policy Directive, NPD 1001.0D, 2022 NASA Strategic Plan. March 2022.
3. NASA Special Publication, NASA/SP-2014-615, Organizational Risk and Opportunity Management: Concepts and Processes for NASA's Consideration. November 2016.
4. Bacon, F. The Essays. 1612.
5. NASA Policy Directive, NPD 1000.0C, NASA Governance and Strategic Management Handbook. January 2020.
6. NASA Policy Directive, NPD 1000.5C, Policy for NASA Acquisition - Updated w/Change 2. July 2020.
7. NASA Policy Directive, NPD 8700.1F, NASA Policy for Safety and Mission Success. July 2022.
8. NASA Policy Statement, NPS 1001.105, NASA Chief Acquisition Officer's Intent, June 18, 2024.
9. NASA Internal Report, Risk Management Tiger Team Report. September 07, 2023.

10. Ellsberg, D., Risk, Ambiguity, and the Savage Axioms, Quarterly Journal of Economics, Vol. 75, No. 4. November 1961.
11. Kaplan, R., and Mikes, A., Risk Management – the Revealing Hand, Harvard Business School Working Paper 16-102. 2016.
12. NASA Special Publication, NASA/SP-2010-576, NASA Risk-Informed Decision Making Handbook, Version 1.0. April 2010.
13. NASA Procedural Requirements, NPR 7120.5F, NASA Space Flight Program and Project Management Requirements w/Change 3. August 2021.
14. NASA Procedural Requirements, NPR 7123.1D, NASA Systems Engineering Processes and Requirements w/Change 1. July 2023.

3 Risk Models, Analysis, and Decision Concepts

Risk is a term with a variety of meanings and is subject to many interpretations. Common definitions of risk include the possibility of loss or injury [1]; someone or something that creates a hazard [1]; a situation involving exposure to danger [2]; and an intentional interaction with uncertainty [3]. Consequently, in the absence of an unambiguous shared understanding of the term and the concepts that surround it, there is a possibility that those who within an organization are tasked with identifying, understanding, managing, and communicating risk will to some extent talk past each other and/or work at cross purposes, therefore possibly making their organization ineffective in achieving its intended goals of performance, safety, reliability, and affordability.

The purpose of this chapter is to present and explain in consistent fashion key definitions and concepts relating to risk and risk management that can be used by organizations within NASA or elsewhere as a basis for establishing, structuring, or refining their risk management processes, tools, and activities. The concepts presented herein reflect much of the thinking that went into several earlier NASA publications addressing topics directly or indirectly related to risk management [4-8], as well as the more recent thinking and developments more specifically conceived to address the requirements of the most recent version of NPR 8000.4 [9] and the recommendations of Ref.[10].

In order to be precise about the concepts presented in this handbook, quantitative mathematical illustrations, involving mathematical functions, coordinate axes, and the like, are routinely used in this chapter. This is not intended to imply that risk management is necessarily only quantitative. Rather, it is intended to communicate the concepts with sufficient clarity that they can be adapted to any level of rigor of risk management, whether in quantitative or corresponding qualitative terms. For example, a quantitatively defined probability distribution function may have a qualitative equivalent, which assigns different degrees of likelihood to a finite set of qualitatively defined consequences of a given risk-relevant event. Regardless of whether risk scenarios and parameters are addressed via quantitative or qualitative formulations, the goal should remain that of supporting the establishment and execution of risk management processes that are internally coherent, and effectively integrated into the project / activity management activities they support.

3.1 Foundational Risk Concepts

The definition of risk that is operationally applicable in the context of all NASA activities is set by Ref. [9]. This definition has been presented and discussed in Section 2.1.1 and is intended to make the achievement or non-achievement of the stated objectives of a project or activity the primary focus of risk and opportunity management. Other sections of Chapter 2, more specifically Sections 2.1.2 through 0, have elaborated further on this fundamental definition and concept, with discussion of the importance of understanding the flow-down of organizational objectives and their allocation to organizational units for execution in different activity domains, the types of risks that may arise in the execution contexts of different domains, and the importance of considering opportunity alongside risk if an overall well-balanced risk posture is to be pursued and achieved.

The material presented in the remainder of this chapter proceeds from the foundational ODRM concepts presented in Chapter 2 and moves the discussion of the risk management framework into the more detailed technical dimension of the operational characterization of risk: both as an

aggregate in relation to each activity objective that is affected by it, and as individual scenarios that arise from the possible occurrence of a “departure” from the baseline of a planned or intended activity execution.

Besides the operational characterization and modeling of different types of risk, the other set of foundational concepts discussed in the chapter concerns the relationship between qualitative declarations of risk posture and the operational, and where possible quantitative, definitions of risk tolerance levels and risk acceptance criteria. This second set of foundational risk framework concepts represents a new development specifically conceived and presented in this handbook as the fundamental instrument of implementation of the risk leadership principles formulated and asserted by Ref. [11].

3.2 Characterization of Risk

This section examines the different basic perspectives from which risk may be considered and accordingly characterized and modeled. Chapter 2 has previously discussed “types of risk,” whereby a classification was considered according to various possible criteria, e.g., by risk source, type of impact, or domain of affected activity. Such a discussion was presented primarily in recognition of the fact that risk has traditionally been labeled in one way or another depending on the context of the discussion and the primary focus of the activity of the classifying entity. In this and following sections of the present chapter, however, “the type of risk” is considered from the technical and practical point of view of how it can be characterized, modeled and assessed, e.g., based on whether it is viewed and analyzed as a single scenario with specific causes and consequences, or as the “aggregate” of multiple contributions to a possible underperformance with respect to a defined objective, or even as an undefined possibility of underestimation of any such aggregate.

3.2.1 Aggregate Risk to an Objective

A definition of Aggregate Risk (AR) has been provided earlier in Section 2.2.2. Such a definition underscores the fact that AR is generally and typically composed of an ensemble of distinct contributions, that may combine in different ways to result in the overall AR profile. Indeed, going back to the basic and more general definition of risk provided in Section 2.1.1, which refers to the “potential for shortfalls” to which a given objective is exposed, it is natural to conclude that such a potential may typically be associated with a multitude of things that can go wrong in the course of an organization’s efforts to achieve an objective of concern. Similarly, there are also a multitude of ways that things can go right, each of which has its own specific probability of occurrence and its own successful outcome by which the objective can be achieved. Collectively, the totality of ways that things can go wrong or right define a probability density function (pdf) over the space of possible outcomes, as shown in Figure 3-1. In the figure, the outcomes to the left of Objective A represent shortfalls in performance relative to Objective A, and the area under the pdf in that region quantifies the probability that a shortfall will occur. In other words, the red shaded area of the figure represents the risk to Objective A, which is the cumulative result of all the ways that things can go wrong with respect to its achievement.

This picture of risk as a probability of failing to meet an objective is consistent with the fact that NASA’s objectives are in most cases expressed in binary form, in that they are associated with distinct thresholds of performance that must be met⁴. Such a threshold model of objectives is especially applicable to requirements, which are formally verified as met or not met according to their established verification protocols. The ability to deliver a payload to a specified orbit or to meet a launch window are additional examples of objectives expressed in binary form.




Figure 3-1. Top-Level Anatomy of Risk

Of course, a binary perspective of risk limited to viewing risk as the probability of falling short of an objective of concern is not the whole story about what can go wrong in pursuit of an objective. The magnitude of the shortfall can also be relevant, especially when the objective is a so-called “soft” objective where there is some arbitrariness to where the “line in the sand” is drawn. For example, for a \$100M program, a shortfall of \$1M is likely to be perceived as more successful than a shortfall of \$50M. Similarly, given a reliability objective of 0.99, an achieved reliability of 0.985 is likely to be perceived as more successful than a reliability of 0.75. Similarly, the extra margin by which an objective is achieved can also be relevant. For example, achievement of a mass-to-orbit objective with large margins can create opportunities to meet additional or upgraded objectives that were not previously considered practical to pursue.

3.2.2 Individual Risk Scenarios

The representation provided in Figure 3-1 is essentially complete as a high-level theoretical expression of risk. However, from the point of view of *managing* risk it is necessary to understand *how* things can go wrong, so that the organization can intervene to the extent practicable, both to

⁴ It is to be noted, however, that the conceptual model of objectives-related risk as a probability of not achieving a given desired level of performance is still valid when the pursued objectives take a distributed, rather than threshold, form. For example, if the stated objective is to be anywhere in a given range of performance, risk can then be expressed by the probability that the level of performance be outside of that range. Even when an objective is expressed in more open-ended terms, such as “maximize revenue,” a probabilistic quantification of risk can still be obtained by first developing a model of what the materially possible and achievable range of revenue may be, then generating a probability distribution for the activity outcomes in that range, and finally estimating the probability of failing to achieve the possible maximum revenue of that range.

prevent things from going wrong and to mitigate the shortfall should they go wrong. To do this, a given risk must be characterized in a way that reveals what organizational and project/program objective(s) may be threatened, what the causes of such an outcome may be, how likely those causes are to be operative, and how those causes can propagate through an organization and/or its work products to disrupt its achievement of its objectives. To this end, NPR 8000.4 elaborates on its definition of risk by *operationally characterizing* risk as a set of triplets:

- a. The scenario(s) leading to degraded performance with respect to one or more performance measures,
- b. Their likelihoods, and
- c. Their consequences (in terms of the magnitude of performance degradation).⁵

Pertinent performance measures may pertain to program or project objectives (e.g., safety, technical performance, cybersecurity, physical security, cost, and schedule metrics), institutional activities (e.g., staffing, facility availability, and supply chain metrics), and/or enterprise activities (e.g., strategic planning, compliance, and acquisition metrics). The concept of a risk triplet is illustrated in Figure 3-2. Each scenario in this set of risk triplets is considered to be an *individual risk scenario*, contributing to, and representing a portion of, the aggregate risk depicted in Figure 3-1. Thus, in practical terms, a representation of aggregate risk such as the one shown in Figure 3-1 is often produced by identifying an “as complete as possible” list of individual risk scenario contributors and calculating by appropriate analytical means their combined contribution to the risk affecting the objective of concern.




Figure 3-2. Risk Operationally Characterized as a Set of Risk Triplets

For completeness of information and to avoid any misunderstandings, the reader must also be made aware that in general risk management literature and guidelines the term *individual risk* is often used to refer to what we have defined here to be an *individual risk scenario*. In the remainder of this handbook the latter term will prevailingly be used, but the two terms should in any case be considered as being equivalent and interchangeable in their meaning and use.

⁵ NPR 8000.4 further specifies that uncertainties are included in the evaluation of likelihoods and identification of scenarios.

In realistic terms, even a scenario that is identified and analyzed as an individual risk may be a partial aggregate of “sub-scenarios” that could themselves be considered as being individual risks at a lower level of modeling indenture: for example, a scenario defined as “mission fails because of launch vehicle attitude control system (ACS) failure” could be broken down into the sub-scenarios “mission fails because of launch vehicle ACS computer failure,” “mission fails because of launch vehicle ACS sensor failure,” “mission fails because of launch vehicle ACS actuator failure,” etc. What is defined in practical terms to constitute an “individual risk scenario / individual risk,” including what constitutes its causes, may therefore be determined by considerations of modeling and analytical convenience, as suggested by the primary intent of making that risk identifiable and explicitly addressable by the risk management function of the organization that is potentially impacted by it.

Based on the above considerations it is possible to provide a definition and characterization of an **individual risk scenario** in the practical operational terms that are necessary for its proper identification, modeling and quantification, both for the purpose of its own characterization as a distinct component of the risk “spectrum” potentially affecting an activity objective and the related performance measure(s), and as a contributor to the overall expression of aggregate risk to that objective. This is provided in the blue-box below. It may be further noted, contrasting the definition of individual risk scenario to that of aggregate risk, that the former implies that an individual risk scenario formulation is a means of characterizing risk in terms of the chain of key events or conditions from and through which an individual risk proceeds and produces its consequences, whereas the latter pertains to the coalescence of multiple scenarios into the specific type of consequence that impacts an activity objective and its quantification by means of the corresponding performance measure(s).

Definition and Characterization of Individual Risk Scenarios

Individual Risk Scenario. A sequence of events or combination of such sequences, originated by an event or condition followed by other events or conditions, which are judged to be unique and defining with respect to how the consequences of concern are produced and impact one or more activity objectives.

An Individual Risk Scenario may be characterized as a “triplet,” as defined and referred to in the general characterization of risk provided by Ref. [9]:

“Risk is operationally characterized as a set of triplets:

- *The scenario(s) leading to degraded performance with respect to one or more performance measures (e.g., scenarios leading to injury, fatality, destruction of key assets; scenarios leading to exceedance of mass limits; scenarios leading to cost overruns; scenarios leading to schedule slippage).*
- *The likelihood(s) (qualitative or quantitative) of those scenarios.*
- *The consequence(s) (qualitative or quantitative severity of the performance degradation) that would result if those scenarios were to occur.*

Uncertainties are included in the evaluation of likelihoods and identification of scenarios.”

As mentioned above, there typically are multiple paths by which things that can go wrong in the course of an organization's efforts to achieve its objectives. In principle, this multitude of possible event sequences is infinitely divisible based on finer and finer distinctions between what are otherwise similar scenarios. So, as a practical matter and reflecting what was discussed above, a given individual risk scenario will generally envelope classes of more finely defined scenarios that share similar causes, propagation pathways, and consequences. The likelihood of the enveloping individual risk scenario will then be the likelihood that *any one or more* scenarios within this envelope will occur. In other words, individual risk scenarios are proxies that are constructed to represent a bounded multitude of underlying undesired possibilities. Construction of individual risk scenarios is an analytical activity subject to human judgement, so different analysts may define somewhat different sets of individual risk scenarios in order to characterize the same aggregate risk.

Additionally, the characterization of risk in NPR 8000.4 explicitly recognizes that a scenario leading to a shortfall with respect to one objective may very well also lead to a shortfall with respect to another objective. In fact, in practice this is usually the case, due to the high degree of correlation among the corresponding performance measures. For example, project schedule risks and cost risks are typically highly correlated because schedule slippages typically result in additional costs. For crewed missions, mission safety risks and mission success risks are also typically highly correlated because catastrophic events can cause both the loss of crew and the loss of the mission. Therefore, individual risk scenarios can, and often do, entail shortfalls with respect to multiple objectives. The need to capture these correlations with appropriate models and metrics shows that risk management should not be stovepiped by objective. It also illustrates that unlike aggregate risk representations (see Section 3.2.1), individual risk scenarios are not domain specific. Instead, a single individual risk scenario can potentially threaten objectives in all of the above-mentioned domains (safety, supply chain, reputational, etc.). The development of individual risk scenarios is addressed in more detail in Chapter 5, which addresses the concept of a *risk statement* that is used as a standardized format for defining and communicating the key characteristics of an individual risk scenario and for identifying individual risk scenarios that are cross-cutting in their potential for impacting the objectives of multiple projects and of the organizational units responsible for their execution.

3.2.2.1 Characterizing an Individual Risk Scenario via a Risk Scenario Diagram

A risk scenario begins with an undesired “*departure event*” that represents a departure from the intended state or path of execution. Subsequent events that are relevant to the evolution of the scenario may (or may not) occur and may have either a mitigating or exacerbating effect on the scenario progression. Depending on this progression, the scenario can result in a reduction in performance with respect to the performance measures of interest. There is no hard constraint on what can constitute a departure event. Departure events can be technical or programmatic (e.g., a cost overrun, schedule slippage, safety mishap, health problem, malicious activity, environmental phenomenon, or failure to achieve a needed scientific or technological objective or success criterion). The resulting risk scenario can be illustrated using a risk scenario diagram (RSD), as

shown in Figure 3-3, which shows the various ways that that an individual risk scenario can propagate to result in degraded performance.⁶




Figure 3-3. Schematic of a Risk Scenario Diagram (RSD)

3.2.2.2 Characterizing the Likelihood of an Individual Risk Scenario

The likelihood that an individual risk scenario will be initiated is the likelihood of the departure event that defines the initial departure of the activity from its intended path. The likelihood that the departure will result in a specific end state that describes the possible consequences of concern and their severity depends on the probabilities of the subsequent pivotal events whose occurrence or non-occurrence leads to that end state. For example, referring to Figure 3-3, the probability that the activity in question will progress to “End State j” is the probability of the departure event, times the probability of Pivotal Event 1, times one minus the probability of Pivotal Event n:

$$P_{\text{End State j}} = P_{\text{DE}} \times P_{\text{Pivotal Event 1}} \times (1 - P_{\text{Pivotal Event n}})$$

Depending on the context and the need for robustness in decision-making, likelihoods can be treated either qualitatively or quantitatively. Methods for determining likelihoods can vary widely depending on the context, but include expert elicitation, empirical (via testing or historical record), fault tree analysis (FTA), and phenomenological modeling.

3.2.2.3 Characterizing the Consequences of an Individual Risk Scenario in Terms of Its Effects on Performance

It can be tempting to characterize the consequences of an individual risk scenario in terms of the worst case or most dramatic event in the scenario, such as explosion, budget cut, technology development failure, or program cancellation. However, in keeping with the definition of risk as a shortfall with respect to an explicitly stated objective, the consequences of an individual risk scenario should be characterized by performance in terms of the performance measures used to

⁶ An RSD is an event sequence diagram (ESD) that illustrates the possible propagation pathways of a departure event.

define risk per Figure 3-1. This means that each end state in the RSD should be expressed in terms of performance across all the objectives and associated performance measures of the organization in question. In general, different end states will result in different levels of performance, but the set of performance measures used to describe that performance should be the same for all end states, as well as for all the different individual risk scenarios to which the organization is exposed.

RSDs are typically supported by performance models that generate performance measures for each path through the RSD, conditioned on the occurrence of the departure event and the occurrence or non-occurrence of the pivotal events along the path. Like likelihood modeling, the character of performance models can vary in rigor, ranging from expert opinion to phenomenological modeling. In all cases, however, performance must be modeled rigorously enough to compare it to the objectives and determine whether or not it falls short. Performance modeling is addressed further in Chapters 4 and 5.

3.2.2.4 Characterizing Individual Risk Scenarios Produced by Hostile Agents

NPR 8000.4 includes provisions to address the special case of risk associated with intentional actions by hostile agents, within the general context of physical security and cybersecurity. An appendix of the NPR [9, Appendix C] specifically illustrates the form that the characterization of this type of risk may take. Figure 3-4 below, reproduced as a combination of the Figures C1 and C2 in that reference, show two summary-form RSDs that represents, respectively, an accidental risk scenario (Figure 3-4a) and an intentionally-triggered one (Figure 3-4b).




Figure 3-4. RSD Representation of “Standard” vs. “Hostile Agent” Risk Scenarios

The two RSDs in Figure 3-4 are boiled-down equivalents of the more detailed form of scenario representation illustrated earlier in Figure 3-3. More specifically: the “Initiating Conditions”

correspond to what Figure 3-3 refers to as the “Departure Event,” the “Impact Event” is one particular “Pivotal Event,” and the “Consequences” are the “End States” of primary interest.

The key take-away from Figure 3-4 is that the two RSDs appear to be essentially the same in their structure, but differ in the way the respective Initiating Condition (or Departure Event in the language of Figure 3-3) should be viewed and treated. In the case of a “standard” individual risk scenario, the Initiating Condition / Departure Event is something that occurs randomly, with a frequency that is relatively stable over time and that therefore can be treated and assessed in probabilistic terms. In the case of a risk initiated by an adversarial and intentional action, on the contrary, the likelihood of the Initiating Condition is driven by dynamically changing conditions (e.g., political, or even openly conflictual), therefore it is more appropriate to assume the condition, if at all credible, as being present (i.e., having probability = 1) and evaluate risk in terms of the conditional probability of the Consequences / End-States, given the presence of such a condition. In the representation of Figure 3-4b, the probability of the undesirable consequences is equal to $P(IE/IC)$, i.e., to the conditional probability of the Impact Event, given the Initiating Condition. Correspondingly referring to Figure 3-3 and the example of probabilistic evaluation given in Section 3.2.2.2, if that risk scenario depiction were for a hostile agent risk, the Departure Event would be the intentional act assumed as occurring and upon which all other events in the RSD are conditional. In such a situation, the probability of “End State j” would be evaluated as a conditional probability and therefore calculated as:

$$P_{\text{End State } j} = P_{\text{Pivotal Event 1}} \times (1 - P_{\text{Pivotal Event } n})$$

From a practical point of view, the conditional probability values that appear in the hostile agent risk scenario depictions discussed here are an expression of how strong and robust the defenses of the system being considered are against the hostile actions of the agent: the stronger the defenses, the smaller the conditional probability of an undesired consequence.

The dynamic and adversarial nature of security-related risks warrants specialized treatment of both cybersecurity and physical security in NASA’s policies and procedures. NPD 2810.1, *NASA Information Security Policy* [12], and NPR 2810.1, *Security of Information and Information Systems* [13], establish NASA’s information security policy and the information security requirements for NASA’s Information Security Program. In particular, NPR 2810.1 aligns roles and responsibilities of IT security personnel to National Institute of Standards and Technology (NIST) Special Publication (SP) 800-37, *Risk Management Framework (RMF) for Information Systems and Organizations: A System Life Cycle Approach for Security and Privacy* [14] and *NIST SP 800-53, Security and Privacy Controls for Information Systems and Organizations* [15]. The RMF provides a disciplined, structured, and flexible process for managing security and privacy risk that includes information security categorization; control selection, implementation, and assessment; system and common control authorizations; and continuous monitoring. The RMF includes activities to prepare organizations to execute the framework at appropriate risk management levels. The RMF also promotes near real-time risk management and ongoing information system and common control authorization through the implementation of continuous monitoring processes; provides senior leaders and executives with the necessary information to make efficient, cost-effective, risk management decisions about the systems supporting their

missions and business functions; and incorporates security and privacy into the system development life cycle.

Physical security is addressed in NASA's policies and procedures by NPD 1600.2, *NASA Security Policy* [16], NPR 1600.1, *NASA Security Program Procedural Requirements* [17], and NPR 1620.3, *Physical Security Requirements for NASA Facilities and Property* [18]. These documents define roles and responsibilities related to physical security and require the establishment of an Enterprise Security Risk Management Program that enhances operational readiness and mission success by providing security support to program/projects throughout the life cycle of a system or activity that is commensurate with the risk and helps ensure mission critical information, technologies, and/or assets are appropriately protected.

3.2.2.5 *Evaluating Likelihood and Consequence Uncertainty*

The preceding discussion of likelihood deals with the various ways that an activity can depart from its intended path and result in degraded performance. It is a matter of chance whether or not the activity does in fact depart from its intended path, and it is also a matter of chance how the departure propagates, and which end state is ultimately realized. This element of chance is known as *aleatory uncertainty*, and describes the randomness of circumstance that is inherent in the activity.

There is another type of uncertainty that also should be taken into account in the characterization of risk, namely the uncertainty associated with imperfect knowledge about the activity itself. This type of uncertainty is known as *epistemic uncertainty* and can result in the inability to specify with certainty what the likelihoods in the risk scenario actually are, or what the levels of performance will be achieved for each of the end states. Epistemic uncertainty is addressed in more detail in Chapters 4 and 5.

3.2.3 Forms and Metrics for Individual Risk Scenario Contributions to Aggregate Risk

Having introduced in this and preceding chapter the concepts and definitions of Aggregate Risk (AR) and Individual Risk Scenarios (IRSs), and also having discussed the basic forms of characterization and modeling used to analyze and assess these related but different manifestations of risk, this section considers other key aspects of their interrelationship that are significant for a correct evaluation of how an ensemble of IRSs contributes to an AR of concern.

A mode of IRS contribution to an AR that is commonly discussed and assessed in the technical literature occurs when all the IRSs in the AR-contributing set, if realized, produce the same undesired consequence. A well investigated example of this type of aggregation is the contribution of individual failure modes to the failure of a given system, for which a representation of the resulting AR is rather straightforwardly given by the pair of parameters that represent, respectively:

- a. System Failure (SF) as the consequence dimension of the AR, and
- b. p_{SF} , the probability that SF will occur.

If the occurrence of any IRS is statistically uncorrelated with the occurrence of any other IRS, and p_{IRS_i} is the probability of occurrence of the i -th IRS, p_{SF} can be calculated from the well-known “OR-gate” formula:

$$p_{SF} = 1 - \prod_{i=1}^N (1 - p_{IRS_i})$$

In more general types of ARs the key characteristic of the above specific case, i.e., that all the contributing IRSs carry the same consequence in both qualitative and material terms, no longer holds. In fact, in more general types of aggregation, such as those concerning the context of program or project risk in relation to a given activity objective, the contributing IRSs, if realized, do not necessarily carry the same consequences. Accordingly, in such situations the IRS undesired consequences are expressed in terms of a performance measure (PM) shortfall value (SV) that varies in absolute magnitude within a given possible range $[0, SV_{MAX}]$.

For such more general situations, which are discussed in some level of detail in [19], if a characterization of the resulting AR is desired in terms of a “classic” pair of probability vs. consequence parameters, this might be done by referring to the probability that any combination of IRSs capable of producing a PM shortfall will occur, and to some representative “point value” – such as a median or mean value – of the SV consequence parameter. If N is the number of such IRSs, p_{IRS_i} is the probability of the i -th IRS occurrence, and it can still be assumed that there is no significant degree of correlation among the IRS occurrences, then the probability portion of the two-parameter AR representation, p_{AR} , is then expressed by a formula analogous to the one provided above for the p_{SF} parameter, i.e.:

$$p_{AR} = 1 - \prod_{i=1}^N (1 - p_{IRS_i})$$

For the consequence side of an AR two-parameter representation, i.e., the median or mean values of the PM shortfall value, respectively indicated with the notations SV_{50} and SV_m , are calculated over the range of possible shortfalls SV_i . For the reader’s benefit, we recall below how this can be done according to the definition of such parameters.

Given a set $\{IRS_i\}$ ($i=\{1,2,...,N\}$) of IRSs that produce associated shortfalls SV_i with probability values p_{IRS_i} , the median shortfall S_{50} is a value of the shortfall in the range $[0, SV_{MAX}]$ for which the sum of the probabilities of the SV_i s that are smaller than SV_{50} is approximately⁷ equal to the sum of the probabilities of the SV_i s that are greater. This also implies that either sum is approximately equal to 0.5. For the same set, the mean value can be defined and calculated as:

⁷ The term “approximately” is used here because of the discrete nature of the IRS set and associated probability mass distribution. In this type of distribution, it is generally not possible to identify a median value such that the sum of the probabilities of distribution items respectively smaller and greater than the median is exactly the same and equal to 0.5.

$$SV_m = \sum_{i=1}^N p_{IRSi} SV_i$$

While the above may be a valid AR characterization for some cases of aggregation, two situations that are often encountered in practical programmatic contexts need to be given consideration, as the AR characterizations provided above may be unsuitable or even misleading for them. These situations are discussed in the below sub-sections.

3.2.3.1 Characterization of Aggregate Risk Produced by Mutually-Exclusive IRSs

The first case of practical interest is one where the consequences of the IRS occurrences are again different in shortfall value as in the generic case discussed above, but are also mutually exclusive of one another. For example, IRSs representing different types of possible faults may cause a planetary lander intended to hit a target spot on a planet surface to miss that target by different amounts of physical distance. The spectrum of IRS negative consequences for such scenarios would thus be represented by the different degrees of lost-science value corresponding to the magnitude of the target-miss. However, because of the one-shot nature of the mission, each of the possible negative outcomes hypothetically possible for it is mutually exclusive with any of the others.

In this situation of mutual exclusivity of the IRSs that may occur, the AR characterization formulas provided earlier still hold if in them the probability terms p_{IRSi} , which in most cases are assessed without taking into account the IRS mutual-exclusivity feature, are replaced by probability terms pex_{IRSi} that do account for it. For small probabilities, the relation between the two is defined and calculated by means of the formula below:

$$pex_{IRSi} = p_{IRSi} \prod_{j \neq i} (1 - p_{IRSj})$$

Once the pex_{IRSi} terms are calculated, a two-parameter AR characterization more appropriate for this type of situation is then definable as:

$$p_{AR} = \sum_{i=1}^N pex_{IRSi}$$

$$SV_m = \sum_{i=1}^N pex_{IRSi} SV_i$$

Of course, if it were possible to directly quantify the pex_{IRSi} terms, i.e. assess IRS probability values as already reflective of the mutual-exclusivity of the possible outcomes, then the formula shown above that maps p_{IRSi} values into pex_{IRSi} ones would not need to be utilized.

3.2.3.2 Characterization of Aggregate Risk Produced by IRSs with Cumulative Effects

The second subcase of interest includes all the situations in which multiple occurrences of the AR-contributing IRSs can be realized, and the corresponding consequences are cumulative. An example of this situation with which any manager is familiar is when the PM of concern is cost or schedule, and the contributing IRSs are the various scenarios whereby a cost or schedule overrun may occur and add itself to those that already may have occurred.

The calculation of the probability and consequence parameters that characterize the AR for this situation is conceptually straightforward, but potentially complex in mathematical terms, depending on how many contributing IRSs are to be considered, and whether the magnitude of their impact on the PM of interest can be assumed to vary continuously in a range, or is approximated by a finite set of discrete values. Examples of this type of AR derivations are presented in Appendix D. In general, however, when situations of this kind are encountered it is recommended that the analytical processes needed to estimate any risk parameter of interest be entrusted to technical staff with expertise in risk and probability modeling. In complex situations, computer assisted Monte Carlo simulations may have to be employed as the most practical means of analysis and assessment.

3.2.4 Unknown and/or Underappreciated Risk

There are many well established methods that can be used to identify individual risk scenarios, and as is elaborated on further in this handbook, the identification of individual risk scenarios is a fundamental risk management activity. However, like most discovery activities, risk identification is vulnerable to *incompleteness*, which means that there is never any guarantee that an organization's individual risk scenarios will collectively address the totality of risk to which the organization is exposed. An organization can and should minimize incompleteness by applying formalism, structure, and expertise to the task of risk identification, but it can never be absolutely certain that it can completely eliminate it by diligently applying such means. An organization like NASA, which operates at the cutting edge of technological innovation and human accomplishment, executes missions for which there exists a very limited basis of experience to draw upon when identifying individual risk scenarios. Thus, the question of adequate comprehensiveness is very pertinent to many NASA activities. In fact, even organizations with long operating histories, which have learned from real-world experience what their risks typically are and how to control them, are still vulnerable to individual risks having frequencies of occurrence at or below the low end of their (cumulative) experience, or which reflect unknown or unaccounted-for changes in the organizations, their activities, or the environments in which they operate.

Therefore, characterization of risk based solely on any set of identified individual risk scenarios is systematically non-conservative, and decisions based on risk characterized solely in this manner tend to unknowingly expose an organization to excessive risk, setting it up for failure. The situation is conceptually illustrated in Figure 3-5, which shows the risk to notional Objective A (from Figure 3-1) partitioned into the *known risk*, i.e., the risk that is collectively accounted for by the organization's individual risk scenarios, and the *unknown and/or underappreciated (U/U) risk* that may elude characterization in the form of specific scenarios but is nevertheless just as real as the

known risk. The figure conceptually shows a quantitative depiction of risk, however the potential existence of U/U risk should be recognized independently of whether it may be conceptualized in quantitative or qualitative form. Thus, it is important, when managing risk, to recognize the potential for U/U risk and include it in the overall management of risk. The assessment and management of U/U risk is particularly challenging, since by definition it cannot be characterized and examined in terms of well-defined event sequences. However, a perspective on its presence can still be gained by having an appreciation of its historical magnitude in the types of activities in which an organization is presently engaged, and how factors such as complexity, novelty, schedule pressure, and adherence to quality principles can mitigate or exacerbate it [7].




Figure 3-5. *Anatomy of Risk: Known Risk and U/U Risk*

Novel and complex uses of artificial intelligence (AI), including machine learning (ML), can contribute to risks that are not easily captured by traditional analysis of known failure modes. Additionally, AI/ML models can be non-deterministic and highly complex resulting in a high potential for U/U risk. Additional information on AI risk management frameworks can be found in the NIST AI Risk Management Framework [20] and the NIST AI Risk Management Framework Playbook [21].

3.2.4.1 *Leading Indicators of U/U Risk*

In a given situation, factors such as those just mentioned as producers of U/U risk can be usually quantified. They can therefore be thought of as *leading indicators* of U/U risk, and their current values and trends can be used to scope the potential magnitude of the U/U risk.

The usefulness of leading indicators to an enterprise like NASA, or any other technical enterprise, is similar to the usefulness of economic leading indicators to the financial world. In the latter case, indices that measure consumer confidence and other leading economic indicators have been used with reasonable success to assess the potential for future economic downturns and even the potential magnitude of the downturn. Their efficacy is predicated on the idea that historical experience can be used to develop correlations between the indicators and the outcomes,

recognizing that although an identified correlation may still correspond to a considerable amount of scatter in the underlying data, the mean trends of the correlations are nevertheless easily recognizable and can still provide useful information. The same is true in the technical world, where leading indicators such as those listed in Appendix B can (and have) been correlated with the potential for accidents resulting in death or injury, performance shortfalls, cost overruns, and schedule slippages [22, 23].

The effective use of leading indicators is a cornerstone of the proactive response dimension of RM discussed in Section 2.2. The challenge of proactive RM is to select the right indicators (a task for which the judgment of experts is helpful), develop data-based correlations between the values of the indicators and the magnitude of the shortfall, observe the amount of scatter surrounding the correlations, and apply the results to estimate a margin by which the probability distribution for each performance measure owing to known risks should be augmented to account for U/U risk. Example leading indicators of U/U risk are discussed in Appendix B. Ref. [23] also provides an extensive discussion of the leading indicators that may be correlated with specific types of risk impacts, and potential shortfalls in specific performance dimensions. Examples of potential leading indicators of AI related U/U risk include algorithm complexity, training data quality/bias, and lack of extensive operational heritage.

3.2.4.2 Characterizing U/U risk via Implied Risk Scenarios

U/U risk can be operationally characterized via implied risk scenarios. For instance, in investigations following various catastrophic accidents spanning space, nuclear, chemical, and other enterprises, it has been found that there have been known leading indicators of potential risks that were overlooked, including deficiencies in the management culture, insufficient attention to defense in depth, and lack of human factors consideration in the design of diagnostics. Because scenarios that could result from these deficiencies were not easily envisioned and depicted, but only generically implied, effective mitigation was not employed. An implied risk scenario, thus, is one for which the specific chain of events leading to an undesired end state and corresponding loss is not known beforehand, although there is an historically based correlation between one or more leading indicators and the occurrence of loss in one or more performance measures. Figure 3-6, reproduced from [24] and addressing robotic spaceflight projects at JPL, provides data that could be used to estimate the potential for U/U risk due to a mismatch between design complexity and funding. The data show that failed and impaired programs tend to be relatively underfunded for their complexities. The data could be used to develop a probabilistic model of program outcome as a function of its funding and design complexity, which in turn could be used to define implied “cost-complexity mismatch” risk scenarios for new programs whose funding falls short of that which has historically resulted in success.

3.2.5 Further Observations on Accidental vs. Adversarial Scenarios

This section adds some additional considerations on the subject of risk scenarios resulting from adversarial actions, which was introduced in Section 3.2.2.4 from the perspective of their modeling and characterization via RSDs. Sequences of events initiated by accidental occurrences, random failures, unplanned delays, and unexpected costs are examples of unintended scenarios that can lead to undesirable consequences. These scenarios may be defined or implied, in accordance with

the definitions for defined and implied scenarios presented in Section 3.2.3.




Figure 3-6. Robotic Spaceflight Project Outcome as a Function of Cost and Design Complexity

Intentional acts that breach physical security or cybersecurity boundaries can also result in sequences of events that lead to undesirable consequences. If the type of adversary that initiates the sequence is identified, the associated scenarios are a class of defined scenarios, rather than implied scenarios, because the plausible events that can unfold can be deduced ahead-of-time. For example, a collection of possible scenarios might be identified by exploring the set of plausible moves by the adversary and countermoves by the defender in much the same manner as a chess-playing algorithm or a war-gaming algorithm explores the set of possible moves and countermoves within its game space.

Cybersecurity scenarios involve threats to and vulnerabilities of cyber-assets – i.e., data, information, control systems, and or information systems stored and/or implemented in digital form – and any related consequences caused by or resulting from unauthorized access, use, disclosure, degradation, disruption, modification, or destruction of information or information systems cyber-assets, including such related consequences caused by an act of terrorism.

Physical security scenarios are similar in character but involve direct attacks on physical, functional, control, information, and human assets capable of causing damage, or of resulting in unauthorized possession or control of such assets. Program, project, mission, and institutional activity security includes interactive aspects of cybersecurity and physical security, as well as aspects of either type of security that may not have been explicitly recognized in earlier traditional definitions, such as the security of computer and network-based control systems. As such, it requires the application of both cybersecurity and physical security provisions and protections.

Given the variety of attack scenarios that an adversary can choose from and the fact that most adversaries favor attacks that are unexpected and thus more likely to succeed, the main difficulty faced by the risk analyst with respect to characterizing intentional scenarios is of predicting the absolute likelihood of the attack. For this reason, the likelihood component in the risk triplet (Section 3.2.2) is better framed for intentional scenarios as a conditional probability rather than an absolute probability. By dealing with conditional probabilities, the analyst presumes that an attack

has occurred, and focuses on the likelihood that the assailant is successful given the attack. The risk triplet, then, consists of the scenario definition, inclusive of the definition of the type of attack, the conditional likelihood of success (given that type of attack), and the resulting consequence(s). The attack likelihood, if considered at all, is better dealt with separately, as something that may change at any time, according to such volatile factors as international relations, political contexts, and the like.

3.2.6 The Organization-Specific Risk Model

A full treatment of risk accounts for all sources, whether explicitly identified and characterized as individual risk scenarios, or more generically inferred from past experience. The risk to which each of the organization objectives are exposed, is addressed not only via the identification and analysis of the individual risks that contribute to it (i.e., the known risk) but also on an assessment of the potential for risk due to unidentified and/or underappreciated causes (i.e., the U/U risk). This totality of risk information constitutes the organization's *risk model*, i.e., the authoritative representation of the risks faced by the organization, developed to a level of detail that enables it to effectively inform decision-making. Figure 3-7 notionally illustrates the concept of risk model for an organization.




Figure 3-7. The Risk Model (of Objective B)

The integration of individual risk scenarios into the risk model can be accomplished in terms of the corresponding RSDs by linking their departure events to a common success path from which

the scenarios depart, as shown conceptually in Figure 3-8.⁸ Each end state of this linked scenario diagram has a likelihood of occurrence that is indicated by the height of the bar on the corresponding performance measure bar chart, and a level of performance that is indicated by the placement of the bar along the x-axis of the chart. Levels of performance that meet the objective are colored green, and those that fall short of the objective are colored red. At the bottom of the figure, the bars are notionally stacked on a common bar chart that shows the full distribution of end states representing the intended activity outcome along with the identified contributing scenarios. This distribution of *known risk* is then augmented by some margin to account for *U/U risk*, using methods that are not illustrated in the figure but are discussed in Section 4.7.3.4 and Appendix H. The result is a performance measure distribution that accounts for all sources of risk, both known and unknown/underappreciated, consistent with Figure 3-5.

Additional detailed discussion of risk modeling can be found in [25].

3.3 Organizational Risk Posture and Tolerance

Up to this point we have focused on the objective characteristics of risk as described by an organization's risk model. In this section we address the organization's overall *posture* towards the risks it is exposed to, and the specific thresholds and associated tolerance levels the organization may set in each of the relevant performance measure dimensions, to ultimately determine whether such risks are acceptable (and if so, by whom), or whether risk management actions should be taken to reduce or better understand them. Further technical considerations on risk evaluation and decision-making can be found in Appendix C.

3.3.1 Definition and Application of Risk Tolerance Levels

The expression of *level of risk tolerance* in quantitative probabilistic terms follows from the definition of Risk Tolerance provided by NPR 8000.4, which has been reproduced in the "blue box" in Section 2.2.3.1. That definition makes it clear that the acceptability or non-acceptability of a risk with respect to a stated objective should be decided on the basis of an identified limit for the probability that the objective be met or not.

Whether an objective is met or not can be expressed in terms of *performance markers (PMKs)*, i.e., specified thresholds in the dimension(s) of the performance measure(s) that represent the degree of achievement of the objective of concern. Chapter 4 further discusses performance markers as part of the RIDM process, but it can be more generally stated here that performance markers are customarily defined, even in contractual or interorganizational formal contexts, as either *performance constraints*, *performance requirements*, or *performance goals*. In this handbook, the term "*performance markers*" is used to refer to both targets of performance that are provisionally identified for analytical purposes, as is the case in the context of a RIDM process carried out in the planning stages of a project or activity, and to the performance thresholds that are set and formally utilized in the execution stages of such a project or activity. The difference between one type of marker and another depends on how strictly the associated level of performance is intended to be pursued and achieved at execution time. For example, a stipulation

⁸ Figure 3-8 is conceptual only and does not illustrate complexities such as the possibility of both individual risk scenarios being realized, or of Individual Risk Scenario 1 being prevented due to realization of Individual Risk Scenario 2.




Figure 3-8. Aggregation of Risk within the Integrated Organizational Risk Model

between an *Acquirer* and a *Provider* organization may define a specific marker as a firm performance threshold against which no or minimal tolerance of violation is permitted, while another marker value is identified more flexibly as a desirable target level of performance for which a less than full achievement can be tolerated, if other stricter requirements are nevertheless met.

With regard to the above, it is also noted that although the performance markers that are formally set as targets of performance for the activity execution stages may not have exactly the same definitions and values as those provisionally used for pre-execution activity planning purposes, the relation between the two sets needs to be clear and well understood. In fact, if execution stage markers were to be chosen in such a way as to substantially differ and be inconsistent with the marker values used in the RIDM analyses carried out to risk-inform the planning stage decision processes, the conclusions of the latter would no longer be applicable and valid for the activity execution stages.

In the probabilistic model context of RM and RA, what has been introduced above translates into the following operational definition of *risk tolerance level*:

- ***Risk Tolerance Level (RTL)***: *The limit value set by an organization for the probability that the performance measure expressing the achievement of an organizational objective may miss an established performance marker.*

Risk tolerance levels are illustrated in Figure 3-9, which assumes that two performance markers have been defined – a *performance requirement* (PMK-R) and a *performance goal* (PMK-G), and correspondingly assigns separate risk tolerance levels to them, in the form of probability threshold values, expressing respectively lower tolerance for a violation of the requirement and higher tolerance for missing the goal. In the figure, these are identified as RTL-R, a limit probability that defines unacceptable risk, in the sense that a violation of the PMK-R requirement with probability greater than that probability value is unacceptable, and RTL-G, indicating that risk is acceptable if the probability of missing the PMK-G goal is smaller than RTL-G while the probability of missing the requirement is also smaller than RTL-R.

The above definition is further illustrated by Figure 3-10, which represents the estimated cumulative distribution function (CDF) of a hypothetical performance measure (PM). In practice, CDFs communicate risk tolerance levels more effectively than pdfs because the risk of not meeting a given level of performance can be read directly off the y-axis, rather than being represented by the area under the pdf in the region of shortfall. The relationship between pdfs and CDFs is illustrated in Figure 3-11. The remainder of this section discussion remains relative to the situation just described. Other situations are possible, both with respect to what *performance markers* may be used and with respect to how RTL may be defined. For example, in some cases only one marker may be defined, and more than one RTL value may be associated with that single marker to identify different levels of risk. More specifically, when only one marker is defined, two RTLs are needed to identify “green,” “yellow,” and “red” levels of risk. Appendix C addresses these alternative situations, and provides more detailed explanations of the technical concepts introduced in this section and of their interconnections.




Figure 3-9. Satisfaction of Risk Tolerance Levels Relative to a Performance Measure PDF




Figure 3-10. Satisfaction of Risk Tolerance Levels Relative to a Performance Measure CDF




Figure 3-11. The Relationship between PDFs and CDFs

If the projected CDF of the performance PM is as shown in the figure, it follows that the depicted situation satisfies the risk tolerance level set for the performance requirement PMK-R. This is because:

$$[P1 = p (PM < PMK-R)] < RTL-R$$

However, the risk tolerance level with respect to the performance goal PMK-G is not satisfied, because:

$$[P2 = p (PM < PMK-G)] > RTL-G$$

As further discussion given below indicates, a situation like the one depicted in Figure 3-10 would normally be viewed as one of “marginal risk” with respect to the objective represented by the performance measure PM.

3.3.2 Risk Classification Based on Risk Tolerance Levels

When dealing with the satisfaction or not of risk tolerance levels set for a combination of performance requirements and performance goals, the resulting determinations of risk “coloring” and possible response, i.e., acceptance without further action or different types of active intervention, can be made according to the classification criteria illustrated by Table 3-I.

Table 3-I. Classification of Risk Based on Satisfaction of Tolerance Levels

RTL-Rs for Performance Requirements (PMK-Rs)	RTL-Gs for Performance Goals (PMK-Gs)	Risk “Acceptability” Classification
Satisfied	Satisfied	Green / “Acceptable”
Satisfied	Not Satisfied	Yellow / “Marginal”
Not Satisfied	n/a	Red / “Unacceptable”

It should be noted that what is shown in the table represents an “in-principle” determination of acceptability, non-acceptability, or in-between “marginal” status. That is, at the end of the day, a risk may indeed be accepted or not by a responsible and accountable decision maker by taking into account its acceptability classification, as initially determined according to the established risk posture and resulting tolerance criteria. However, other factors may also play an important role in the ultimate decision, e.g., factors resulting from more specific risk management determinations, such as the feasibility and cost of risk mitigation measures. Such factors are discussed in the following chapters of this handbook.

Given all that has been presented thus far regarding the definition and setting of risk tolerance levels, some important implications follow and need also to be discussed:

- A. The definition of an RTL is equally applicable to cases in which a performance objective is expressed over a possible continuum of outcomes (as in Figure 3-10), and to cases where it is expressed in binary terms. E.g., if a performance requirement is defined as having to land a spacecraft in a defined area of a planet (a “binary requirement” in terms of possible outcome, in the sense that the landing in the target area is either achieved or not achieved) a corresponding RTL is in such case definable as a maximum value of probability that the spacecraft may fail to land there.
- B. In the cases where a probability parameter is used as a “surrogate” performance measure, e.g., as for making a binary performance measure (landing / non-landing) into a continuous one (e.g., the probability of landing), the risk assessor or manager may express risk tolerance levels in terms of the probability that markers on the surrogate measure may not be met. This may appear to be confusing, as one would be considering the “probability of a probability.” A better way to frame this situation is to view and treat the surrogate (probability) parameter as just another performance measure, with its own pre-defined performance markers to be met or not met at some level of risk tolerance. In other words, risk tolerance levels are a way of expressing levels of tolerance for the uncertainty affecting the outcome of a performance measure, whatever that may have been chosen to be. When setting performance markers on a probability being used as a surrogate parameter, one is considering whether some estimation process adopted as a means of “measuring” such a probability parameter will show that the latter meets

or does not meet the performance marker being considered. That is, in such situations a probability estimate is itself treated as surrogate evidence, taking on the same role that direct measurements or observations may have in the definition of whether a physical performance parameter meets or not its performance markers.

- C. Given a performance measure relative to an objective, more than a single risk tolerance level may be associated with it, if more than one performance marker is established. More specifically, if the performance measure has both a *performance requirement* and a *performance goal* established for it, as is the case in the example of Figure 3-10, two different *risk tolerance levels* should correspondingly be defined. The one defined for the *performance requirement* marker must be lower (usually significantly lower) than the one for the *performance goal* marker.
- D. Defining risk tolerance – and consequently, risk acceptability – in terms of limits on the probability of under-performance with respect to performance markers, as NPR 8000.4 indicates, differs from the common past practice of determining and displaying risk tolerability or acceptability directly on the basis of where a specific risk falls in a probability-vs-severity risk matrix: the former model utilizes the concept of distributed probability for a level of performance, whereas the latter defines the tolerability or acceptability of individual risk scenarios in terms of point-values or discrete categories of performance shortfalls and associated probabilities. Although equivalences may be established by ad-hoc “mappings” between the two perspectives, the means and criteria recommended by NPR 8000.4 are what constitutes the first basis of consideration for risk acceptance or non-acceptance, and this can be communicated and displayed in even simpler form than by the traditional matrix format, i.e., directly via a three-tier “risk bar” as illustrated below by Figure 3-12.




Figure 3-12. Risk-Bar Classification for Case of Increasing PM Direction of Goodness

Figure 3-12 shows a “risk-bar” – or, discrete “risk-meter” – set up with the traditional three risk colors: Green to show risks that are classified as “Acceptable,” Yellow for risk that are classified as “Marginal,” and Red for risk that are classified as “Unacceptable.” The “coloring criteria” illustrated by the figure are intended for the purpose of classifying and displaying risk(s) relative to a performance objective which is expressed via a performance measure PM and two performance markers, i.e., more specifically, a performance requirement PMK-R with an associated risk tolerance level RTL-R, and a performance goal PMK-G with an associated risk

tolerance level RTL-G. To complete the example, a risk R_{XY} , representing the same hypothetical risk depicted by Figure 3-10, is placed in the “Yellow” area of the matrix, consistently with its classification of “Marginal Risk” per the criteria shown in the figure.

It is again brought to the reader’s attention that the risk classification criteria expressed by Figure 3-12 apply to the case where the performance “direction of goodness” coincides with increasing values of the Performance Measure (PM). For the opposite case, i.e., when performance goodness increases with decreasing values of PM, what matters in terms of *risk tolerance* is the probability of exceedance (rather than the probability of shortfall) of PM with respect to *performance markers* such as PMK-R and PMK-G. In this latter situation, the assessment part of the risk evaluation process is better accomplished with the use of a display of the PM probability distribution in Complementary Cumulative Distribution Function (CCDF) form, rather than the CDF (Cumulative Distribution Function) form shown in Figure 3-12. A CCDF display directly plots the probability of exceedance (i.e., the probability that PM is greater than some value X) instead of the probability of non-exceedance of that value, as a CDF does. Figure 3-13 shows an example of PM CCDF for an assessment of performance risk that is complementary to the case shown in Figure 3-10.




Figure 3-13. Satisfaction of Risk Tolerance Levels Relative to a Performance Measure CCDF

For the alternative case illustrated by Figure 3-13, Figure 3-14 illustrates the risk classification criteria that correspondingly apply to this type of situation, where “goodness” increases with decreasing values of the PMs that quantify the activity objectives. Obvious examples of such conditions are when addressing cost and budget objectives and attending risk(s), or when using a probability parameter as a measure of safety performance (e.g., as with the Probability of Loss of Crew (P(LOC)) in human operated space mission).

Risk Classification Criteria for Case of Decreasing PM ≡ "Direction of Performance Goodness"		
Risk Classification based on comparison of PM probability with RTL	Risk of Performance Shortfall	
	$P1 > RTL-R$	"High" Probability of PM Exceedance w respect to PMK-R ≡ Unacceptable Risk
	$[(P1 \leq RTL-R)$.AND. $(P2 > RTL-G)]$	R_{xy} "High" Probability of PM Exceedance w respect to PMK-G ≡ Marginal Risk
	$[(P1 \leq RTL-R)$.AND. $(P2 \leq RTL-G)]$	"Low" Probability of PM Exceedance ≡ Acceptable Risk

LEGEND

Unacceptable Risk

Marginal Risk

Acceptable Risk

PMK-R = Performance Requirement
 PMK-G = Performance Goal
 P1 = Probability that PM > PMK-R
 P2 = Probability that PM > PMK-G
 R_{xy} : Hypothetical Risk classified according to the probabilities X and Y identified in Figure 3-13

Figure 3-14. Risk-Bar Classification for Case of Decreasing PM Direction of Goodness

3.3.3 Use of Risk Tolerance Level Thresholds in Decisions for Risk Acceptance

The preceding discussion has addressed how the guidance of NPR 8000.4 for the expression and definition of risk tolerance requires:

- The definition of performance threshold values – which this handbook refers to as **performance markers**;
- The correspondent definition of maximum probability limits on whether such markers may not be satisfied by the future actual performance outcome – which this handbook refers to as **risk tolerance levels**.

Risk tolerance levels can be defined and utilized in different degrees of formality within a given organization. For example, the risk tolerance level for crew safety is characterized in very formal fashion in terms of probability of loss of crew, P(LOC). Specifically, NPR 8705.2, *Human Rating Requirements for Space Systems* [26], requires the establishment of risk tolerances for crew safety in terms of Agency-level safety goals and thresholds that define long-term targeted and maximum tolerable levels of risk to the crew. This serves as guidance to developers in evaluating "how safe is safe enough" for a given type of mission.

The overall organizational risk posture, articulated into the specific expression and definition of levels of risk tolerance that an organization has towards risks affecting performance in various pertinent performance dimensions, is what should ultimately inform the management acceptance of its actual risk. The actual determination of whether a risk is acceptable or not should be made by responsible and accountable managers, consistently with both project-specific definitions of risk posture and general principles established and observed across the organization. NPR 8000.4 defines risk acceptability criteria under which a NASA organizational entity has the authority to accept a risk. For NASA, whose activities often involve coordination among multiple organizational units, effective risk communication is essential, so "yellow" risks in the "marginal risk region" can be accepted by the organizational entity but must also be reported up to the higher-level organizational unit.

3.3.4 Risk Tolerance Levels for Individual Risk Scenarios

Risk tolerance levels (RTLs) are anchored to an organization's objectives, consistent with the definition of risk in NPR 8000.4 as "the potential for shortfalls with respect to achieving explicitly

established and stated objectives.” As such, the tolerability of an organization’s risk profile is fundamentally concerned with the aggregate risk (AR) to each of its objectives, as discussed in Sections 3.3.1 and 3.3.2. However, as a practical matter, the management of aggregate risk is primarily accomplished via the management of the *individual risk scenarios (IRSs)* that contribute to that aggregate, so there needs to be a rational procedure for defining RTLs for such individual risks that does not lose sight of the overall purpose of ensuring that the quantitative definition of risk tolerances at the risk management operational level is done consistently with the organization’s declarations of risk posture relative to each project / activity objective of concern.

RTLs for individual risk scenarios should have the following properties:

- a. They should be derived from RTLs for aggregate risk, such that there is reason to believe that the aggregate risk is tolerable if each of the contributing individual scenarios is tolerable.
- b. They should be set in a manner that takes into account the various ways by which individual scenarios can, by single occurrence of an associated set of events and consequences or in combination with other IRSs and their associated events, compromise the achievement of an organization’s objectives. For this purpose it is helpful to keep in mind the considerations made in Section 3.2.3 with regard to how the coalescence of IRSs into an AR can actually be defined and assessed – i.e., whether the contributions of the IRSs to the AR can happen on a mutually exclusive, cumulative, correlated, or uncorrelated basis in terms of the ultimate effect on the affected objective and performance measure(s). With regard to the establishment of IRS RTLs (referred to in the following as Individual Risk Scenario Risk Tolerance Levels (IRTLs)) and their relationship to the corresponding RTLs for aggregate risk (referred to as Aggregate Risk Tolerance Levels (ARTLs)), it can be established that the definition criteria should, as a minimum, take into due account the following alternative types of IRSs that are of concern:
 - 1) High-Consequence IRSs (HCIRSs) that, if they occur, directly result in the failure to meet an objective. Risks that threaten safety objectives in an intrinsically binary way are typically of this type. For example, if crew safety is an objective, then an individual scenario whose consequence is expressed in binary terms as the loss of crew (LOC) consequence is sufficient on its own to produce the undesired violation of the safety objective. Even when the measure of achievement of an objective is not intrinsically binary but represented on a continuum by a performance measure (PM), and the achievement or not of the objective is therefore expressed “in binary mode” by reference to a performance marker (PMK) set within that continuum, an IRS with sufficiently large negative impact is potentially capable of compromising the objective without concurring contributions by other IRSs.
 - 2) Cumulative-Consequence Individual Risk Scenarios (CCIRSs) that on their own cannot produce a failure to meet an objective, but in sufficient numbers can. Risks that threaten cost objectives are typically of this type. For example, if an activity’s baseline cost is at some margin below a cost cap, then an individual scenario whose consequence is an erosion of margin cannot on its own push the cost of the activity

beyond its cap, but if sufficient numbers of such risks are realized then the cost cap will be breached and the activity will fail to meet its cost objective. Similar situations can occur when the PM of concern is project / activity schedule.

- c. They should enable individual risk scenarios to be managed individually, i.e., in relative independence from what the “big picture” of aggregate risk might be assessed to be.

Note: This does not eliminate the separate need to assess the tolerability of aggregate risk and to respond accordingly.

3.3.4.1 Steps of IRTL Definition and IRS Classification Process

The derivation and utilization of an operationally effective IRTL value that can be valid for both the HCIRS and CCIRS types of scenarios can be accomplished via a simple heuristic process, as defined by the steps described below:

- A. Estimate the number N of IRSs that are deemed to be significant contributors to the AR relative to a given objective.
- B. For each of the performance markers (PMK_i) identified for the assessment and management of aggregate risk, utilize the aggregate risk RTL values ($ARTL_i$) associated with it and the N value estimated in Step A to set an initial value of $IRTL_i$ to be used in relation to that marker as: $IRTL_i = ARTL_i / N$.
- C. Assess each identified IRS according to whether its contribution over the “baseline” PM estimate (see Section 3.3.4.2 below for definition) results or not in a violation of a PMK with a probability (P) higher than the corresponding IRTL. For example, in the case of a PM such as Cost, for which the “direction of goodness” is from higher to lower values, the questions of concern will be, in relation to each PMK_i , whether:
$$P(PM > PMK_i) > IRTL_i$$
- D. Assign a “stop-light” risk classification, i.e., “Acceptable / GREEN,” “Marginal / Yellow,” or “Unacceptable / RED” to each IRS, according to the criteria also used for the classification of aggregate risk, as previously discussed in Section 3.3.2.

Some of the key elements of the above stepwise process require further discussion. This is provided in the sections that follow below.

3.3.4.2 Assessment of IRSs with Respect to Baseline Estimates and Performance Markers

Step C of above Section 3.3.4.1 implies that for an individual risk scenario, the performance CDF or CCDF for a given performance measure should be referred to an expected “*baseline estimate of performance*,” i.e., the level of performance expected for the project / activity of concern when everything goes according to plan and no major risks become realized. The outcomes of PM beyond such a baseline (or below it, depending on the “direction of goodness”) should only reflect the inclusion of the effect of that single individual scenario on the PM. Such a CDF / CCDF therefore differs from the CDFs and CCDFs of Figure 3-10 and Figure 3-13, respectively, which on the contrary reflect the inclusion of all known individual risk scenario contributions.

Figure 3-15 below notionally shows the CCDF for an individual risk scenario in both a fully detailed quantitative form (Figure 3-15a), and in an approximate discrete stepwise form (Figure 3-15b). This is in recognition that the assessment of probability distributions may result in either form of results, depending on the amount and quality of the related information, and to emphasize that the procedures illustrated in the following can be executed in either case with sufficient proficiency to support the IRS assessment of interest.

3.3.4.3 Estimation of the Number of Individual Risk Scenarios

As indicated by Step A in Section 3.3.4.1, in order to derive IRTLs for individual risk scenarios from ARTLs for aggregate risk, it is first necessary to estimate the expected number of significant individual scenarios that threaten the objective in question. Past experience has shown that the number of such risks tends to be relatively low, on the order of ten or less [19]. As in Section 3.3.4.1, this number is referred in this section as N .

At the beginning of an activity, N can be estimated based on the number of key individual risks of major impact potential that have been carried over from Activity-Planning RIDM, augmented with some margin to account for the possibility of additional risk identification during CRM.

As the activity proceeds, N may change as some risk scenarios are closed and new ones are identified. Consequently, IRTLs for individual risk scenarios are not necessarily fixed throughout an activity but may change with the evolving risk profile and the point in the life cycle of the activity.




Figure 3-15. CCDF for an Individual Risk Scenario in Continuous and Approximate Discrete Form

3.3.4.4 Consideration of the Effect of Cumulative-Consequence Individual Risk Scenarios

Given a performance marker, an aggregate RTL for the performance marker, and an estimate N of the number of individual risk scenarios that threaten performance, Section 3.3.4.1 has provided the heuristic rule for the determination of individual risk scenario RTLs. In general, such a rule works equally well for HCIRSs of a non-cumulative nature and for CCIRSs. Given that the nature of the latter may be of special concern, in the sense that the potential magnitude of the consequences of

a multiple materialization of CCIRSs may appear to be almost unbounded, it is useful to present some examples that show how the heuristic IRTL-setting rules of Section 3.3.4.1 work for representative cases of CCIRS cumulation, in terms of the resulting Aggregate Risk profiles.

To illustrate such applications of the heuristic IRTL rules defined above, a detailed example that considers a realistic situation concerning a project Cost is provided in Appendix D.

It must be noted that, while the two illustrations of the heuristic IRTL-setting criteria provided in Appendix D correspond to realistic situations of possible IRS cumulative effects on AR and show the practical usefulness of the criteria for those situations, they nevertheless cannot be interpreted as being a proof of general validity of the criteria under any other alternative circumstances of IRS aggregate effects. As previously mentioned, the assessment and evaluation of IRSs according to criteria that consider them individually must be complemented by a parallel and well organized assessment of Aggregate Risk for each of the organizational objectives and performance dimensions of concern.

3.3.5 Risk Classification of Individual Risk Scenarios

The classification of individual risk scenarios can be conducted using criteria similar to those presented in relation to the classification of Aggregate Risk. Given a Performance Requirement (PMK-R) and associated ARTL-R, and a Performance Goal (PMK-G) and associated ARTL-G, corresponding individual risk scenario RTLs can be defined (i.e., an Individual Risk scenario IRTL-R and an Individual Risk scenario IRTL-G, using the simple heuristic criteria discussed in Section 3.3.4.1. Using these IRTL values, the logic of Table 3-I can be directly adapted to the classification of individual risk scenarios, as shown in Table 3-II.

Table 3-II. Classification of Individual Risk Scenarios

IRTL-R	IRTL-G	Individual Risk Scenario “Acceptability” Classification
Satisfied	Satisfied	Green / “Acceptable”
Satisfied	Not Satisfied	Yellow / “Marginal”
Not Satisfied	n/a	Red / “Unacceptable”

Based on Table 3-II, the individual risk scenario shown in Figure 3-16 would be classified as “Marginal”/YELLOW because, while its probability of violation of the PMK-R marker satisfies the IRTL-R limit, its probability of violation of the PMK-G marker exceeds the IRTL-G limit.

3.3.5.1 Alternative Criteria for Classification of Individual Risk Scenarios

The IRS classification scheme presented in Table 3-II is intuitively consistent from an overall perspective with the definition of performance markers and of the associated RTLs in terms of requirements and goals. However, independently of the objectively defined and assessed values of PMs and associated probabilities that identify risks, once such an assessment has been made, the

classification and “coloring” of such risks remains a matter of context and perspective for which individual project or activity managers may apply criteria that deviate from the general guidelines represented in Table 3-II. Examples of alternative and more conservative risk classification criteria are discussed in the following section, and also utilized in the application examples discussed in Part 2 Chapter 4.




Figure 3-16. Individual Risk Scenario CCDF Relative to the IRTLs

3.3.6 Display of Individual Risk Scenarios in Traditional Risk Matrix Format

As stated upfront in Chapter 1, a consistent application of the Risk Leadership (RL) and Objectives-Driven Risk Management (ODRM) principles requires the distinct application of *risk assessment* and *risk classification* processes and means. The assessment processes are accordingly rooted in the utilization of rigorous and objectives assessment metrics, and the downstream classification processes are based on criteria that, proceeding from the objectively derived risk metrics and probability distributions, relate these to *risk tolerance levels* that represent the practical and operational expression of the risk leadership principles established by the leaders and managers of the organization.

In the above perspective, the common past practice of conflating risk assessment and risk classification into the single step of placing each individual risk scenario into one of the quadrants of a “5x5 risk matrix,” purely on the basis of mostly qualitative assessment criteria, can no longer be considered valid. However, once an IRS has been assessed and classified according to the means and criteria discussed in this and following chapters, such a risk can, if so desired for any practical reasons, still be displayed and communicated by means of a 5x5 matrix. However, for the purpose and mode of a matrix display to be consistent with the ODRM framework of risk assessment and classification, the matrix discrete bins and quadrants must be referred to and anchored by the PMK and IRTL values that underpin the framework.

The following presents examples of matrix definitions based on the above guideline of discrete binning anchored to PMK and IRTL values. To illustrate that IRS classification criteria may vary according to the more or less conservative disposition of a project or activity, the examples cover, besides the use of a matrix bin classification scheme consistent with the basic IRS classification

criteria discussed in Section 3.3.5 and summarized in Table 3-II, cases where a more conservative definitions of “risk color” is adopted. Some degree of extra-conservatism in the classification and management of individual risk scenarios is in fact practically admissible as a reflection of an explicit and deliberate choice made by project and activity managers. The caveat, when choosing a set of IRS classification criteria that deviate from the definitions first discussed in Section 3.3.5, is that any such deviations should not result in an application of **aggregate risk posture** that is at odds with the general risk posture directives provided by the top leadership of the organization to which the project or activity of concern belongs.

3.3.6.1 IRS Matrix Definition Consistent with Table 3-II Classification Criteria

The examples provided in this section illustrate two cases of matrix definition. The first case concerns the situation, previously discussed in Section 3.3.5, where an activity objective is represented via a performance measure scale with two performance markers, i.e., a requirement PMK-R and a goal PMK-G, identified along with the corresponding risk tolerance levels, IRTL-R and IRTL-G. The second case addresses the alternative situation, the possibility of which was earlier discussed in Section 3.3.1, where only one marker, simply referred to as “PMK,” is used along with its corresponding risk tolerance level RTL.

The 5 x 5 matrix definition corresponding to the first of the two cases is shown in Figure 3-17.




Figure 3-17. Risk Matrix for Case of Two Performance Markers

To obtain the 25-quadrant granularity of a 5 x 5 matrix, additional threshold values of performance and probability, besides the anchoring PMK and IRTL values, appear in the two respective dimensions. In the figure example, the additional probability thresholds are defined as three-fold

multiples or fractions of IRTL values, but different multiplier and/or divider values could be selected, depending on context and project/activity managers' preferences. The additional thresholds used for performance bin definition are also to be selected at the project/activity managers' discretion, and their general definition and meaning is as follows:

- A value X representing a % deviation from the PMK-R value and defining the boundary between performance being “marginally better” and “substantially better” than PMK-R
- A value Y representing a % deviation from the PMK-G value and defining the boundary between “marginally better” and “substantially better” than PMK-G

The additional performance and probability threshold values should be selected according to the degree of conservatism that the project or activity context may suggest for the day-to-day management of IRSs. It is noted in this regard that the classification of risk into the traditional “green,” “yellow,” and “red” categories provided in this example is somewhat more conservative than what would result from a literal application of the criteria defined in Table 3-II. There is no contradiction in this, as the definition of matrix bin boundaries, in finer granularity than what can be obtained by the sole use of PMK and IRTL values as anchoring points, is a project/activity choice based on the practical objective of communicating and tracking individual risk scenario within the management context of that project or activity, and within CRM-type of risk management executions. To illustrate this, different definitions of risk matrices are used in some of the examples presented in other parts of this handbook. However, to maintain consistency with the risk posture expressed via performance marker and risk tolerance definitions, any ultimate decisions on risk acceptability should be based, for individual risk scenarios, on the IRTL criteria represented in Table 3-II and for aggregate risk on the ARTL criteria represented in Table 3-I.

The second matrix example, relative to a situation where just one PMK and associated IRTL value are utilized, is illustrated by Figure 3-18.

The figure shows a matrix similar in its quadrant coloring to the one presented in Figure 3-17. This is consequent from the assumption made in constructing the matrix definition that the performance-bin threshold values Y and Z appearing in the figure may be in the order of 10% and 20% or less, respectively. Selection of different values for these thresholds, as well as for the multiplier and divider values defining probability thresholds other than the pre-established RTL value, may result in a different definition of the matrix quadrant risk coloring.

In the context of Figure 3-18, and as noted within the figure itself, the meaning of the X, Y, and Z values used in the performance bin definitions is as follows:

- X is a % deviation from the PMK value that defines the boundary between “substantially worse” and “marginally worse”
- Y is a % deviation from the PMK value that defines the boundary between “marginally better” and “moderately better”
- Z is a % deviation from the PMK value (larger than Y) that defines the boundary between “moderately better” and “substantially better”




Figure 3-18. Risk Matrix for Case of a Single Performance Marker

3.3.7 Stepwise Recap and Example of Risk Leadership and Objectives-Driven Risk Management Application

In order to further assist the readers' understanding of the conceptual and operational flow of the Objectives-Driven Risk Management (ODRM) action and implementation processes, Appendix E provides a stepwise recap and example of risk leadership, posture, and risk tolerance application. The example highlights the key steps that are involved in linking together the fundamental concepts introduced in Chapters 2 and 3. It should be noted that this example does not aim to cover all the intricacies and details of an RM execution, but rather focuses on the implementation of the principal concepts of interest.

3.4 References for Chapter 3

1. Merriam-Webster Dictionary, <https://www.merriam-webster.com/dictionary/risk>.
2. Oxford Dictionary, <https://en.oxforddictionaries.com/definition/risk>.
3. Wikiversity, <https://en.wikiversity.org/wiki/Risk>.
4. NASA Special Publication, NASA/SP-2010-576, NASA Risk-Informed Decision Making Handbook, Version 1.0. April 2010.
5. NASA Special Publication, NASA/SP-2011-3422, NASA Risk Management Handbook, Version 1. November 2011.
6. NASA Special Publication, NASA/SP-2010-580 Version 1.0, NASA System Safety Handbook, Volume 1: System Safety Framework and Concepts for Implementation. November 2011.

7. NASA Special Publication, NASA/SP-2014-612, NASA System Safety Handbook, Volume 2: System Safety Concepts, Guidelines, and Implementation Examples. November 2014.
8. NASA Special Publication, NASA/SP-2014-615, Organizational Risk and Opportunity Management: Concepts and Processes for NASA's Consideration. November 2016.
9. NASA Procedural Requirements, NPR 8000.4C, Agency Risk Management Procedural Requirements. April 2022.
10. NASA Internal Report, Risk Management Tiger Team Report. September 07, 2023.
11. NASA Policy Directive, NPD 1000.0C, NASA Governance and Strategic Management Handbook. January 2020.
12. NASA Policy Directive, NPD 2810.1, NASA Information Security Policy, January 2022.
13. NASA Procedural Requirements, NPR 2810.1, Security of Information and Information Systems, January 2022.
14. NIST, NIST SP 800-37, Risk Management Framework (RMF) for Information Systems and Organizations: A System Life Cycle Approach for Security and Privacy, December 2018.
15. NIST, NIST SP 800-53, Security and Privacy Controls for Information Systems and Organizations, December 2020.
16. NASA Policy Directive, NPD 1600.2, NASA Security Policy, February 2024.
17. NASA Procedural Requirements, NPR 1600.1, NASA Security Program Procedural Requirements, August 2013.
18. NASA Procedural Requirements, NPR 1620.3, Physical Security Requirements for NASA Facilities and Property, May 2019.
19. Guarro, S., Risk Assessment for Decision Support and Resilience in Space Systems, Aerospace Corporation Report ATR-2020-00733. July 2020.
20. NIST, NIST AI 100-1, Artificial Intelligence Risk Management Framework, January 2023.
21. NIST, NIST Artificial Intelligence Risk Management Framework Playbook. 2025.
22. JWST Independent Comprehensive Review Panel (ICRP) Final Report. October 2010.
23. NASA Common Leading Indicators Detailed Reference Guide, NASA Office of the Chief Engineer, January 2021.
24. Bearden, D., Perspectives on NASA Mission Cost and Schedule Performance Trends, NASA Goddard Space Flight Center System Engineering Symposium, June 2008.
25. NASA Special Publication, NASA/SP-2011-3421 Probabilistic Risk Assessment Procedures Guide for NASA Managers and Practitioners, Second Edition. December 2011.
26. NASA Procedural Requirements, NPR 8705.2C, Human-Rating Requirements for Space Systems. July 2017.

4 Risk-Informed Decision Making

The content of this chapter presents the framework and step-wise processes that permit the execution of Risk-Informed Decision Making (RIDM), as one of the two principal tools by which risk management is implemented and carried out at NASA.

Earlier chapters of this Handbook have introduced RIDM as a tool that can be applied in any of the NASA activity domains – i.e., the enterprise, program/project, and institutional domains – and also in different stages of an activity or project progression, wherein the objectives can qualitatively differ. More specifically, Section 2.2.5, and in particular Figure 2-8, has addressed at an introductory level the conceptual distinction between the application of RIDM at “Activity-Planning,” for selection of activity paths and solutions, at “Activity-Rebaseline,” for rebaselining of requirements and risk tolerances should the activity decision authority determine that such rebaselining is necessary, and at “Activity-Execution,” for deciding upon a risk response during the CRM Plan step.

Because the typical contexts of RIDM application span a wide range, the RIDM processes and analytical components can be adapted and tailored to the objectives and complexity of such potentially different contexts. In line with this consideration, the notion of a “graded approach” and tailoring of RIDM analytical features, according to the nature and prominence of its application, is explicitly addressed and discussed in Sections 4.1.1 and 4.11. However, rather than attempting to define and describe separate versions of RIDM adaptation and tailoring, this chapter first focuses on providing a discussion of RIDM steps and components in their most general and complete form. This description defines a fully expanded reference flow of RIDM process and steps, typically corresponding to the type of application that is appropriate and recommended at the “Activity-Planning” stage of a high-value and complex activity. This comprehensive definition and discussion of RIDM elements constitutes a large portion of the contents of this chapter and, besides providing guidance for the corresponding RIDM context of application, it also serves as a “RIDM baseline” reference when other types of RIDM applications are discussed in following sections of the chapter. In these further descriptions, RIDM adaptations for different application contexts are described in comparative terms with respect to the steps and analytical components of the RIDM baseline definition articulated in the first part of the chapter.

Part 2 of this handbook is dedicated to examples of RM application, including RIDM applications, in different domains and contexts. The reader will find in those examples an illustration of some of the tailored and graded forms of application that the RIDM processes and analyses may take in domain-specific contexts.

4.1 Range and Objectives of RIDM Application

The range of activity contexts and conditions in which RIDM processes can be applied is broad, but the principal shared characteristic of these conditions is that they involve the need, by a deliberating entity, to identify an optimal path forward considering the risks associated with each alternative. The “activity” and “paths” of concern may vary considerably, depending on the associated domain and context. E.g., in the enterprise domain they may involve the global trade space for definition of a portfolio of space missions and other initiatives that satisfy the Agency’s overall objectives and goals within the constraints of assigned budgets and timelines, while in the

program/project domain they may typically concern alternative system design selections, and in the institutional domain the identification of infrastructure developments necessary to achieve some desired level of technical capability.

Regardless of the domain context, RIDM is applicable, at any stage of the activity or project of concern, whenever a choice must be made among activity strategies, technical solutions, or rebaselining alternatives that carry weight and consequence. By applying RIDM, the deliberating authority seeks a solution for achieving the activity objectives that balances the expenditure of organizational resources and technical effort with the potential risks and opportunities that each possible execution path, technical solution, and/or rebaselining alternative presents.

In essential terms, RIDM has the objectives of making explicit the representation of risk profiles for the parameters and variables by which the achievement of an activity or project objectives can be represented and judged, and of allowing the selection of a particular solution to be informed by this awareness of the degree of risk that can be identified, at solution-selection time, as being associated with each solution.

4.1.1 Adaptive and Graded Approach to RIDM Application

The RIDM steps and analyses are discussed in their possible full scope and details in ensuing sections of this chapter; however, as noted earlier, their actual application can, and should, be tailored in its breadth and depth to the context, relevance, and priority of the decision for which RIDM is invoked. It is the responsibility of the authority that is calling for the application of RIDM to make sure that a suitable adaptation and/or graded approach in the level of rigor applied is followed by the team of analysts that is asked to tailor and execute the RIDM steps and supporting techniques for a specific decision context and scope.

In the section discussing decision analysis, the NASA Systems Engineering Handbook identifies certain characteristics of a decision context that generally correlate with the need for formality and rigor in its execution. In the decision context of a RIDM application, the same considerations can suggest criteria for the degree of breadth and depth to be used in the execution of RIDM steps and analyses. The following descriptions of such characteristics are taken verbatim from the Rev.2, 2016 version of the SE Handbook:

- *Complexity: The actual ramifications of alternatives are difficult to understand without detailed analysis.*
- *Uncertainty: Uncertainty in key inputs creates substantial uncertainty in the ranking of alternatives and points to risks that may need to be managed.*
- *Multiple Attributes: Greater numbers of attributes cause a greater need for formal analysis.*
- *Diversity of Stakeholders: Extra attention is warranted to clarify objectives and formulate technical performance measures.*

Besides the above general principles, when considering the adaptation of RIDM processes to a specific application context, the top-level characterizing factors of the latter that are most relevant are Activity Class, the stage of development and execution of the activity, and the type of decision

being supported by RIDM. These factors, which can be interdependent, are briefly discussed below in Subsections 4.1.1.1 and 4.1.1.2, to set the stage for more specific guidance on the shaping and tailoring of RIDM steps in contexts of interest. This further guidance is provided in Section 4.11, after the complete RIDM baseline process is defined and discussed in Sections 4.2 through 4.10.

4.1.1.1 Effect of Activity Classification on RIDM Tailoring

NPR 8705.4 classifies NASA robotic missions and instruments into four “risk tolerance classes” (Class A thru D), according to a set of factors that among others include cost, national significance, criticality to Agency strategic plan, and complexity. For each of the classes, it then provides guidance for implementation of certain Safety and Mission Assurance (SMA) processes.

A classification scheme similar to the one defined by NPR 8705.4 may be adopted to tailor the application of RIDM to activities of different levels of criticality and importance executed in any NASA domain. This is possible by extending, with any needed adaptations, the risk tolerance classes of the NPR to activities that are carried out inside or outside the flight-mission domain. The classification that appears applicable for a given activity may then be used to inform and guide the tailoring of a RIDM application to that activity.

A tabulation that suggests possible levels of tailoring, graded according to an activity classification scheme that mirrors the NPR 8705.4 definitions of risk tolerance classes, is provided and discussed later in this chapter (see Section 4.11.2).

4.1.1.2 Effect of Decision Context and Activity Stage on RIDM Tailoring

Chapter 2 has identified three basic types of RIDM processes that may be initiated and executed in the life cycle of an activity or project:

- A. Activity-Planning RIDM
- B. Activity-Rebaseline RIDM
- C. Activity-Execution RIDM

The RIDM process type also has a significant bearing on the RIDM application scope of breadth and depth, in a manner that is for the most part orthogonal to the activity-class considerations presented above.

Activity-Planning RIDM is in general the RIDM process type of greatest breadth to be carried out for an activity. This is because this process seeks, before activity execution is actually initiated, the identification of activity performance risk dimensions, associated risk tolerances, and performance targets, for the ultimate purpose of selecting a preferred risk-informed activity execution alternative and system concept (as may be applicable for activities that involve the development and execution of a system design).

Activity-Execution RIDM and Activity-Rebaseline RIDM are processes that take place after activity execution is initiated, if certain trigger conditions occur at some subsequent stage of an activity development and execution life cycle. It is useful to recall that, as indicated in Chapter 2, Activity-Rebaseline RIDM is invoked for any given activity only if: a) the activity RM process determines that established requirements and corresponding risk tolerances cannot be met by

application of any feasible and cost-effective risk control measures and plans, and b) the Activity Decision Authority invokes the RIDM process to help determine what rebaselining of activity requirements and execution plans may be necessary as a result.

If an Activity-Rebaseline RIDM is invoked for a given activity, this is more likely to occur in the early stages of activity development and execution when rebaselining is less disruptive and less costly. By its very nature the RIDM scope will be limited in breadth to the reassessment of a subset of performance objectives and performance measures, i.e., the subset which is impacted by the performance risk believed to be beyond tolerance limits. The RIDM process steps and analyses will be accordingly tailored and more narrowly focused than in an Activity-Planning RIDM execution.

The scope and breadth of Activity-Execution RIDM will in general be more narrowly defined than for the other two types of RIDM process, but the depth of the RIDM execution may be greater. The more narrow focus follows again from the intrinsic nature of this type of RIDM process, whereby the selection to be made is among a set of possible alternative risk responses that have been identified as means of addressing a specific major risk issue. The greater depth, if deemed to be needed on the basis of importance considerations, follows from the greater level of maturity achieved in the development of the design or solution as the execution proceeds in time, and the corresponding need and ability to consider details down to a lower level.

The adaptation and tailoring of process steps and analyses for Activity-Execution RIDM will have to be consistent with the narrower focus of the selection to be carried out, which, as just mentioned, is relative to the identification of an optimal risk response concerning individual risk scenarios or a specific performance measure. Larger redefinition or rebaselining of activity requirements, plans, or system designs at an activity-wide level are usually beyond the scope of Activity-Execution RIDM.

4.2 Overview of the RIDM Process

As specified in NPR 8000.4, the RIDM process itself consists of the three parts shown in Figure 4-1. This section provides an overview of the process and an introduction to the concepts and terminology established for its implementation. An overview of the steps associated with each part of the process can be found in Section 4.4, and a detailed exposition follows in Sections 4.5 through 4.10.

Throughout the RIDM process, interactions take place between the **stakeholders**, the **risk analysts**, the **subject matter experts (SMEs)**, the **Technical Authorities**, and the **decision-maker** to ensure that objectives, values, and knowledge are properly integrated and communicated into the deliberations that inform the decision. Figure 4-2 notionally illustrates the functional roles and internal interfaces in an RM and RIDM supported deliberation. As shown in the figure, for a successful execution of the process it is essential that the risk-informed analysis of alternatives incorporate the objectives of the various stakeholders as an integral part of the deliberation criteria that are to be balanced and satisfied. Accordingly, the deliberation-supporting analyses are performed by, or with the support of, subject matter experts in the domains spanned by the objectives. The completed risk-informed analyses are deliberated upon, along with other considerations, and the decision-maker selects a decision alternative for implementation (with the

concurrence of the relevant Technical Authorities). The risk profile associated with the selected decision alternative becomes then the central focus of CRM, which works to control and mitigate its contributors during implementation, thus avoiding performance shortfalls in the actual outcomes.




Figure 4-1. The RIDM Process




Figure 4-2. Functional Roles and Information Flow in RIDM Deliberations

The RIDM process is portrayed in this handbook primarily as a linear sequence of steps, each of which is conducted by individuals in their roles as stakeholders, risk analysts, subject matter experts, and decision-makers. The linear step-wise approach is used for instructional purposes only. In reality, some portions of the processes may be conducted in parallel, and steps may be

iterated upon multiple times before moving to subsequent steps. In particular, Part 2, Analysis of Alternatives, is internally iterative as analyses may be refined to meet decision needs in accordance with a graded approach, and Part 2 is iterative with Part 3, Risk-Informed Alternative Selection, as stakeholders and decision-makers may iterate with the risk analysts in order to develop a sufficient technical basis for robust decision making. Additionally, decisions may be made via a series of downselects, each of which may be made by a different decision-maker who has been given authority to act as proxy for the responsible decision authority.

RIDM Functional Roles*

Stakeholders - A stakeholder is an individual or organization that is materially affected by the outcome of a decision or deliverable but is outside the organization doing the work or making the decision [1]; e.g., Center Directors (CDs), Mission Support Offices (MSOs).

Risk Analysts – A risk analyst is an individual or organization that applies probabilistic methods to the characterization of performance with respect to the mission execution domains of safety, technical, cybersecurity and mission security, cost, and schedule, or with respect to other execution domains applicable to enterprise and institutional activities.

Subject Matter Experts – A subject matter expert is an individual or organization with expertise in one or more topics within the execution domains listed in the preceding item.

Technical Authorities – A technical authority is an individual within the Technical Authority process who is funded independently of a program or project and who has formally delegated Technical Authority traceable to the Administrator. The three organizations who have Technical Authorities are Engineering, Safety and Mission Assurance, and Health and Medical [1].

Decision-Maker – A decision-maker is an individual with responsibility for decision making within a particular organizational scope.

*Not to be interpreted as official job positions but as functional roles.

4.2.1 Part 1, Identification of Alternatives

In Part 1, *Identification of Alternatives*, stakeholder expectations, which in general may be multifaceted and qualitative, are identified and developed into a distinct set of well-defined **performance objectives**, each of which is associated with a **performance measure** that quantifies the degree to which the performance objective is addressed by a given decision alternative. In general, a performance measure has a “direction of goodness” that indicates the direction of increasingly beneficial performance measure values. A comprehensive set of performance measures is considered collectively for decision making, reflecting stakeholder interests and spanning the associated activity execution domains, which typically include the principal dimensions of:

- Safety (e.g., avoidance of injury, fatality, or destruction of key assets)
- Technical (e.g., thrust or output, amount of observational data acquired)
- Physical Security and Cybersecurity (e.g., protection against physical attack or cyberattack)

- Cost (e.g., execution within allocated cost)
- Schedule (e.g., meeting milestones)

Other performance dimensions, such as political, legal, compliance with Government requirements, and/or public support impact, may be considered as well, depending on their relevance in the context of a particular deliberation.

Performance Objectives, Performance Measures, and Imposed Constraints

In RIDM, stakeholder expectations, which may be multifaceted and qualitative, are decomposed into a set of **performance objectives**, each of which is implied by the stakeholder expectations, and which cumulatively encompass all the facets of the stakeholder expectations. Each performance objective relates to a single facet of the stakeholder expectations, and is quantifiable. These two properties of performance objectives enable quantitative comparison of decision alternatives in terms of capabilities that are meaningful to the RIDM participants. In the case of a deliberation concerning the architecting of a human space flight mission, examples of possible performance objectives are:

• Maintain Astronaut Health and Safety	• Minimize Cost
• Maximize Payload Capability	• Maximize Public Support

A performance measure is a metric used to quantify the extent to which a performance objective is fulfilled. In RIDM, at least one performance measure is associated with each performance objective, and it is through performance measure quantification that the capabilities of the proposed decision alternatives are assessed. Examples of possible performance measures, corresponding to the above performance objectives, are:

• Probability of Loss of Crew (P(LOC))	• Cost (\$)
• Payload Capability (kg)	• Public Support Ranking (1 – 5)

Note that, in each case, the performance measure is the means by which the associated performance objective is assessed. For example, the ability of a proposed decision alternative to Maintain Astronaut Health and Safety (performance objective) may be measured in terms of its ability to minimize the Probability of Loss of Crew, P(LOC) (performance measure).

Although performance objectives relate to single facets of the stakeholder expectations, this does not necessarily mean that the corresponding performance measure is directly measurable. For example, P(LOC) might be used to quantify Maintain Astronaut Health and Safety, but the quantification itself might entail an assessment of vehicle reliability and abort effectiveness in the context of the defined mission profile.

An imposed constraint is a limit on the allowable values of the performance measure with which it is associated. Imposed constraints are minimum performance requirements that are pre-defined and negotiated between NASA organizational units in order to define the task to be performed. In order for a proposed decision alternative to be feasible it must comply with the imposed constraints. A hard limit on the minimum payload capability that is acceptable is an example of a possible imposed constraint.

Objectives whose performance measure values must remain within defined limits for every feasible decision alternative give rise to ***imposed constraints*** that reflect those limits. An example of a hard limit in the program/project activity domain that must be respected in all cases is the mass-to-orbit value required for a launch vehicle being developed to support the launch of a certain type of payload, whose mass characteristics are defined and known. In general, a hard performance limit is the value of a performance measure below or above which no useful functionality of the system or product of the activity or mission being considered is delivered to its users and/or stakeholders. Objectives and imposed constraints form the basis around which decision alternatives are compiled, and performance measures are the means by which their ability to meet imposed constraints and satisfy objectives is quantified.

4.2.2 [Part 2, Analysis of Alternatives](#)

In Part 2, *Analysis of Alternatives (AoA)*, the performance measures of each alternative are quantified, taking into account any significant uncertainties that stand between the selection of the alternative and the accomplishment of the objectives. Given the presence of uncertainty, the actual outcome of a particular decision alternative will be only one of a spectrum of forecasted outcomes, depending on the occurrence, nonoccurrence, or quality of occurrence of intervening events. Therefore, it is incumbent on risk analysts to model each significant possible outcome, accounting for its probability of occurrence, in terms of the scenarios that produce it. This results in the definition of a distribution of outcomes for each alternative, as may be characterized by probability density functions (pdfs), or, in the case of a simplified approach, by discretized probability histograms, over the performance measures (see Figure 4-3).

RIDM is conducted using a graded approach, i.e., the depth of analysis needs to be commensurate with the stakes and complexity of the decision situations being addressed. Thus, for example, the type of representation the analysts seek to obtain for the risk/uncertainty profiles of the performance measures that are relevant in the AoA – e.g., development of high resolution probability density functions (pdfs), vs approximate histograms – would be selected accordingly, as risk analysts conduct RIDM at a level sufficient to support robust selection of a preferred decision alternative, without needlessly applying effort above and beyond such a level. If the uncertainty on one or more performance measures is preventing the decision-maker from confidently assessing important differences between alternatives, then the risk analysis may be iterated in an effort to reduce uncertainty. The analysis stops when the technical case is made; if the level of uncertainty does not preclude a *robust decision* from being made then no further uncertainty reduction is warranted. (See the blue box in Section 2.2.3.2 for a definition of *robustness*.)

The principal product of the risk analysis is the ***Technical Basis for Deliberation (TBfD)***, a document that catalogues the set of candidate alternatives, summarizes the analysis methodologies used to quantify the performance measures, and presents the results. The TBfD is the input that risk-informs the deliberations that support decision making. The presence of this information does not necessarily mean that a decision is risk-informed; rather, without such information, a decision is not risk-informed. Appendix F contains a template that provides guidance on TBfD content. It is expected that the TBfD will evolve as the risk analysis iterates.




Figure 4-3. *Uncertainty of Performance Outcomes Due to Uncertainty of Determining Conditions across the Enterprise, Program/Project, and Institutional Activity Domains*

4.2.3 [Part 3, Risk-Informed Alternative Selection](#)

In Part 3, *Risk-Informed Alternative Selection*, deliberation by the decision-maker(s) takes place, preceded by the identification of risk tolerance levels that can be applied and compared with the performance measure uncertainty and risk representations, obtained in the analysis step in the form of distribution functions and/or histograms. This requires input from the stakeholders of the activity or project of concern, and the application by the decision maker of the risk leadership principles that the responsible authorities wish to see reflected in the set-up and execution of the activity or project addressed in the deliberation. The criteria formulated by the decision maker for the setting of risk tolerance and performance measure thresholds implement the organizational risk leadership principles, defining the risk posture that is applied in the RIDM deliberation and selection steps.

After the driving criteria are identified and declared by the decision maker, the risk-informed deliberation may actually take place in iterative steps; i.e., in some cases the decision-maker may first cull the set of alternatives and ask for further scrutiny of the remaining alternatives, or even ask for new alternatives to be defined and analyzed, before an alternative is ultimately selected for implementation.

To facilitate deliberation, risk-informed **performance markers** are identified by the decision maker or his/her proxies as a means of applying risk posture criteria with respect to the performance measures associated with each alternative. Performance markers may typically include both “hard” constraints on performance, i.e., **performance constraints** that are imposed by physical and/or otherwise predefined conditions, alongside **performance targets**, i.e., markers of performance that can be set at levels chosen by the decision-maker or suggested by other stakeholders. Either type

of marker represents performance levels that an alternative is capable of meeting with some probability of exceedance/non-exceedance.

The above concepts are illustrated by Figure 4-4, which is similar to Figure 3-9 (Section 0), but refers to **performance markers** that, as pointed out in the below blue-box and further explained in the following Section 4.2.4, have a different meaning and function than their Section 0 counterparts. The figure limits the illustration to just one of the performance measures, but the concepts that it addresses apply across the set of performance measures considered in the AoA and selection process.

Risk-Informed Performance Markers

A **risk-informed-performance marker** is a performance measure value (or, in some cases, a qualitative definition) corresponding to a **performance constraint** or to a **performance target** level, and associated with a given risk tolerance / probability to be met which is acceptable to the decision-maker, for an alternative to be selected. Performance markers are used within the RIDM process in order to:

- Allow comparisons of decision alternatives in terms of performance capability at the specified risk tolerances of each performance measure (i.e., risk normalized).
- Serve as the starting point for requirements development, so that a linkage exists between the selected alternative, the risk tolerance of the decision-maker, and the requirements that define the objective to be accomplished. Performance markers are not themselves performance requirements. Rather, they represent achievable levels of performance that are used to risk-inform the development of credible performance requirements as part of the overall systems engineering process.

What type of performance measure markers are utilized to represent organizational Risk Attitude in the achievement of activity objectives is context and domain dependent. Figure 4-4 notionally shows two markers, i.e., a performance constraint (PC) and a performance target (PT), for a given performance measure X (PMX) within the set of performance measures being considered in the Analysis of Alternatives. The performance measure is characterized by a probability density function (pdf), due to uncertainties that affect the analyst's ability to forecast a precise value. The risk levels for not meeting a given PC or PT value are represented, respectively, by the areas corresponding to "A" and "A+B".

If the decision maker's overall risk posture is defined in terms of a set of risk attitudes represented by the definition of maximum risk levels that can be accepted for meeting the PC and PT performance marker values, then, in evaluating alternatives, the decision maker may carry out a risk-informed comparison of the selection alternatives, as discussed in Section 4.9.2.

It is noted that an important practical distinction exists between performance markers used in the context of an Activity-Planning or Rebaselining RIDM AoA – i.e., before formal requirements are negotiated and defined between an *Acquirer* and a Provider organization – and performance markers formally used in the execution stages of an activity and discussed in Chapter 3 in relation to the translation of qualitative definitions of risk posture into operational and quantitative risk tolerance levels (RTLs) – see also more discussion of this subject in Section 4.2.4 below.




Figure 4-4. Risk Posture Expressed by Risk-Informed Performance Markers

Deliberation and decision making might take place in a number of venues over a period of time or tiered in a sequence of downselects. Depending on the desired level of documentation deemed appropriate by the decision authorities in relation to the subjects being decided, the rationale for the selected decision alternative may be documented in a detailed Risk-Informed Selection Report (RISR), or in a more succinct decision memo. The rationale should be based on such factors as:

- The risk deemed acceptable for each performance measure;
- The risk information contained in the TBfD; and
- The pros and cons of each contending decision alternative, as discussed during the deliberations.

Guidance for the RISR is provided in Appendix G. This assures that deliberations involve discussion of appropriate risk-related issues, and that they are adequately addressed and integrated into the decision rationale.

4.2.4 Performance Markers in Activity Planning vs. Activity Execution Stages

As pointed out in the blue-box in the preceding section, performance markers identified in the pre-execution stages of an activity or project, for the purpose of identifying preferred risk-informed activity path and system design selections, have a meaning and role that differs from those of the performance markers discussed in Chapter 3 in the context of the definition of risk tolerance levels (RTLs) to be operationally used in the RM processes applied during an activity or project execution stages. The former are preliminarily defined by the team and organization(s) performing the RIDM

analyses and deliberations taking into account both known **performance constraints** and **performance targets** identified and communicated by leaders and stakeholders. They have the function of permitting a meaningful risk-informed evaluation and selection of activity and system design solution among a set of alternatives. The latter are defined later, just before the formal start of the selected activity or project, and formalized for use within that activity or project execution stages. As discussed in Chapter 3, they usually take the form of **performance requirements** or **performance goals**, and have a specific formal role in the contractual or memorandum-of-understanding definition of the agreement between the *Acquirer* and Provider parties for the design, production and delivery of the activity or project product(s). More specifically, Chapter 3 discusses how such execution-stage performance markers also provide the context for the definition and utilization of RTLs in the set-up and implementation of the RM processes applied during activity or project execution.

While a distinction between planning-stage and execution-stage performance markers needs to be made, it is equally important to note that the two types are related in their definitions. Since the preferred course of an activity or project is suggested by RIDM via a comparison of the assessed ability of candidate alternatives to meet the planning-stage performance markers, the execution-stage performance markers eventually formalized for the selected alternative should not be defined in a way that makes significantly different or even incompatible with the planning-stage preliminary version. If this were to be the case, the whole RIDM risk-informed comparisons and deliberations could be invalidated, which clearly is an undesirable situation outcome that would place the whole ODRM framework and processes on the wrong track from the very execution-stage start.

4.3 Accounting for U/U Risks

It is important to recognize and account for the fact that there is a distinction between the risk estimates in key dimensions of the activity or project that are produced by consideration of the initial spectrum of known risks and projections that apply when considering the acceptability of the *total risk*, inclusive of both *known risks* and *unknown and/or underappreciated (U/U) risks*. As an example, suppose that simulation analyses produced an estimate of the probability of loss of crew, or P(LOC), for a particular mission in the order of one in a hundred, or 0.01, but that the mission were believed to be of particularly high complexity and to be executed by a system design of a particularly high degree of novelty (e.g., due to the use of new technology). In such circumstances, historical experience indicates that the contribution of U/U risks to the total risk may be as much as four times the contribution from known risks [2]. In other words, the total risk (known plus U/U) might be a factor of five times the known risk, i.e., in the order of 0.05. Under such a situation, to account for the anticipated contribution from U/U risks, the *risk tolerance level (RTL)* defined per stakeholders' indications of *risk posture* and to be used in the TBfD would have to be applied to a PLOC distribution adjusted, with respect to the simulation estimates, by a multiplier factor that is intended to account for the postulated potential contribution of U/U risk. More detailed guidance on estimation of risk margins to account for U/U risks can be found in Chapter 5.

With respect to Figure 4-4, therefore, the statement in the upper right corner of the figure that “*risk posture concerning performance measure X is expressed by setting risk tolerances for Area A and Area (A + B)*” is to be considered as applicable to a PM distribution adjusted with an appropriate margin factor intended to account for the contribution of *U/U risks*.

Similar considerations apply to risk tolerances applied to performance measure distributions estimated in relation to other mission execution domains, as well as in institutional and enterprise domains. For example, many of the factors that promote *U/U risk* in the safety mission execution domain (such as design complexity, time pressures, inadequate quality control, and deficient management processes) have a direct bearing as well on cost, schedule, security, workforce availability, and legal liability. Thus, risk estimates in all these areas must include margin factors to account for the potential contribution of *U/U risks* to *total risk*.

More discussion on the estimation of margin factors to account for *U/U risks* in the evaluation of PM risk profiles against established risk tolerance levels (RTLs) will appear later in this chapter and in Chapter 5.

4.4 The RIDM Process Steps

The RIDM process outlined above and depicted in Figure 4-1 can conceptually be expanded into six main implementation steps, as illustrated below by Figure 4-5. The figure shows that the RIDM process may be initiated not only by the initial set-up or baselining of an activity that requires the selection of its specific “path forward” solutions, but also by the need to rebaseline the requirements and risk tolerances as a result of an inability to meet the present ones during activity execution, by the need to select risk response options out of a range of major alternatives identified within the execution of the CRM “Plan Stage” of risk management (as was discussed previously in Section 2.2.5 and depicted in Figure 2-8, or by the identification of an opportunity to more effectively achieve the organizational objectives. The level of actual effort necessary in each RIDM step may vary significantly from case to case in specific executions, depending on the activity domain context and the reason for execution.

The figure also shows that a typical RIDM execution also feeds back, at conclusion of its cycle of steps, into the CRM steps of risk management, and provides to CRM the risk profile information relative to the selected alternative for the activity of concern, so that the corresponding risk can be effectively managed from the start in the CRM process.

As mentioned earlier, although Figure 4-1 and Figure 4-5 depict the RIDM process as a linear sequence of steps, in practice it is expected that some steps could overlap in time and that the process may also become in part iterative. In the more complex contexts of deliberation, information from latter steps may feed back into progressively more refined execution of previous steps, until stakeholder issues are adequately addressed and the decision-maker has sufficient information, at a sufficient level of analytical rigor, to make a robust risk-informed decision. The primary issues that may drive the need for iteration are discussed in the following subsections, in the context of the RIDM process steps in which they arise.




Figure 4-5. RIDM Process Steps

The RIDM process has been informed by current theoretical and practical work in decision analysis and analytic-deliberative processes (see, for example, [3], [4] and [5]). Some methodological tools and techniques, generally applicable to structured, rational decision making, such as objectives hierarchies, performance measures, and deliberation, have been directly adopted into the RIDM process. Others, such as analytic hierarchy process (AHP) and multi-attribute utility theory (MAUT), are formally applicable to rational decision making but also present practical challenges in the context of requirements development within a complex organizational hierarchy having its own highly developed program management policies and practices. It is therefore left to the discretion of the practitioner to determine, on a case-by-case basis, whether or not such techniques will aid in deliberation and selection of a decision alternative.

The sections that follow discuss the main activities that support each RIDM step.

4.4.1 Steps in Part 1, Identification of Alternatives

As indicated in NPR 8000.4 and discussed earlier in this handbook, decision alternatives are identifiable only in the context of the objectives they are meant to satisfy. Therefore, identification of alternatives begins with the process of understanding stakeholder expectations and organizational objectives for the activity or project that is the subject of the desired RIDM deliberations. From the identification of expectations and top-level declared objectives, a basis for evaluating decision alternatives is developed by decomposing the former, often qualitative formulations, into more detailed and quantifiable objectives that enable comparison among any candidate alternative activity solutions. Only then, after an appropriate context has been established, is it possible to identify and define a set of feasible alternatives that address the objectives. Figure 4-6 illustrates this part of the process, which is delineated in subsequent subsections.




Figure 4-6. RIDM Process Flowchart: Part 1, Identification of Alternatives

4.4.2 Steps in Part 2, Analysis of Alternatives

Once a definite set of activity and/or system alternatives is identified RIDM carries out an Analysis of Alternatives (AoA) that quantifies and evaluates the performance measures of the possible alternatives from a risk perspective. Risk analysis consists of performance assessment supported by probabilistic modeling. It links the uncertainties inherent in a particular decision alternative to uncertainty in the achievement of objectives, were that decision alternative to be pursued. Performance is assessed in terms of the performance objectives developed in Part 1 Step 1. The performance measures established for these objectives provide the means for quantifying performance uncertainty and risk, so that alternatives can be effectively compared.

Figure 4-7 illustrates Part 2 of the RIDM process, Analysis of Alternatives. In Step 3, risk analysis methodologies are selected for each analysis domain represented in the objectives, and coordination among the analysis activities is established to ensure a consistent, integrated evaluation of each alternative. In Step 4, the risk analysis is conducted, which entails probabilistic evaluation of each alternative's performance measure values, iterating the analysis at higher levels of resolution as needed to clearly distinguish performance among the alternatives. Then the TBfD is developed, which provides the primary means of risk-informing the subsequent selection process.




Figure 4-7. RIDM Process Part 2, Risk Analysis of Alternatives

4.4.3 Steps in Part 3, Risk-Informed Alternative Selection

Figure 4-8 illustrates Part 3 of the RIDM process, Risk-Informed Alternative Selection. In Step 5 (“Deliberate”), **performance targets** are identified, and related to any **performance constraints** representing minimum allowable performance limits via **performance margins**. Performance constraints and targets identified for each relevant performance measure represent performance markers for that measure, with which **risk tolerance** values can be associated to express the organizational **risk attitude** with regard to the corresponding performance objectives. This requires careful consideration, as it constitutes the means by which consistent levels of risk tolerance are applied and represented across alternatives, translating the risk leadership principles of the organization into a practical definition of risk posture in the specific context of the deliberative selection process.

In Step 5, relevant stakeholders and risk analysts deliberate the relative merits and drawbacks of each alternative, given the information in the TBfD. This step is iterative and may involve additional risk analysis or other information gathering as the participants strive to fully assess the alternatives and identify those that they consider to be reasonable contenders, worthy of serious consideration by the decision-maker.

In the following Step 6 (“Select an Alternative and Accept the Associated Risk”) a reduced set of alternatives is identified as being worthy of consideration, and the decision-maker, or his/her proxy, may be more directly involved at this stage to help cull the number of alternatives from the original broader set (a.k.a. downselecting). Once a set of contending alternatives has been identified, the decision-maker integrates the issues raised during deliberation into a rationale for the selection of an alternative. In parallel to the selection of the preferred alternative, the decision-maker finalizes the identification of performance targets in all relevant performance dimensions, and of the associated risk levels that are correspondingly accepted.




Figure 4-8. RIDM Process Part 3, Risk-Informed Alternative Selection

When the decision maker specifies the requirements to be levied upon the *Provider*, that decision maker considers the RIDM-provided performance targets as indicators of what is achievable. In addition, the formulation of requirements by the decision maker normally includes other considerations such as lessons learned from historical experience, corresponding best practices, stakeholder expectations, and *Provider* preferences. These additional considerations may, under some circumstances, not produce an optimized set of requirements for achieving the specific performance objectives of the activity or project under consideration, since lessons learned, best practices, and human expectations and preferences informed by historical experience may not necessarily apply in the present case. The RIDM team, therefore, will normally be asked to provide an analysis of the decision maker's selected requirements leading to the possibility of recommendations that the decision maker may find useful.

In the last sub-steps of Step 6, the actual execution of the selected alternative along with the selected requirements are assigned to an organizational unit, and the decision rationale is documented, in accordance with any existing project management directives, in a Risk Informed Selection Report (RISR). For pedagogical purposes, the process is here initially described as if alternative selection involves a single decision that is made once deliberations are complete. However, as discussed in the Section 4.7.3.6 guidance on sequential analysis and downselection, decisions are often made in stages and in a number of forums that may involve a variety of proxy decision-makers.

Additionally, this handbook refers to the participants in deliberation as deliberators. This is also for pedagogical purposes, as in any given context deliberators may be drawn from any of the sets of stakeholders, risk analysts, SMEs, and decision-makers.

4.5 Details of RIDM Step 1 (Part 1), Identify Objectives and Performance Measures

As described in Section 4.4.1 and illustrated in Figure 4-6, development of a full understanding of **stakeholder expectations** and declared **organizational goals** relative to the activity being considered is a key part of the step of identifying and defining its more detailed and technically-defined objectives, as well as the associated metrics, i.e., the **performance measures** by which the achievement of such objectives can be gauged and verified.

The identification of any imposed constraints that condition an activity, and the development of unambiguous performance measures that permit the objective representation of stakeholder expectations and organizational goals, as well as of such constraints, is the foundation of sound decision making. Paragraph 3.2.1 of NPR 7123.1A establishes systems engineering process requirements for stakeholder expectations definition, and Section 4.1 of the NASA Systems Engineering Handbook provides further guidance on understanding stakeholder expectations.

Typical inputs needed for the organizational goals and stakeholder expectations definition process include:

- **External Stakeholder Expectations:** The expectations that are provided by individuals or organizations that are materially affected by the outcome of a decision or deliverable but are outside the organization doing the work or making the decision.
- **Upper-Level Management Expectations:** These would be the expectations (e.g., needs, wants, desires, capabilities, constraints, external interfaces) that are being flowed down to a particular activity of interest from a higher level within the organization (e.g., Agency, Directorate, Center, Program, Project, etc.).

As implied in the above bullets, a variety of organizations, both internal and external to NASA, may have a stake in a particular decision. Besides internal stakeholders like NASA Headquarters (HQ), the NASA Centers, and NASA advisory committees, external stakeholders might also exist at all levels, including the White House, Congress, the National Academy of Sciences, the National Space Council, and many other groups in the science and space communities.

Stakeholder expectations, the vision of a particular stakeholder individual or group, result when stakeholders specify what is desired as an end state or as an item to be produced and then put bounds upon the achievement of the goals. These bounds may encompass expenditures (resources), time to deliver, performance objectives, or other less obvious quantities such as organizational needs or geopolitical goals.

Typical outputs for capturing stakeholder expectations include the following:

- **Top-Level Expectations:** These would be the top-level needs, wants, desires, capabilities, constraints, and external interfaces for the product(s) to be developed.

- **Top-Level Conceptual Boundaries and Functional Milestones:** When the activity which is the subject of an AoA process involves the design and engineering of a system, these subjects cover and describe how the system will be operated during the life-cycle phases to meet stakeholder expectations. They address the system characteristics from an operational perspective and help facilitate an understanding of the system goals. This is usually accomplished through use-case scenarios, design reference missions (DRMs), and concepts of operations (ConOps).

In the terminology of RIDM, the definitions that constitute the outputs of this step consist of top-level objectives and imposed constraints. Top-level objectives state what the stakeholders and users hope to achieve or obtain from the activity. They are typically qualitative and multifaceted, reflecting competing sub-objectives (e.g., more data vs. lower cost). Imposed constraints represent the top-level success criteria for the undertaking, outside of which the top-level objectives are not achieved. For example, if an objective is to put a satellite of a certain mass into a certain orbit, then the ability to lift that mass into that orbit is an imposed constraint, and any proposed solution that is incapable of doing so is infeasible.

In general, decision alternatives cannot be directly assessed relative to multifaceted and/or qualitative top-level objectives. Although the top-level objectives state the goal to be accomplished, they may be too complex, as well as vague, for any operational purpose. To deal with this situation, objectives are decomposed, using an **objectives hierarchy**, into a set of conceptually distinct lower-level objectives that describe the full spectrum of necessary and/or desirable characteristics that any feasible and attractive alternative should have. When these objectives are quantifiable via performance measures, they provide a basis for comparing proposed alternatives.

4.5.1 Constructing an Objectives Hierarchy

An objectives hierarchy is constructed by subdividing an objective into lower-level objectives of more detail, thus clarifying the intended meaning of the general objective. Decomposing an objective into precise lower-level objectives clarifies the tasks that must be collectively achieved and provides a well-defined basis for distinguishing between alternative means of achieving them.

An objectives hierarchy is shown notionally in Figure 4-9. At the first level of decomposition the top-level objective is partitioned into the NPR 8000.4C performance domains of safety, mission success (technical), physical security, cybersecurity, cost, and schedule. This enables each performance measure and, ultimately, performance or process requirement, to be identified as relating to a single domain. Below each of these domains the objectives are further decomposed into sub-objectives, which themselves are iteratively decomposed until appropriate quantifiable performance objectives are generated.

There is no prescribed depth to an objectives hierarchy, nor must all performance objectives reside at the same depth in the tree. The characteristics of an objectives hierarchy depend on the top-level objective and the context in which it is to be pursued. Furthermore, a unique objectives hierarchy is not implied by the specification of an objective; many different equally legitimate objectives hierarchies could be developed.

Judgment must be used to decide where to stop by considering the advantages and disadvantages of further decomposition. Things to consider include:

- Are all facets of each objective accounted for?
- Are all the performance objectives at the levels of the hierarchy quantifiable?
- Is the number of performance objectives manageable within the scope of the decision-making activity?




Figure 4-9. Notional Objectives Hierarchy

One possibility is to use a “test of importance” to deal with the issue of how broadly and deeply to develop an objectives hierarchy and when to stop. Before an objective is included in the hierarchy, the decision-maker is asked whether he or she feels the best course of action could be altered if that objective were excluded. An affirmative response would obviously imply that the objective should be included. A negative response would be taken as sufficient reason for exclusion. It is important when using this method to avoid excluding a large set of attributes, each of which fails the test of importance but which collectively are important. As the decision-making process proceeds and further insight is gained, the test of importance can be repeated with the excluded objectives to assure that they remain non-determinative. Otherwise, they must be added to the hierarchy and evaluated for further decomposition themselves until new stopping points are reached.

The decomposition of objectives stops when the set of performance objectives is operationally useful and quantifiable, and the decision-maker, in consultation with appropriate stakeholders, is satisfied that it captures the expectations contained in the top-level objective. It is desirable that the performance objectives have the following properties. They should be:

- Complete – The set of performance objectives is complete if it includes all areas of concern embedded in the top-level objective.
- Operational – The performance objectives must be meaningful to the decision-maker so that he or she can understand the implications of meeting or not meeting them to various degrees. The decision-maker must ultimately be able to articulate a rationale for preferring one decision alternative over all others, which requires that he or she be able to ascribe

value, at least qualitatively, to the degree to which the various alternatives meet the performance objectives.

- Non-redundant – The set of performance objectives is non-redundant if no objective contains, or significantly overlaps with, another objective. This is not to say that the ability of a particular alternative to meet different performance objectives will not be correlated. For example, in application, *maximize reliability* is often negatively correlated with *minimize cost*. Rather, performance objectives should be conceptually distinct, regardless of any solution-specific performance dependencies.
- Solution independent – The set of performance objectives should be applicable to any reasonable decision alternative and should not presuppose any particular aspect of an alternative to the exclusion of other reasonable alternatives. For example, an objectives hierarchy for a payload launch capability that had *Minimize Slag Formation* as a performance objective would be presupposing a solid propellant design. Unless solid propellant was specifically required based on a prior higher-level decision, *Minimize Slag Formation* would not reflect an unbiased decomposition of the top-level objective.

Guidance on developing objectives hierarchies can be found in Clemen [3] and Keeney and Raiffa [4], as well as on websites such as Comparative Risk Assessment Framework and Tools (CRAFT) [6].

4.5.2 Fundamental vs. Means Objectives

When developing an objectives hierarchy it is important to use ***fundamental objectives*** as opposed to ***means objectives***. Fundamental objectives represent *what* one wishes to accomplish, as opposed to means objectives, which represent *how* one might accomplish it. Objectives hierarchies decompose high-level fundamental objectives into their constituent parts (partitioning), such that the fundamental objectives at the lower level are those that are implied by the fundamental objective at the higher level. In contrast, means objectives indicate a particular way of accomplishing a higher-level objective. Assessment of decision alternatives in terms of fundamental objectives as opposed to means objectives represents a performance-based approach to decision making, as recommended by the Aerospace Safety Advisory Panel (ASAP) in emphasizing “early risk identification to guide design, thus enabling creative design approaches that might be more efficient, safer, or both.” [7].

The difference between fundamental objectives and means objectives is illustrated in Figure 4-10, reproduced from a tutorial example formulated in [8], which shows a hierarchy of fundamental objectives on the top and a means objectives network on the bottom.

The first thing to notice is that the objectives hierarchy is just that, a hierarchy. Each level decomposes the previous level into a more detailed statement of what the objectives entail. The objective, *Maximize Safety*, is decomposed (by partitioning) into *Minimize Loss of Life*, *Minimize*

Serious Injuries, and *Minimize Minor Injuries*. The three performance objectives explain what is meant by *Maximize Safety*, without presupposing a particular way of doing so.⁹




Figure 4-10. Fundamental vs. Means Objectives [8]

In contrast, the means objectives network is not a decomposition of objectives, which is why it is structured as a network instead of a hierarchy. The objective, *Educate Public about Safety*, does not explain what is meant by any one of the higher-level objectives; instead, it is a way of accomplishing them. Other ways may be equally effective or even more so. Deterministic standards in general are means objectives, as they typically prescribe techniques and practices by which fundamental objectives, such as safety, will be achieved.

4.5.3 Performance Measures

As mentioned in Section 2.1, once an objectives hierarchy is completed that decomposes the top-level objective into a complete set of quantifiable performance objectives, a performance measure is assigned to each as the metric by which its degree of fulfillment is quantified. The appropriate performance measure to use is most often self-evident from the objective, but in some cases, the choice may not be as evident, and effort must be made to assure that the objective is not only quantifiable, but that the performance measure used to quantify it is adequately representative of the objective.

⁹ NASA has developed quantitative *safety goals* and associated *thresholds* (akin to imposed constraints) to be used to guide risk acceptance decisions [9]. An example of a quantitative safety goal would be: the risk to an astronaut from the ascent phase of a launch to LEO should be less than <a specified value>.

Objectives that have natural unit scales (e.g., *Minimize Cost*, *Maximize Payload*,) are generally easy to associate with appropriate performance measures (e.g., *Total Cost* or *Cost Overrun* [\$], *Payload Mass* [kg]). Other objectives might not have an obvious or practical natural unit scale, thereby requiring the development of either a *constructed scale* or a *proxy performance measure*.

A constructed scale is typically appropriate for measuring objectives that are essentially subjective in character, or for which subjective or linguistic assessment is most appropriate. An example of such an objective might be *Maximize Stakeholder Support*. Here, stakeholder support is the attribute being measured, but there is no natural measurement scale by which an objective assessment of stakeholder support can be made. Instead, it might be reasonable to construct a scale that supports subjective/linguistic assessment of stakeholder support (see Table 4-I). Constructed scales are also useful as a means of quantifying what is essentially qualitative information, thereby allowing it to be integrated into a quantitative risk analysis framework.

Table 4-I. A Constructed Scale for Stakeholder Support (Adapted from [3])

Scale	Value	Description
5	Action-oriented Support	Two or more stakeholders are actively advocating and no stakeholders are opposed.
4	Support	No stakeholders are opposed and at least one stakeholder has expressed support.
3	Neutrality	All stakeholders are indifferent or uninterested.
2	Opposition	One or more stakeholders have expressed opposition, although no stakeholder is actively opposing.
1	Action-oriented Opposition	One or more stakeholders are actively opposing.

Alternatively, it may be possible to identify an objective performance measure that *indirectly* measures the degree of fulfillment of an objective. In the previous paragraph the objective, *Maximize Stakeholder Support*, was assessed subjectively using a *Stakeholder Support* performance measure with a constructed scale. Another strategy for assessing the objective might be to define a proxy for stakeholder support, such as the average number of stakeholders attending the bi-weekly status meetings. In this case, the proxy performance measure gives an indication of stakeholder support that might be operationally adequate for the decision at hand, although it does not necessarily correlate exactly to actual stakeholder support.

The relationship between natural, constructed and proxy scales is illustrated in Figure 4-11 in terms of whether or not the performance measure directly or indirectly represents the corresponding objective, and whether the assessment is empirically quantifiable or must be subjectively assessed. Additionally, the figure highlights the following two characteristics of performance measures:

- The choice of performance measure type (natural, constructed, proxy) is not a function of the performance measure alone. It is also a function of the performance objective that the

performance measure is intended to quantify. For example, $P(LOC)$ can be considered a natural performance measure as applied to astronaut life safety, since it directly addresses astronaut casualty expectation. However, in some situations $P(LOC)$ might be a good *proxy* performance measure for overall astronaut health, particularly in situations where astronaut injury and/or illness are not directly assessable.

- There is seldom, if ever, a need for an indirect, subjective performance measure. This is because performance objectives tend to be intrinsically amenable to direct, subjective assessment. Thus, for objectives that do not have natural measurement scales, it is generally productive to ask whether the objective is better assessed directly but subjectively, or whether it is better to forego direct measurement in exchange for an empirically-quantifiable proxy performance measure. The first case leads to a constructed performance measure that is direct but perhaps not reproducible; the second to a performance measure that is reproducible but may not fully address the corresponding performance objective.




Figure 4-11. Types of Performance Measures

A performance measure should be adequate in indicating the degree to which the associated performance objective is met. This is generally not a problem for performance measures that have natural or constructed scales, but can be a challenge for proxy performance measures. In the *Maximize Stakeholder Support* example above, it is possible that a stakeholder who perceives the activity to be an obstacle to his or her real objectives might attend the meetings in order to remain informed about potential threats. Thus, the average number of stakeholders attending the status meetings might not be an accurate representation of stakeholder support, and in this case may have a contraindicative element to it.

Figure 4-12 illustrates the relationship between performance objectives and performance measures. A performance measure has been established on each of the performance objectives based on the objective's natural measurement scale, a constructed scale that has been developed for subjective quantification, or via a proxy performance measure.

Although it is preferable that a performance measure be directly measurable, this is not always possible, even for objectives with natural measurement scales. For example, safety-related and security-related risk metrics such as *Probability of Loss of Mission*, $P(LOM)$, *Probability of Loss of Crew*, $P(LOC)$, and *Probability of Hacker Intrusion Success* are typically used to quantify the objectives *Avoid Loss of Mission*, *Maintain Astronaut Health and Safety*, and *Protect Against Cyber Threats*. These performance measures are the product of modeling activities as opposed to

direct measurement, involving the integration of numerous parameters within an analytical model of the alternative under consideration. In cases such as these, where modeling methods are integral to the resultant performance measure values, the modeling protocols become part of the performance measure definition. This assures that performance measures are calculated consistently.




Figure 4-12. The Relationship between Performance Objectives and Performance Measures

One proxy performance measure of particular importance to many NASA decisions is *Flexibility*. *Flexibility* refers to the ability to support more than one current application. A technology choice that imposes a hard limit on the mass that can be boosted into a particular orbit has less flexibility than a choice that is more easily adaptable to boost more. The objective, *Maximize Flexibility*, allows this type of issue to be addressed systematically in decision making. However, since *Maximize Flexibility* refers to potential capabilities that are as yet undefined, there is no natural measurement scale that can be used for quantification.¹⁰ A constructed scale is possible, although it requires subjective assessment. A proxy performance measure for flexibility can be constructed

¹⁰ In such applications, *Flexibility* is a surrogate for certain future performance attributes. This idea is discussed more extensively by Keeney [8] and Keeney and McDaniels [10].

by, for example, assessing the capability of the alternative to support a selected set of alternative objectives, such as boosting a larger mass into orbit.

4.5.4 Risk Minimization Is Not a Performance Objective

It is sometimes the practice in decision analyses and trade studies to treat *Minimize Risk* as a distinct performance objective, which is then decomposed into domains such as technology, programmatic, cost, and schedule, resulting in performance measures such as *technology risk*, *programmatic risk*, *cost risk*, and *schedule risk*. However, in NPR 8000.4, risk is the potential for shortfalls with respect to performance requirements (which in a RIDM context translates operationally into shortfalls with respect to performance targets). Therefore, *Minimize Risk* is not a distinct objective in the objectives hierarchy. Rather, risk is an attribute of every performance objective, as measured by the probability of falling short of its associated performance target, and the task of risk management is to make sure the risk of achieving the target performance is within the risk posture.

For example, if a certain payload capability is contingent on the successful development of a particular propulsion technology, then the risk of not meeting the payload performance target is determined in part by the probability that the technology development program will be unsuccessful. In other words, the risk associated with technology development is accounted for in terms of its risk impact on the performance targets (in this case, payload). There is no need to evaluate a separate Technology Risk metric.¹¹

4.5.5 Example Performance Measures

Performance measures should fall within the mission execution domains of safety, technical, security, cost and schedule. Table 4-II contains a list of typically important kinds of performance measures for planetary spacecraft and launch vehicles. Note that this is by no means a comprehensive and complete list. Although such lists can serve as checklists to assure comprehensiveness of the derived performance measure set, it must be stressed that performance measures are explicitly derived from top-level objectives in the context of stakeholder expectations, and cannot be established prescriptively from a predefined set.

4.6 Details of RIDM Step 2 (Part 1), Identify Decision Alternatives

The objective of Step 2, as discussed in Section 4.4.1 and portrayed in Figure 4-6, is to identify a comprehensive list of feasible decision alternatives through the consideration of a reasonable range of initially compiled alternatives. The result is a set of alternatives that can potentially achieve objectives and warrant the investment of resources required to analyze them further.

4.6.1 Compile an Initial Set of Alternatives

Decision alternatives developed under the activity solution and plan definition process [11] are the starting point. These may be revised, and unacceptable alternatives removed after deliberation by stakeholders based upon criteria such as violation of technical or safety standards, etc. Any listing of alternatives will by its nature produce both practical and impractical alternatives. It would be of little use to seriously consider an alternative that cannot be adopted; nevertheless, the initial set of

¹¹ Unless *Engage in Technology Development* is a performance objective in its own right.

proposed alternatives should be conservatively broad in order to reduce the possibility of excluding potentially attractive alternatives from the outset. Keep in mind that novel solutions may provide a basis for the granting of exceptions and/or waivers from deterministic standards, if it can be shown that the intents of the standards are met, with confidence, by other means. In general, it is important to avoid limiting the range of proposed alternatives based on prejudgments or biases.

Table 4-II. Performance Measures Examples for Planetary Spacecraft and Launch Vehicles

Performance Measures for Planetary Spacecraft	Performance Measures for Launch Vehicles
• End-of-mission (EOM) dry mass • Injected mass (includes EOM dry mass, baseline consumables and upper stage adaptor mass) • Consumables at EOM • Power demand (relative to supply) • Onboard data processing memory demand • Onboard data processing throughput time • Onboard data bus capacity • Total pointing error	• Total vehicle mass at launch • Payload mass (at nominal altitude or orbit) • Payload volume • Injection accuracy • Launch reliability • In-flight reliability • For reusable vehicles, percent of value recovered • For expendable vehicles, unit production cost at the n^{th} unit

Defining feasible alternatives requires an understanding of the technologies available, or potentially available, at the time the system is needed. Each alternative should be documented qualitatively in a description sheet. The format of the description sheet should, at a minimum, clarify the allocation of required functions to that alternative's lower-level components. The discussion should also include alternatives which are capable of avoiding or substantially lessening any significant risks, even if these alternatives would be more costly. If an alternative would cause one or more significant risk(s) in addition to those already identified, the significant effects of the alternative should be discussed as part of the identification process.

Stakeholder involvement is necessary when compiling decision alternatives, to assure that legitimate ideas are considered and that no stakeholder feels unduly disenfranchised from the decision process. It is expected that interested parties will have their own ideas about what constitutes an optimal solution, so care should be taken to actively solicit input. However, the initial set of alternatives need not consider those that are purely speculative. The alternatives should be limited to those that are potentially fruitful.

4.6.2 Identify Viable Decision Alternatives by Use of a Trade Tree or Matrix

One way to represent decision alternatives under consideration is by a trade tree or trade matrix. Initially, a trade tree or matrix contains a number of high-level decision alternatives representing high-level differences in the strategies used to address objectives. It is then developed in greater detail by determining a general category of options that are applicable to each strategy. Trade tree or matrix development continues iteratively until the identified and possible trade alternatives are well enough defined to allow quantitative evaluation via risk analysis (see Section 4.4.2).

Along the way, branches of the trade tree (or rows of the trade matrix) containing unattractive categories can be pruned, as it becomes evident that the alternatives contained therein are either *infeasible* (i.e., they are incapable of satisfying imposed constraints) or categorically inferior to

alternatives on other branches. An alternative that is inferior to some other alternative with respect to every performance measure is said to be *dominated* by the superior alternative. At this point in the RIDM process, assessment of performance is high-level, depending on simplified analysis and/or expert opinion, etc. When performance measure values are quantified, they are done so as point estimates, using a conservative approach to estimation in order to err on the side of inclusion rather than elimination.

Trade trees and trade matrices are completely equivalent in their utilization and information content, as an example presented in Part 2 Section 3.2.4.2 illustrates. Figure 4-13 below presents an example of launch vehicle trade tree from the Exploration Systems Architecture Study (ESAS) [12]. At each node of the tree the alternatives were evaluated for feasibility within the cost and schedule constraints of the study's ground rules and assumptions. Infeasible options were pruned (shown in red), focusing further analytical attention on the retained branches (shown in green). The key output of this step is a set of alternatives deemed to be worth the effort of analyzing with care. Alternatives in this set have two key properties:

- They do not violate imposed constraints
- They are not known to be dominated by other alternatives (i.e., there is no other alternative in the set that is superior in every way).

Alternatives found to violate either of these properties can be screened out.

4.7 Details of Step 3 (Part 2), Conduct Integrated Risk Analysis of Each Alternative

This step of the RIDM process, as discussed in Section 4.4.2 and portrayed in Figure 4-7, is concerned with how domain-specific analyses, conducted in accordance with existing methodological practices, are integrated into a multidisciplinary framework to support decision making under uncertainty. In general, each mission execution domain has a suite of analysis methodologies available to it that range in cost, complexity, and time to execute, and which produce results that vary from highly uncertain rough order-of-magnitude (ROM) estimates to the detailed simulations. The challenge for the risk analysts is to establish a framework for analysis across mission execution domains that:

- Operates on a common set of (potentially uncertain) **performance parameters** for a given alternative (e.g., the cost model uses the same mass data as the lift capacity model);
- Consistently addresses uncertainties across mission execution domains and across alternatives (e.g., budget uncertainties, meteorological variability);
- Preserves correlations between performance parameters (discussed further in Chapter 5); and
- Is transparent and traceable.

The means by which a given level of performance will be achieved is alternative specific, and accordingly, the analyses that are required to support quantification are also alternative specific. For example, one alternative might meet the objective of *Minimize Crew Fatalities* by developing a high reliability system with high margins and liberal use of redundancy, eliminating the need for

an abort capability. Since the high mass associated with the high margins of this approach impacts the objective, *Maximize Payload Capacity*, a different alternative might address the same crew safety objective by combining a lighter, less reliable system with an effective crew abort capability. For these two alternatives, significantly different analyses would need to be performed to quantify the probability $P(LOC)$ of accomplishing the crew safety performance measure. In the first case, $P(LOC)$ is directly related to system reliability. In the second case, reliability analysis plays a significant part, but additional analysis is needed to quantify abort effectiveness, which involves analysis of system responsiveness to the failure, and survivability given the failure environment.




Figure 4-13. Example Launch Vehicle Trade Tree from ESAS

Performance Parameters

A **performance parameter** is any value needed to execute the models that quantify the performance measures. Unlike performance measures, which are the same for all alternatives, performance parameters typically vary among alternatives, i.e., a performance parameter that is defined for one alternative might not apply to another alternative.

Example performance parameters related to the performance objective of lofting X lbs into low Earth orbit (LEO) might include propellant type, propellant mass, engine type/specifications, throttle level, etc. Additionally, performance parameters also include relevant environmental characteristics such as meteorological conditions.

Performance parameters may be uncertain. Indeed, risk has its origins in performance parameter uncertainty, which propagates through the risk analysis, resulting in performance measure uncertainty.

4.7.1 Set the Analytical Framework

For a given alternative, the relationship between performance measures and the analyses needed to quantify them can be established and illustrated using a means objectives network (introduced in Section 4.5.2). Figure 4-14, adapted from [13], illustrates the idea. This figure traces Performance Parameter 1 through the risk analysis framework, showing how it is used by multiple risk analyses in multiple mission execution domains.

For example, Performance Parameter 1 is a direct input to a risk analysis in the Cost and Schedule mission execution domains (which have been combined in the figure for convenience). This analysis produces outputs that are used as inputs to two other Cost and Schedule risk analyses. One of these risk analyses produces a value for Performance Measure 1, whereas the other risk analysis produces an output that is needed by a risk analysis in the Safety mission execution domain. This Safety risk analysis ultimately supports quantification of Performance Measure **n**.

Each of the **m** performance parameters that defines Alternative **i** can be similarly traced through the risk analysis framework.

Figure 4-14 illustrates the need for coordination among the organizations conducting the analyses to assure that:

- There is an organization responsible for the quantification of each performance measure;
- The data requirements for every risk analysis are understood and the data sources and destinations have been identified;
- All data are traceable back through the risk analysis framework to the performance parameters of the analyzed alternative.

4.7.1.1 Configuration Control

It is important to maintain consistency over the definition of each analyzed alternative to ensure that all involved parties are working from a common data set. This is particularly true during the earlier phases of the program/project life cycle where designs may be evolving rapidly as decisions are made that narrow the trade space and extend it to higher levels of detail. It is also true when decisions are revisited, such as during requirements rebaselining (as discussed in Section 2.2), in which case the complete definition of the alternative may be distributed among various organizational units at different levels of the NASA hierarchy. In this case it is necessary for the organization at the level of the decision to be made to consolidate all relevant alternative data at its own level, as well as levels below, into a configuration managed data set.

Additionally, the risk analysis framework itself must be configuration controlled, in terms of the analyses (e.g., version number) and data pathways.

4.7.2 Choose the Analysis Methodologies Using a Graded Approach

The selection of the appropriate analytical methodologies depends both on the domain of execution gauged by means of the correspondingly identified performance measure(s), and by the level of depths and detail required for a robust deliberation and alternative selection, at the stage of activity development when the AoA process is executed. The general principle for taking this into account is the application of a graded approach in methodology selection.




Figure 4-14. Risk Analysis Framework (Alternative Specific)

The spectrum of analysis disciplines involved in the risk analysis of alternatives is as broad as the spectrum of performance measures that are apt for gaging the performance of the products of an activity or project, spanning the activity execution domains: safety, technical, physical security, cybersecurity, cost, and schedule, plus any others that may be relevant (e.g., those listed in Figure 4-3). It is not the intent of this handbook to provide detailed guidance on the conduct of domain-specific analyses. Such guidance is available in domain-specific documents like the NASA Cost Estimating Handbook [14], the NASA Systems Engineering Handbook [11], and the NASA Probabilistic Risk Assessment Procedures Guide [15].

Depending on activity or project scale, and stage of activity development when the AoA is executed, etc., different levels of analysis are appropriate. The rigor of analysis should be enough to assess compliance with imposed constraints and support selection between alternatives. Iteration is to be expected as part of the analysis process, but as a general rule of thumb, the rigor of analysis should increase with the progress made in the level of definition of the alternative solutions being evaluated for the activity or project of concern.

If RIDM is being applied for identification of risk-control solutions in the course of execution of an activity involving the development and implementation of a solution or a project involving the formal design of a system, the application of the AoA may, at least in theory, happen at any of the successive activity stages or program/project life-cycle phases. Accordingly, the level of detail

required in the analysis may vary to reflect the need to reach a robust decision at that particular stage of the solution or product design development.

For any and all of the RIDM types identified in Figure 2-8 (Activity-Planning, Activity-Rebaseline, and Activity-Execution), the level of rigor in the analysis should also increase with the importance of the scenario being evaluated. Regardless of the time during the life cycle, certain scenarios will not be as important as others in affecting the performance measures that can be achieved for a given alternative. Scenarios that can be shown to have very low likelihood of occurrence and/or very low impacts on all the mission execution domains do not have to be evaluated using a rigorous simulation methodology or a full-blown accounting of the uncertainties. A point-estimate analysis using reasonably conservative simulation models and input parameter values should be sufficient for the evaluation of such scenarios.

The RIDM process does not imply a need for a whole new set of analyses. In general, some of the necessary analyses will already be planned or implemented as part of the systems engineering, cost estimating, and safety and mission assurance (S&MA) activities. Risk analysis for RIDM should take maximum advantage of existing activities, while also influencing them as needed in order to produce results that address objectives, at an appropriate level of rigor to support robust decision making.

The details of the graded approach for Activity-Planning RIDM and Activity Rebaseline RIDM will be covered later in this chapter within Section 4.11. On the other hand, because the analyses performed for RIDM during activity execution are conducted as part of CRM, the graded approach aspects of these analyses are covered in Chapter 5.

4.7.3 Conduct the Risk Analysis

Once the risk analysis framework is established and risk analysis methods determined, performance measures can be quantified. As discussed previously, however, this may be part of an iterative process of successive analysis refinement driven by stakeholder and decision-maker needs (see Part 3 of the RIDM process).

4.7.3.1 Probabilistic Modeling of Performance

If there were no uncertainty, the question of performance assessment would be one of quantifying point value performance measures for each decision alternative. In the real world, however, uncertainty is unavoidable, and the consequences of selecting a particular decision alternative cannot be known with absolute precision. When the decision involves a course of action there is uncertainty in the unfolding of events, however well planned, that can affect the achievement of objectives. Budgets can shift, overruns can occur, technology development activities can encounter unforeseen phenomena (and often do). Even when the outcome is realized, uncertainty will still remain. Reliability and safety cannot be known absolutely, given finite testing and operational experience. The limits of phenomenological variability in system performance can likewise not be known absolutely nor can the range of conditions under which a system will have to operate. All this is especially true at NASA, which operates on the cutting edge of scientific understanding and technological capability.

For decision making under uncertainty, risk analysis is necessary, in which uncertainties in the values of each alternative's performance parameters are identified and propagated through the analysis to produce uncertain performance measures. Moreover, since performance measures might not be independent, correlation must be considered. For example, given that labor tends to constitute a high fraction of the overall cost of many NASA activities, cost and schedule tend to be highly correlated. High costs tend to be associated with slipped schedules, whereas lower costs tend to be associated with on-time execution of the program/project plan.

One way to preserve correlations is to conduct all analysis within a common Monte Carlo “shell” that samples from the common set of uncertain performance parameters, propagates them through the suite of analyses, and collects the resulting performance measures as a vector of performance measure values [16]. As the Monte Carlo shell iterates, these performance measure vectors accumulate in accordance with the parent joint pdf that is defined over the entire set of performance measures. Figure 4-15 notionally illustrates the Monte Carlo sampling procedure as it would be applied to a single decision alternative (Decision Alternative *i*).




Figure 4-15. Risk Analysis Using a Monte Carlo Sampling Procedure

4.7.3.2 Treatment of Aleatory and Epistemic Uncertainties

Uncertainties are distinguished by two categorical groups: aleatory and epistemic [17], [18]. Aleatory uncertainties are random or stochastic in nature and cannot be reduced by obtaining more knowledge through testing or analysis. Examples include:

- The room-temperature properties of the materials used in a specific vehicle.

- The scenario(s) that will occur on a particular flight.

In the first case, there is random variability caused by the fact that two different material samples will not have the same exact properties even though they are fabricated in the same manner. In the second case, knowing the mean failure rates for all the components with a high degree of certainty will not tell us which random failures, if any, will actually occur during a particular flight.

On the other hand, epistemic uncertainties are not random in nature and can be reduced by obtaining more knowledge through testing and analysis. Examples include:

- The properties of a material at very high temperatures and pressures that are beyond the capability of an experimental apparatus to simulate.
- The mean failure rates of new-technology components that have not been exhaustively tested to the point of failure in flight environments.

In both cases, the uncertainty is caused by missing or incomplete knowledge or by limitations in the models used to make predictions.

A caveat about epistemic uncertainty should be noted here. When discussing epistemic uncertainty in the context of synthetic risk analysis, as in the preceding paragraphs of this subsection and in the paragraphs that follow, explicit treatment of epistemic uncertainty usually considers only the uncertainties that affect the assessment of “known risk.”¹² These are sometimes referred to as *known unknowns*. However, U/U risk is that part of the risk that eludes treatment in the synthetic analysis of risk. This distinction is similar to the distinction between *known unknowns* and *unknown unknowns*, which purported to have first been made by former Secretary of Defense Donald Rumsfeld in reference to the lack of evidence linking the government of Iraq with the supply of weapons of mass destruction to terrorist groups [19]. It is also noted that, while in other contexts the acronym “UU” is used to refer to the above mentioned “*unknown unknowns*,” in this handbook the acronym “U/U” refers instead to “*unknown and/or underappreciated*” risk.

It has become common in risk analysis to separate these two contributions to uncertainty by using the term *risk* to reflect the variability caused by aleatory uncertainties alone, and the term *risk uncertainty* or simply *uncertainty* to reflect the impreciseness of our knowledge of the risk caused by epistemic uncertainties alone. This distinction is useful for deciding whether additional research is worth the cost that it would entail, but is not always crucial for distinguishing between different architectural or design alternatives. Therefore, for purposes of the RIDM process, we speak only of uncertainties in the broad sense and do not distinguish between their aleatory and epistemic parts. However, the analyst always has the option of keeping aleatory and epistemic uncertainties separate from one another if he or she desires to do so, and in CRM where mitigation options are considered, this separation can be essential.

Further arguments about the relative advantages of combining aleatory and epistemic uncertainties in the *known unknown* category versus keeping them separate may be found in [20].

¹² This is typically done by characterizing model parameter values in terms of uncertainty distributions, i.e., characterizing the *parameter uncertainty* within a set of equations that are considered applicable to the modeling of risk for the activity in question.

4.7.3.3 Use of Qualitative Information in RIDM

As discussed in the preceding section, uncertainties in the forecasted performance measures are caused by uncertainties in the input performance parameters and in the models that are used to calculate the outcomes. These parameters and modeling uncertainties may be expressed in either quantitative or qualitative terms. If a parameter is fundamentally quantitative in nature, it is represented as having an uncertainty distribution that is expressed in terms of numerical values. For example, the date that a part is delivered is a quantitative performance parameter because it is defined in terms of the number of days between a reference date (e.g., the project's initiation) and the delivery date. The date has a discrete numerical distribution because it changes in 24-hour increments. Most performance parameters, such as the cost of the part or its failure rate, have continuous numerical distributions.

A performance parameter can often also be expressed in terms of a constructed scale that is qualitative in nature. For example, the technology readiness level (TRL) at the time of project initiation is a qualitative parameter because it is defined in terms of ranks that are based on non-numerical information. A TRL of 1, for example, is defined by terms such as: "basic principles observed and reported," "transition from scientific research to applied research," "essential characteristics and behaviors of systems and architectures," "descriptive tools are mathematical formulations or algorithms." Such terms are not amenable to quantitative analysis without a significant amount of interpretation on the part of the analysts.

While the performance parameter may be either quantitative or qualitative, the probability scale for the uncertainty distribution of the performance parameter is generally defined in a quantitative manner. The probability scale may be either continuous or discrete (although in most cases it is continuous). For example, a five-tiered discretization of probabilities on a logarithmic scale might be based on binning the probabilities into the following ranges: 10^{-5} to 10^{-4} for level 1, 10^{-4} to 10^{-3} for level 2, 10^{-3} to 10^{-2} for level 3, 10^{-2} to 10^{-1} for level 4, and 10^{-1} to 100 for level 5. It could be argued that the probability levels could also be defined in verbal terms such as "very unlikely to happen," "moderately likely to happen," and "very likely to happen." While these definitions are not numerical as stated, it is usually possible to ascertain the numerical ranges that the analyst has in mind when making these assignments. Thus, the probability should be relatable to a quantitative scale.

Various types of quantitative and qualitative uncertainty distributions for the input parameters and conditions are shown in Figure 4-16. Three of these (the top left and right charts and the lower right chart within the first bracket) are types of probability density functions, whereas the fourth chart (lower left) is a form of a complementary cumulative distribution function (CCDF). Either form of distribution (density form or cumulative form) may be used to express uncertainty. The choice is governed by whichever is the easier to construct, based on the content of the uncertainty information.

As depicted in the figure, the values of the output performance measures, as opposed to the values of the input performance parameters, are always quantitative in that they are defined in terms of numerical metrics. The output uncertainty distributions are expressed in the form of a histogram

representation of output values obtained from Monte Carlo sampling of the input values and conditions.




Figure 4-16. Uncertain Performance Parameters Leading to Performance Measure Histograms

Because the numerically based models are set up to accept numerical inputs, execution of the models for calculating the output performance measures is in general easier if all the performance parameters are defined in terms of quantitative scales, whether continuous or discrete. Caution should be used where one or more of the inputs are defined in terms of a qualitative, or constructed, scale. In these cases, the calculation of the performance measures may require that different models be used depending on the rank of the qualitative input. For example, the initial TRL for an engine might depend upon whether it can be made out of aluminum or has to be made out of beryllium. In this case, an aluminum engine has a higher TRL than a beryllium engine because the former is considered a heritage engine and the latter a developmental engine. On the other hand, a beryllium engine has the potential for higher thrust because it can run at higher temperatures. The model for calculating performance measures such as engine start-up reliability, peak thrust, launch date, and project cost would likely be different for an aluminum engine than for a beryllium engine.

4.7.3.4 Evaluation of U/U Risk

U/U risk is accounted for as a margin that is applied to the known risk, based on historical precedent with similar activities and informed by an assessment of leading indicator values for those leading indicators known to correlate with U/U risk. Figure 4-17 illustrates the application of a U/U risk margin to a CCDF that has been calculated for known risk. As the figure suggests,

U/U risk can be a significant portion of the total risk, especially for novel activities. The development of U/U risk margin and its relationship to leading indicators is discussed further in Appendix H.




Figure 4-17. Application of Risk Margin to the Known Risk to Account for U/U Risk

4.7.3.5 Risk Analysis Support of Robust Decision Making

Because the purpose of risk analysis in RIDM is to support decision making, the adequacy of the analysis methods must be determined in that context. The goal is a robust decision, where the decision-maker is confident that the selected decision alternative is actually the best one, given the state of knowledge at the time. This requires the risk analysis to be rigorous enough to discriminate between alternatives, especially for those performance measures that are determinative to the decision.

Figure 4-18 illustrates two hypothetical situations, both of which involve a decision situation having just one performance measure of significance. The graph on the left side of the figure shows a situation where Alternative 2 is clearly better than Alternative 1 because the bulk of its pdf is to the left of Alternative 1's pdf. Thus, the decision to select Alternative 2 is robust because there is high probability that a random sample from Alternative 1's pdf would perform better than a random sample from Alternative 2's pdf. In contrast, the graph on the right side of the figure shows a situation where the mean value of Alternative 1's performance measure is better than the mean value of Alternative 2's, but their pdfs overlap to a degree that prevents the decision to select Alternative 1 from being robust; that is, unless the pdfs for Alternatives 1 and 2 are highly correlated, there is a significant probability that Alternative 2 is actually better. The issue of correlated pdfs will be taken up later in this section.

For decisions involving multiple objectives and performance measures, it is not always possible to identify *a priori* which measures will be determinative to the decision and which will only be marginally influential. It is possible that some performance measures would require extensive analysis in order to distinguish between alternatives, even though the distinction would ultimately

not be material to the decision. Consequently, the need for additional analysis for the purpose of making such distinctions comes from the deliberators and the decision-maker, as they deliberate the merits and drawbacks of the alternatives. The judgment of whether uncertainty reduction would clarify a distinction between contending decision alternatives is theirs to make; if it would be beneficial and if additional analysis is practical and effective towards that purpose, then the risk analysis is iterated and the results are updated accordingly.




Figure 4-18. Robustness and Uncertainty

4.7.3.6 Sequential Analysis and Downselection

While the ultimate selection of any given alternative rests squarely with the decision maker, he or she may delegate preliminary downselection authority to a local proxy decision-maker, in order to reduce the number of contending alternatives as early as practical in the decision-making process. There is no formula for downselection; it is an art whose practice benefits from experience. In general, it is prudent to continuously screen the alternatives throughout the process.

Downselection often involves the conduct of sequential analyses, each of which is followed by a pruning of alternatives. In this way, alternatives that are clearly unfavorable due to their performance on one (or few) performance measures can be eliminated from further analysis once those values are quantified. To optimize the process of analysis and downselection, within the constraints of the analytical dependencies established by the risk analysis framework set in the previous step, it may be prudent to order the conduct of domain-specific analyses in a manner that exploits the potential for pruning alternatives prior to forwarding them for additional analysis. There is no hard rule for an optimal ordering; it depends on the specific decision being made, the alternatives compiled, and the analysis methods employed. It is recommended that opportunities for sequential analysis and downselection be looked for as alternatives are analyzed, and that the

ordering of analyses be adjusted as appropriate to facilitate downselection, depending on which performance measures can be used as a basis for pruning.

It is important to document the basis for eliminating alternatives from further consideration at the time they are eliminated. Two such bases that were discussed in Section 4.6.2 are infeasibility and dominance. Additional discussion of downselection is presented in Section 4.7.3.6. In all cases, downselection requires active collaboration among the risk analysts, the deliberators and the decision maker, as it is not the role of the risk analysts to eliminate alternatives by his/her own sole judgment, except on the grounds of clear infeasibility. Sequential downselection, like all decision making, must be done in the context of stakeholder values and decision-maker responsibility and accountability.

4.7.3.7 Model Uncertainty and Sensitivity Studies

As is the case with all modeling activities, risk modeling typically entails a degree of model uncertainty to the extent that there is a lack of correspondence between the model and the alternative being modeled. Model uncertainty is a form of epistemic uncertainty, but rather than being caused by uncertainty in the parameters that are input to the model (the subject of Section 4.7.3.2 through 4.7.3.4), the source of the uncertainty stems from limitations in the accuracy and/or applicability of the model itself.

The usual approach to assuring that decisions are robust with respect to model uncertainty is to conduct sensitivity studies over ranges of credible model forms and/or parameter values. Sensitivity studies are particularly pertinent for models that produce point value performance measure results, even when the performance measure is known to be uncertain. In these cases, it is valuable to determine the sensitivity of the decision to bounding variations in the risk model assumptions. Figure 4-19 notionally presents the results of such a study. It shows how the preferred alternative varies as a function of assumptions about contractor support cost rate and payload mass. For example, if the contractor support cost rate is 120 and the payload mass is 18, then Alternative A is the preferred alternative. If, however, the assumed payload mass is 4, then Alternative B is preferable. More generally, if “Alternative B” is preferred for all reasonable values of contractor support cost rate and payload mass, then the decision is robust in favor of Alternative B (with respect to these parameters), without the need for additional rigor in determining the actual contractor support cost rate or payload mass. Likewise, if the reasonable range of these parameters falls entirely within the region “Alternative A,” then the decision is robust for Alternative A. Only when the reasonable range of values straddles more than one region is more rigorous characterization of contractor support cost and payload mass needed for robust decision making.

4.7.3.8 Analysis Outputs

Like the variation in risk analysis methods, the analysis results presentation for RIDM may vary, depending on the nature of the problem being evaluated. Consequently, there can be no one standard analysis output. Instead, the results are tailored to the problem and the needs of the deliberation process. Consideration should be given for providing a variety of results, including:

- Scenario descriptions

- Performance measure pdfs and statistics
- Risk results (e.g., risk of not meeting imposed constraints)
- Uncertainty analyses and sensitivity studies




Figure 4-19. Notional Depiction of Decision Sensitivity to Input Parameters

It is important to note that the risk analysis results are expected to mature and evolve as the analysis iterates with the participation of the stakeholders and the decision-maker. This is not only due to increasing rigor of analysis as the stakeholders and the decision-maker strive for decision robustness. Additionally, as they establish firm performance targets, it becomes possible to evaluate the analysis results in the context of those targets. For example, prior to the development of performance targets, it is not possible to construct a risk list that is keyed to the performance measures (except with respect to imposed constraints, which are firmly established prior to analysis).

4.7.3.9 Assessing the Credibility of the Risk Analysis Results

In a risk-informed decision environment, risk analysis is just one element of the decision-making process, and its influence on the decision is directly proportional to the regard in which it is held by the deliberators. A well-done risk analysis whose merits are underappreciated might not influence a decision significantly, resulting in a lost opportunity to use the available information to the best advantage. Conversely, an inferior risk analysis held in overly high regard has the ability to produce poor decisions by distorting the perceived capabilities of the analyzed alternatives. In

order to address this potential, an evaluation of the credibility of the risk analysis is warranted prior to deliberating the actual results.

NASA-STD-7009, Standard for Models and Simulations [21], provides the decision maker with an assessment of the modeling and simulation (M&S) results against key factors that:

- Contribute to a decision-maker’s assessment of credibility and
- Are sensibly assessed on a graduated credibility assessment scale (CAS).

Table 4-III (which reproduces NASA-STD-7009 Table 1) presents a high-level summary of the evaluation criteria. These are explained in greater detail in Section B.3 of the standard. Table 4-III by itself is not intended to be the whole story regarding credibility assessments. Rather, it is to be used in conjunction with the detailed level definitions in the standard.

Table 4-III. Key Aspects of Modeling and Simulation Credibility Assessment Levels

Level	Verification	Validation	Input Pedigree	Results Uncertainty	Results Robustness	Use History	M&S Management	People Qualifications
4	Numerical errors small for all important features.	Results agree with real-world data.	Input data agree with real-world data.	Non-deterministic & numerical analysis.	Sensitivity known for most parameters; key sensitivities identified.	De facto standard.	Continual process improvement.	Extensive experience in and use of recommended practices for this particular M&S.
3	Formal numerical error estimation.	Results agree with experimental data for problems of interest.	Input data agree with experimental data for problems of interest.	Non-deterministic analysis.	Sensitivity known for many parameters.	Previous predictions were later validated by mission data.	Predictable process.	Advanced degree or extensive M&S experience, and recommended practice knowledge.
2	Unit and regression testing of key features.	Results agree with experimental data or other M&S on unit problems.	Input data traceable to formal documentation.	Deterministic analysis or expert opinion.	Sensitivity known for a few parameters.	Used before for critical decisions.	Established process.	Formal M&S training and experience, and recommended practice training.
1	Conceptual and mathematical models verified.	Conceptual and mathematical models agree with simple referents.	Input data traceable to informal documentation.	Qualitative estimates.	Qualitative estimates.	Passes simple tests.	Managed process.	Engineering or science degree.
0	Insufficient evidence.	Insufficient evidence.	Insufficient evidence.	Insufficient evidence.	Insufficient evidence.	Insufficient evidence.	Insufficient evidence.	Insufficient evidence.
	M&S Development		M&S Operations			Supporting Evidence		

To assist in the application of the evaluation criteria set forth in NASA-STD-7009, Figure 4-20 presents a matrix indicating how the “level” of analysis identified in Figure 4-19 relates to various estimation methods. Each of the estimation methods in Figure 4-20 will be discussed in Section 5.3.1.1.

The assessment of credibility levels per Table 4-III is provided along with the results of the risk analysis to render perspective on the robustness of the evidence used in the technical basis for deliberation.

<u>Cost Estimating Method</u>				<u>Technical Estimating Method</u>				<u>Safety, Reliability, & Operations Estimating Method</u>			
<u>Level</u>	Analogy	Parametric	Engineering Build Up	First-Order	Detailed Simulation	Operating Experience	Testing	Similarity	First-Order Parametric	Detailed Logic Modeling	Statistical
0											
1	X			X				X			
2		X			X				X		
3			X			X				X	
4							X				X

Figure 4-20. Analysis Level Matrix

4.8 Details of Step 4 (Part 2), Develop the Technical Basis for Deliberation

As discussed in Section 4.4.2 and portrayed in Figure 4-7, the TBfD (see Appendix F) specifies the minimum information needed to risk-inform the selection of a decision alternative. The content of the TBfD is driven by the question, "What information do the deliberators and decision-makers need in order for their decision process to be fully risk-informed?"

Graphical tools are recommended, in addition to tabular data, as a means of communicating risk results. At this point in the process, the imposed constraints are the only reference points with respect to which shortfalls can be determined, so they are the only things "at risk" so far. Figure 4-21 presents a notional color-coded chart of imposed constraint risk. In the figure, Alternative 7 is relatively low in known risk for every listed performance measure (i.e., those with imposed constraints on the allowable values), as well as for all constrained performance measures collectively (the "Total" column). Alternatives 12 and 3 have a mix of performance measure risks, some of which are high, resulting in a high risk of failing to meet one or more imposed constraints.

To assist the deliberators and decision-maker in focusing on the most promising alternatives, with an awareness of the relative risks to imposed constraints, the imposed constraints risk matrix has been:

- Sorted by risk, with the least risky alternatives at the top; and
- Colored on a relative basis from low risk (the blue-green end of the spectrum) to high risk (the orange-red end of the spectrum).

When presenting the performance measure pdfs themselves, "band-aid" charts can be used, which show the mean, 5th percentile, and 95th percentile values (and often the median as well). Figure 4-22 shows a notional example of a band-aid chart. Unlike the imposed constraints matrix, which includes only those performance measures that have imposed constraints, band-aid charts can be made for every performance measure in the risk analysis, thereby giving a comprehensive picture of the analyzed performance of each alternative.

When using charts such as the band-aid chart of Figure 4-22, it is important to know the degree of correlation among the different alternatives. For example, in the figure, the pdfs of Alternative 1 and Alternative 2 overlap to an extent that it may seem that the chances of either one having the

higher performance are about the same. Indeed, this is true if the pdfs are independent. However, if they are correlated, then it might not be the case. For example, suppose the alternatives are identical except for some small design difference that slightly increases the value of Performance Measure X for Alternative 2. Then, although the performance of both alternatives is uncertain, the performance difference between them is known and constant.

Alternative	Imposed Constraints					
	PM ₁	PM ₂	PM ₃	...	PM _n	Total
	Constraint (> C ₁)	Constraint (< C ₂)	Constraint (> C ₃)		Constraint (< C _n)	
7	0.7%*	0.09%	0.04%		0.1%	0.9%
1	2%	0.8%	0.2%		0.01%	4%
12	0.4%	5%	20%		0.1%	27%
...				...		
3	15%	0.9%	8%		0.01%	56%

*The probability of not meeting the imposed constraint

Figure 4-21. Notional Imposed Constraints Risk Matrix




Figure 4-22. Notional Band Aid Chart for Performance Measure X

A direct representation of the difference between design alternatives, including the associated uncertainty, can supplement the information provided by band-aid charts, allowing for a better

ability to make comparisons under uncertainty. A possible representation is shown in Figure 4-23 [22]. The figure shows performance measure pdfs for two alternatives whose performance measure values are correlated. A third, dotted, curve shows the pdf of the performance difference between the two alternatives. This curve indicates that despite the significant overlap between the two performance measure pdfs, Alternative 2 is unequivocally superior to Alternative 1, at least for the performance measure shown.




Figure 4-23. Comparison of Uncertainty Distributions

4.9 Details of Step 5 (Part 3), Deliberate

As discussed in Section 4.4.3 and portrayed in Figure 4-8, the deliberation step brings together the analytical information developed in Part 2 of the RIDM process, and documented in the TBfD, with risk leadership and risk posture criteria that are the guiding principle by which decision makers ultimately can risk-inform their selection alternative decisions.

The RIDM process invests the decision-maker with the authority and responsibility for critical decisions. While ultimate responsibility for alternative selection rests with the decision-maker, alternative evaluation can be performed within a number of deliberation forums which may be held before the final selection is made. As partial decisions or “down-selects” may be made at any one of these deliberation forums, they are routinely structured around a team organizational structure identified by the decision-maker. It is important to have a team with broad-based expertise to perform sufficient analysis to support a recommendation or decision. At the top of the structure may be the decision-maker or a deliberation lead appointed by the decision-maker. If a deliberation lead is appointed this individual should be an experienced manager, preferably one with an analytical background.

As the deliberation process brings together multiple sources of information as well as constraints and performance objectives for the alternatives being considered, it may involve relatively complex interactions and iterations. These are included in sub-steps discussed in the remainder of the section.

4.9.1 [Convene a Deliberation Forum](#)

Deliberation forums address the major aspects of the deliberation process leading to a decision. The use of these forums helps ensure that a responsible person leads each important area of analysis. The focus of these forums will vary with the type of study.

Depending on circumstances, forums can be split (e.g., into separate safety, security, and technical), or functions can be combined (e.g., cost and schedule), or entirely new forums can be created (e.g., test, requirements or stakeholder). The final choice of forum structure belongs to the decision-maker. At a minimum, the forums should mirror the major aspects of the study. Thus, the creation of forums offers an important early opportunity to contemplate the effort processes and goals. Every forum must have enough members to achieve a “critical mass” of knowledge, interest and motivation. Typically, a small group with critical mass is more productive than a larger group with critical mass. This suggests starting with a small forum and adding members as necessary.

Members of a deliberation forum should ideally be selected based on their qualifications. Consideration should be given to those with relevant experience, knowledge, and interest in the subject matter. These individuals are frequently referred to as SMEs. In some cases, they have an organizational charter to support the process and in other cases they participate because they are heavily invested in the outcome of the deliberation. When the most qualified are not available, the next most qualified should be sought.

People with diverse viewpoints on controversial issues should also be enlisted to participate in deliberations. They should represent the diversity of stakeholder interests. Partisans, by their nature, will defend their ideas and detect flaws in the ideas of their competition. This allows issues to be raised and resolved early that might otherwise lie in wait. A formal tracking system should be employed throughout the process to track items to closure.

Additional information on deliberative processes can be found in [23].

4.9.2 [Develop Performance Targets and Risk Tolerances for Individual Performance Measures](#)

In Section 0 and its subsections, the concept of performance markers and performance marker risk tolerances was introduced. Two types of performance markers were defined: a *performance constraint* representing pre-set conditions assumed to be non-negotiable, and a *performance target* corresponding to a performance goal identified in accordance with stakeholders’ and/or decision-maker’s preferences, but not necessarily set in rigid terms. It may be useful for the reader to review this material before proceeding.

The evaluations described below establish a **risk-normalized performance target (RPT)** for each performance dimension for each of the decision alternatives being considered. A risk-normalized performance target is a performance measure value set at a particular percentile of the performance measure probability distribution, so as to anchor the decision-maker’s perspective to that performance measure value as if it would be the activity / project goal to achieve, were the decision maker to select that alternative. For a given performance measure, the percentile value is set at the same level for all alternatives, so that the probability of failing to meet the different alternative target values is the same across alternatives.

The inputs to RPT development are:

- The performance measure pdfs or CDFs for each decision alternative;
- An ordering of the performance measures; and
- A risk tolerance for each performance measure, expressed as a percentile value.

4.9.2.1 Risk-Normalized Evaluation of Candidate Alternatives without Pre-Conceived Performance Constraints or Targets

The *risk-normalized* RIDM evaluation of candidate alternatives with respect to a performance measure uses the relation between performance and risk tolerance by setting an RPT that reflects the organization's overall risk posture and the decision maker's risk attitude with respect to each of the performance measures being compared for the competing alternatives being evaluated. In this mode of evaluation, illustrated by Figure 4-24, the decision maker sets a risk value as his/her "level of tolerance" for known risk and the competing alternatives are evaluated in terms of the performance measure values that their uncertainty distribution functions indicate as achievable, as a minimum, at that risk tolerance level.




Figure 4-24. Evaluation of Risk of Alternatives with Respect to a Risk Tolerance Level for Known Risk

This mode of deliberation can be directly applied when a *performance constraint* has not been pre-established, either because the performance trade space is wide-open, or because it has already been determined that the alternatives being considered satisfy any existing performance constraints with ample margins, so that any further performance goal can be identified in the form of a *performance target* as part of the AoA deliberation and selection process itself. In this type of

situation, the deliberation sub-steps that can be followed, in relation to each of the performance measures being considered, may be defined as follows:

- a. A risk tolerance value is selected by the decision maker as an initial deliberation target with respect to the performance measure PMX of concern.
- b. Using performance measure CDF information as in Figure 4-24, the performance levels of each alternative corresponding to the risk tolerance value set in step a) are identified, and alternatives are ranked according to the respective performance levels. Once all performance measures are considered and alternatives ranked, performance target values consistent with the organization risk posture and preferred risk tolerance limits are identified.

For the example case shown in Figure 4-24 Alternative 2 is preferable as it presents a better RPT value than Alternative 1 (as “direction of goodness” is from lower to higher PM values).

4.9.2.2 Risk-Normalized Evaluation of Candidate Alternatives with Pre-Established Performance Constraints or Targets

In some cases, the deliberation and selection process must account for the existence, in a given performance measure dimension, of a *performance constraint*, or of a *performance target* associated with a constraint via a pre-established *design margin*, as hypothetically assumed in the illustrations given previously in Figure 4-4 and Figure 3-9.

In these situations, if the decision maker has identified a *risk tolerance level* that they would like to use for a risk-normalized comparison of alternatives (as outlined in the preceding discussion), the deliberation process may be in a situation such as the one depicted in Figure 4-25.




Figure 4-25. Consideration of Performance Constraints in Risk-Normalized Evaluation of Alternatives

Figure 4-25 represents a hypothetical situation where:

- a) Three competing alternatives are being evaluated and compared with respect to a performance measure (PM);
- b) A performance constraint (PC) for PM pre-exists the deliberation (the term constraint is used here also for a pre-existing “target” derived as “constraint + design-margin”);
- c) The decision maker has, independently from the above, identified a risk tolerance level (RTL) that they judge appropriate for the risk-normalized evaluation of alternatives and for the establishment of *risk-normalized performance targets*.

The figure shows that, if the CDFs assessed for the three alternatives are as indicated, it is not possible for Alternative 1 to satisfy at the same time the target RTL and the pre-established performance constraint limits. Under such circumstances the decision maker, or the deliberation forum to which they may delegate the task, may choose one of the following ways to proceed with the risk-normalized evaluation approach that has been discussed earlier:

- A. **Cull Alternatives:** Eliminate alternatives that cannot satisfy both RTL and performance constraint limits, and continue with the risk-normalized evaluation of remaining alternatives. In the situation described by Figure 4-25, this would mean to eliminate Alternative 1 from further consideration, and continue the evaluation of Alternatives 2 and 3 (which of course would be based not only on the performance measure, but on all the other AoA relevant performance measures).
- B. **Relax Risk Tolerance Level:** Increase the RTL to a value that permits continuing the evaluation of all alternatives on the table. For the case described by the figure, to achieve this the RTL would have to be elevated to a value greater than the Alternative 1 CDF value corresponding to the performance constraint value.
- C. **Renegotiate the Performance Constraint:** Discuss with the parties that have originated the constraint whether its value may be relaxed. This of course may or may not be possible. If it is, it may be so if the decision maker is contemporarily willing to relax the RTL value, as per B above. In the case depicted by the figure, to keep Alternative 1 in the deliberation would require moving the performance constraint value to the left of the abscissa point labeled in the figure as the “Risk-normalized PM Value for A1.”

Once any potential conflicts between decision maker’s risk tolerance preferences and pre-established performance constraints are identified and resolved via one of the means described above, all de-conflicted alternatives can continue to be deliberated upon according to the RIDM risk-normalized evaluation processes.

4.9.3 Sequentially Establish Risk-Normalized Performance Targets for All Performance Measures

For each alternative, each risk-normalized performance target (RPT) is established by sequentially determining, based on a selected performance measure ordering, the value that corresponds to the stated risk tolerance, conditional on meeting previously-defined RPTs. This value becomes the

RPT for the current performance measure, and the process is repeated until all RPTs have been established for all performance measures.

Figure 4-26 illustrates the process. In the figure, there are only two performance measures, PM_1 and PM_2 . If, for example, PM_1 is mass-to-orbit (MTO) and PM_2 is cost, then, the risk analysis results can be shown as a scatter plot on the MTO-Cost plane (see Figure 4-26a), where each point represents the output from a single iteration of a Monte Carlo shell. If the ordering of the performance measures is mass-to-orbit first and cost second, mass-to-orbit would be the first performance measure to have a performance target established for it (see Figure 4-26b). This is done by determining the value of mass-to-orbit whose probability of exceedance equals the defined risk tolerance. That value becomes the mass-to-orbit performance target.¹³ The process is repeated for cost, *conditional on the mass-to-orbit performance target being met*. Thus, the points on the scatter plot that exceed the mass-to-orbit performance target have been removed from consideration and the cost performance target is established solely on the basis of the remaining data (see Figure 4-26c). The result is a set of performance targets for the mass-to-orbit and cost performance measures that reflects the risk tolerances of the deliberators and decision-maker (see Figure 4-26d). This procedure can be extended to any number of performance measures.




Figure 4-26. Establishing Risk-Normalized Performance Targets

¹³ If the “direction of goodness” of the performance measure were reversed, the RPT would be at the value whose probability of exceedance equals one minus the risk tolerance.

In general, different decision alternatives will have different risk-normalized performance targets, but the probability of meeting each RPT will be the same (namely, one minus the risk tolerance of that performance measure), given that prior RPTs in the performance measure ordering have been met:

$$\begin{aligned} P(\text{RPT } i \text{ is met}) &= 1 - PM_i \text{ Risk Tolerance} \\ &= 1 - P(\text{RPT } i \text{ is unmet} \mid \text{RPT } j < i \text{ are met}) \end{aligned}$$

Moreover, the probability of meeting all RPTs is identical for all alternatives, and is calculated as:

$$P(\text{All RPTs Met}) = \prod_{i=1}^{\# PMs} (1 - PM_i \text{ Risk Tolerance})$$

4.9.3.1 Establishing Risk Tolerance Levels for the Performance Measures

The RIDM process calls for the specification of a risk tolerance for each performance measure, along with a performance measure ordering, as the basis for RPT development. These risk tolerance values have the following properties:

- The risk tolerance for a given performance measure is the same across all alternatives, and
- Risk tolerance may vary across performance measures, in accordance with the stakeholders' and decision-maker's attitudes towards risk for each performance measure.

Risk tolerances, and their associated RPTs, play multipurpose roles within the RIDM process:

- Uniform risk tolerance across alternatives normalizes activity risk, enabling deliberations to take place that focus on performance capabilities on a risk-normalized basis.
- The risk tolerances that are established during the RIDM process indicate the levels of acceptable initial risk that the CRM process sets out to manage during implementation. (Note: The *actual* initial risk is not established until performance requirements are agreed upon as part of the overall systems engineering process, and not explicitly addressed until the CRM process is initialized. More information on CRM initialization can be found in Section 5.1.)
- RPTs based on risk tolerance enable point value comparison of alternatives in a way that is also appropriate for situations that involve pre-established thresholds (e.g., imposed performance constraints). By comparing a performance target to a threshold, it is immediately clear whether or not the risk of crossing the threshold is within the established risk tolerance. In contrast, if a value such as the distribution mean were used to define RPTs, the risk with respect to a given threshold would not be apparent.

Issues to consider when establishing risk tolerances include:

- *Relationship to imposed Performance Constraints* – In general, deliberators have a low tolerance for noncompliance with imposed constraints. Imposed constraints are akin to the success criteria for top-level objectives; if imposed constraints are not met, then objectives are not met and the endeavor fails. By establishing a correspondingly low risk tolerance on

performance measures that have imposed constraints, stakeholders and decision-makers have assurance that if an alternative's RPTs exceed the associated imposed constraints, there is a high likelihood of program/project success.

- *High-priority objectives* – It is expected that deliberators will also have a low risk tolerance for objectives that have high priority, but for which imposed constraints have not been set. The lack of an imposed constraint on a performance measure does not necessarily mean that the objective is of less importance; it may just mean that there is no well-defined threshold that defines success. This could be the case when dealing with quantities of data, sample return mass capabilities, or operational lifetimes. It is generally the case for life safety, for which it is difficult to establish a constraint *a priori*, but which is nevertheless always among NASA's top priorities.
- *Low-priority objectives and/or "stretch goals"* – Some decision situations might involve objectives that are not crucial to program/project success, but which provide an opportunity to take risks in an effort to achieve high performance. Technology development is often in this category, at least when removed from a project's critical path. In this case, a high risk tolerance could be appropriate, resulting in RPTs that suggest the alternatives' performance potentials rather than their established capabilities.
- *Rebaselining issues* – Requirements on some performance measures might be seen as difficult to rebaseline. For these performance measures, deliberators might establish a low risk tolerance in order to reduce the possibility of having to rebaseline.

Risk tolerance values are up to the deliberators and decision maker, and are subject to adjustment as deliberation proceeds, opinions mature, and sensitivity excursions are explored. In particular, it is recommended that sensitivity excursions be explored over a reasonable range of risk tolerances, not only for the purpose of making a decision that is robust with respect to different risk tolerances, but also in order to find an appropriate balance between program/project risk and the performance that is specified by the RPTs.

4.9.3.2 Ordering the Performance Measures

Because of possible correlations between performance measures, RPTs are developed sequentially. As discussed earlier, RPTs are defined at the value of a performance measure that corresponds to the defined risk tolerance, conditional on meeting previously defined RPTs. In general, RPT values depend on the order in which they are developed.

Qualitatively, the effect that performance measure order has on RPT values is as follows:

- If performance measures are independent, then the order is immaterial and the RPTs will be set at the defined risk tolerances of the performance measures' marginal pdfs.
- If performance measures are positively correlated in terms of their directions of goodness, then the RPTs that lag in the ordering will be set at higher levels of performance than would be suggested by their marginal pdfs alone. This is because lagging performance measures will have already been conditioned on good performance with respect to leading

performance measures. This, in turn, will condition the lagging performance measures on good performance, too, due to the correlation.

- If performance measures are negatively correlated in terms of their directions of goodness, then the RPTs that lag in the ordering will be set at lower levels of performance than would be suggested by their marginal pdfs alone. Figure 4-26 shows this phenomenon. In Figure 4-26c, the PM₂ RPT is set at a slightly lower performance than it would have been if the data points that exceed the PM₁ RPT were not “conditioned out.”
- The lower the risk tolerance, the lower the effect of conditioning on subsequent RPTs. This is simply because the quantity of data that is “conditioned out” is directly proportional to risk tolerance.

These general effects of performance measure ordering on RPTs suggest the following ordering heuristics:

- Order performance measures from low risk tolerance to high risk tolerance. This assures a minimum of difference between the risk tolerances as defined on the conditioned pdfs versus the risk tolerances as applied to the marginal pdfs.
- Order performance measures in terms of the desire for specificity of the performance measure’s risk tolerances. For example, the RPT for the first performance measure in the ordering is precisely at its marginal pdf. As subsequent RPTs are set, dispersion can begin to accumulate as conditioning increases.

Once the RPTs are developed, each alternative can be compared to every other alternative in terms of their RPTs, with the deliberators’ understanding that the risk of not achieving the levels of performance given by the RPTs is the same across alternatives. Additionally, the RPTs can be compared to any imposed constraints to determine whether or not the possibility that they will not be satisfied is within the risk tolerance of the deliberators, and ultimately, the decision maker. Figure 4-27 notionally illustrates a set of RPTs for each of three competing alternatives. Note that Alternative A does not satisfy the imposed constraint on payload capability within the risk tolerance that has been established for that performance measure.

4.9.4 Pare Down the Contending Alternatives

After the performance targets have been generated, they are used to pare down the set of decision alternatives to those that are considered to be legitimate contenders for selection by the decision-maker. This part of the process is a continuation of the pruning activity begun in Step 2. At this point, however, the deliberators have the benefit of the TBfD and the identification of the RPTs, as well as the subjective, values-based input of the deliberators themselves. Rationales for elimination of non-contending alternatives include:

- **Infeasibility** – Performance targets are inconsistent with the imposed constraints. In this case, imposed constraints cannot be met within the risk tolerance of the decision-maker.
- **Dominance** – Other alternatives exist that permit the establishment of superior performance targets on every performance measure, and substantially superior

performance on some.¹⁴ In this case, an eliminated alternative may be feasible, but nonetheless is categorically inferior to one or more other alternatives.

- **Inferior Performance in Key Areas** – In general, in any decision involving multiple objectives, some objectives will be of greater importance to deliberators than others. Typically, important objectives include crew safety, mission success, payload capability, and data volume/quality. Alternatives that are markedly inferior in terms of the performance targets they permit to be established in key areas can be eliminated on that basis, in recognition of stakeholder and decision-maker values.




Figure 4-27. RPTs and Risk Tolerances for Three Alternatives

Section 4.7.3.6 discusses sequential analysis and downselection, in which non-contending alternatives are identified and eliminated in parallel with risk analysis, thereby reducing the analysis burden imposed by the decision-making process. Sequential analysis and downselection represents a graded approach to the identification of contending alternatives, and is another example of the iterative and collaborative nature of the RIDM process.

¹⁴ When eliminating alternatives on the basis of dominance, it is prudent to allow some flexibility for uncertainty considerations beyond those captured by the performance targets alone (discussed in the next subsection). Minor performance target shortfalls relative to other alternatives do not provide a strong rationale for elimination, absent a more detailed examination of performance uncertainty.

4.9.4.1 Uncertainty Considerations

The guidance above for identifying contending alternatives is primarily focused on comparisons of performance targets. This facilitates comparisons between alternatives (and against imposed constraints), and the elimination of non-contenders from further consideration. However, performance targets do not capture all potentially relevant aspects of performance, since they indicate the performance at only a single percentile of each performance measure pdf. Therefore, alternatives identified as contenders on the basis of their performance targets are further evaluated on the basis of additional uncertainty considerations relating to their performance at other percentiles of their performance measure pdfs. In particular, performance uncertainty may give rise to alternatives with the following characteristics:

- **They offer superior expected performance** – In many decision contexts (specifically, those in which the decision-maker is *risk neutral*¹⁵), the decision-maker's preference for an alternative with uncertain performance is equivalent to his or her preference for an alternative that performs at the mean value of the performance measure pdf. When this is the case, expected performance is valuable input to decision making, as it reduces the comparison of performance among alternatives to a comparison of point values.

However, in the presence of performance thresholds, over-reliance on expected performance in decision making has the potential to:

- Introduce potentially significant probabilities of falling short of imposed constraints, thereby putting objectives at risk, even when the mean value meets the imposed constraints
- Contribute to the development of derived requirements that have a significant probability of not being achievable

Since direction-setting, requirements-producing decisions at NASA typically involve performance thresholds, expected performance should be considered in conjunction with performance targets, to assure that the decision is properly risk informed.

- **They offer the potential for exceptionally high performance** – For a given performance measure pdf, the percentile value at the decision-maker's risk tolerance may be unexceptional relative to other contending alternatives. However, at higher risk tolerances, its performance may exceed that of other alternatives, to the extent that it becomes attractive relative to them. This may be the case even in the presence of inferior performance targets on the same, or different, performance measures.

An example of this is shown notionally in Figure 4-28. In this figure, Alternative 2's performance target is at a worse level of performance than Alternative 1's; however,

¹⁵ A risk-neutral decision maker is indifferent towards a decision between an alternative with a definite performance of *X*, versus an alternative having an uncertain performance whose mean value is *X*. In other words, a risk-neutral decision maker is neither disproportionately attracted to the possibility of exceptionally high performance (*risk seeking*) nor disproportionately averse to the possibility of exceptionally poor performance (*risk averse*).

Alternative 2 offers a possibility of performance that is beyond the potential of Alternative 1. In this case, stakeholders and decision-makers have several choices. They can:

- Choose Alternative 1 on the basis of superior performance at their risk tolerance;
- Choose Alternative 2 on the basis that its performance at their risk tolerance, though not the best, is acceptable, and that it also has the potential for far superior performance; or
- Set their risk tolerance such that the performance target for both alternatives are the same thus making this performance measure a non-discriminator between the two options.




Figure 4-28. An Example Uncertainty Consideration: The Potential for High Performance

In the second case, the decision-maker is accepting a higher program/project risk, which will lead to the development of more challenging requirements and increased CRM burden regardless of which alternative is selected.

- **They present a risk of exceptionally poor performance** – This situation is the reverse of the situation above. In this case, even though the likelihood of not meeting the performance target is within the decision-makers' risk tolerance, the consequences may be severe, rendering such an alternative potentially unattractive.

Another uncertainty consideration, which is addressed below in the discussion of the iterative nature of deliberation, is whether or not a performance measure's uncertainty can be effectively reduced, and whether or not the reduction would make a difference to the decision. This issue is mentioned here because wide pdfs can lead to poor performance targets relative to other alternatives, and it would be unfortunate to discard an alternative on this basis if additional analysis could be done to reduce uncertainty. Note that if two attractive alternatives present themselves and time and resources are available, it may be advantageous to proceed with, at least, partial prototyping (that is, prototyping of some of the critical components) of both to provide the necessary data for reducing key performance measure uncertainties such that a robust decision can be made.

4.9.4.2 Other Considerations

Depending on the decision situation and proposed alternatives, a variety of other risk-based, as well as non-risk-based, considerations may also be relevant. These include:

- **Sensitivity of the performance targets to variations in risk tolerance** – Performance targets are directly related to risk tolerance. Therefore, it is prudent for the deliberators to explore the effects of variations in the specified risk tolerances, to assure that the decision is robust to variations within a reasonable range of tolerances.
- **Risk disposition and handling considerations** – The risks that exist relative to performance targets are ultimately caused by undesirable scenarios that are identified and analyzed in the risk analysis. Because of the scope of risk analysis for RIDM (i.e., the necessity to analyze a broad range of alternatives), risk retirement strategies may not be fully developed in the analysis. Deliberators' expertise is therefore brought to bear on the relative risk-retirement burdens that different alternatives present. For example, deliberators might feel more secure accepting a technology development risk that they feel they can influence, rather than a materials availability risk they are powerless to control.
- **Institutional considerations** – Different alternatives may have different impacts on various NASA and non-NASA organizations and institutions. For example, one alternative might serve to maintain a particular in-house expertise, while another alternative might help maintain a regional economy. These broad-ranging issues are not necessarily captured in the performance measures, and yet they are of import to one or more stakeholders. The deliberation forum is the appropriate venue for raising such issues for formal consideration as part of the RIDM process.

4.9.4.3 Deliberation Is Iterative

As illustrated in Figure 4-8, deliberation is an iterative process that focuses in on a set of contending alternatives for consideration by the decision-maker. Iteration during deliberation has both qualitative and quantitative aspects:

- **Qualitative** – A deliberator may have a particular issue or concern that he or she wishes to reach closure on. This might require several rounds of deliberation as, for example, various subject matter experts are called in to provide expertise for resolution.
- **Quantitative** – One or more performance measures might be uncertain enough to significantly overlap, thereby inhibiting the ability to make a robust decision. Moreover, large uncertainties will, in general, produce poor performance targets, particularly when risk tolerance is low. Therefore, before a set of contending alternatives can be chosen, it is important that the deliberators are satisfied that particular uncertainties have been reduced to a level that is as low as reasonably achievable given the scope of the effort. It is expected that the risk analysis will be iterated, under the direction of the deliberators, to address their needs.

4.9.5 Communicating the Contending Alternatives to the Decision Maker

There comes a time in RIDM when the remaining alternatives all have positive attributes that make them attractive in some way and that make them all contenders. The next step is to find a way to clearly state for the decision-maker the advantages and disadvantages of each remaining alternative, especially how the alternatives address imposed constraints and satisfy stakeholder expectations. It is important that the process utilized by the deliberators affords him or her with ample opportunity to interact with the deliberators in order to fully understand the issues. This is particularly true if the decision-maker has delegated deliberation and downselection to a proxy. The information and interaction should present a clear, unbiased picture of the analysis results, findings, and recommendations. The more straightforward and clear the presentation, the easier it becomes to understand the differences among the alternatives.

Some of the same communication tools used in the TBfD can be used here as well, applied to the contending alternatives forwarded for the decision-maker's consideration. The imposed constraints risk matrix (Figure 4-21) summarizes what is among the most critical risk information. Additionally, information produced during deliberation should be summarized and forwarded to the decision-maker. This includes:

- **Risk tolerances and risk-normalized performance targets** – The deliberators establish risk tolerances on the performance measures, for the purpose of generating performance targets that can serve as the primary basis for comparison of alternatives. These tolerances and the resulting performance targets are key pieces of information for the decision-maker. They strongly influence requirements development and the corresponding program/project risk that is to be accepted going forward. A notional performance target chart is shown in Figure 4-29.
- **Pros and cons of each contending alternative** – An itemized table of the pros and cons of each alternative is also recommended for the contending alternatives. This format has a long history of use, and is capable of expressing qualitative issues. It enables conflicting opinions to be documented and communicated to the decision-maker, so that he or she is aware of contentious issues and/or competing objectives among stakeholders.
- **Risk lists** – Each alternative will have different contributors to its performance target risks. Correspondingly, each contending alternative will have a risk list written for it that identifies the major scenarios that contribute to risk. Each scenario has the potential to impact multiple performance measures over multiple mission execution domains.

Figure 4-30 presents a notional example of a RIDM risk list. Each row of Figure 4-30 represents a “risk,” as the term is used in the CRM process. Each risk is articulated in a risk statement, which identifies an existing *condition* (e.g., “A” for Risk #1) that indicates a possibility of some future *consequence* (“B” for Risk #1) that contributes to one or more performance targets not being met. The magnitude of the contribution is indicated in stoplight format (red/yellow/green) on a performance target basis, as well as on a holistic basis. The basis for determining the magnitude depends on the form of the risk assessment and the criteria established in the risk management plan (RMP), if one exists. For example, analyses that use detailed logic modeling might express risk contributions in terms of importance measures such as the Birnbaum, Fussell-Vesely or Risk

Reduction Worth (RRW) importance measures [15]. Alternatively, an approach that relies on a so-called normalized objective-based scenario consequence, or NOSC may be used. All of these approaches (Birnbaum, Fussell-Vesely, RRW, and NOSC) rely on the use of quantitative risk analysis methods to assess the effects of individual risk scenarios on aggregate performance measures. Less detailed analyses might use more qualitative criteria. Whatever method is used, consistency between the RIDM and CRM processes in this respect aids in the initialization of CRM for the selected alternative.

Alternative	Risk Tolerance	Risk-Normalized Performance Targets					Performance Legend
		PM ₁	PM ₂	PM ₃	...	PM _n	
		10%	1%	30%		5%	
1		100	0.007	0.01		0.4	Lower
2		500	0.0002	0.15		5	Medium
3		350	0.004	0.08		20	Higher
...							
m		400	0.0001	0.25		0.1	

Figure 4-29. Notional Performance Target Chart

Alternative X – RIDM Risk Analysis – Risk List							
Risk #	Risk Statement	Performance Commitments					Total
		PM ₁	PM ₂	PM ₃	...	PM _n	
1	Given A there is a possibility of B						
2	Given C there is a possibility of D						
3	Given E there is a possibility of F						
...							
m	Given Y there is a possibility of Z						

Figure 4-30. Notional Risk List for Alternative X

Regardless of how well the risk information is summarized or condensed into charts or matrices, the decision-maker should also always be presented with the raw risk results, namely the performance measure pdfs, upon request. Only by having these fundamental analysis results can the decision-maker bring his or her full judgment to bear on the selection of an alternative. Band-aid charts, as shown in Figure 4-22, are appropriate communication tools for communicating this information to the decision-maker in condensed format.

4.10 Details of Step 6 (Part 3), Select an Alternative and Accept the Associated Risk

Once the decision-maker has been presented with enough information for risk-informed decision making, he or she is ready to select a decision alternative for implementation. As discussed in Section 4.4.3 and portrayed in Figure 4-8, the decision itself consists of two main ingredients: the selection of the decision alternative and finalization of the performance targets.

4.10.1 Select a Decision Alternative

The RIDM process is concerned with assuring that decisions are risk-informed, and does not specify a particular process for selecting the decision alternative itself. Decision-makers are empowered to use their own methods for decision making. These may be qualitative or quantitative; they may be structured or unstructured; and they may involve solitary reflection or the use of advisory panels. Regardless of the method used for making the decision, the decision-maker formulates and documents the decision rationale in light of the risk analysis.

4.10.1.1 *Alternative Selection Process Is Iterative*

Just as risk analysis and deliberation iterate until the deliberators are satisfied that their issues and concerns have been satisfactorily addressed, alternative selection also iterates until the decision-maker is satisfied that the information at his or her disposal is sufficient for making a risk-informed decision. This is especially true in situations where the decision-maker has delegated much of the activity to others, and is exposed to the issues mainly through summary briefings of analyses and deliberations conducted beforehand. Iteration might consist of additional focused analyses, additional subject matter expert input, consideration of alternate risk tolerances (and associated performance targets) for some performance measures, etc.

4.10.2 Finalize the Performance Targets and Assist the Decision Authority's Deliberation on Requirements

In the requirements-based environment of the NASA program/project life cycle, decisions are essentially defined by the requirements they produce. Performance targets capture the performance characteristics that the decision-maker expects from the implemented alternative, and also establish the initial risk that the decision-maker is accepting and calling on the CRM process to manage. When RIDM is applied at the onset of a project or activity definition process, where the project or activity is intended to eventually design and assemble a system or process to achieve the intended objectives, this information is transferred into the next stage, which is typically a negotiation between “Acquirer” and “Provider” within which **performance requirements** are formally defined. In less formal contexts, the latter may be substituted by less rigidly defined **performance goals**, but in either case the identification of the RIDM produced combinations of

performance targets and corresponding **risk tolerances** represent the input that permits a risk-informed definition of requirements and/or goals.

As discussed in Section 4.9.2, performance targets are produced by the deliberators as a result of establishing risk tolerances on the performance measures. This facilitates deliberation of alternatives in terms of point value estimates of performance that reflect the deliberators' risk attitudes. The decision-maker may choose to keep the risk tolerances and performance targets established by the deliberators, or they may choose to modify them in accordance with his/her own risk tolerances. In situations where the decision-maker's risk tolerances differ significantly from those established by the deliberators, the decision-maker may ask for additional deliberation in light of the modified targets. In turn, the deliberators may ask the risk analysts for a revised risk list that reflects the new situation.

In addition to performance requirements, it is typical for the decision maker to levy other requirements that are motivated by other considerations such as adherence to best practices and lessons learned from past experience, or possibly legal, reputational, or political concerns. These requirement levies tend to come under the category of **process requirements**, rather than performance requirements, since they tend to be mandated apart from performance considerations. As mentioned in Section 4.4.3, the RIDM team should expect to be asked to provide an analysis of the decision maker's levied process requirements to determine whether they act in favor of or counter to the organization's ability to achieve its performance requirements. If the net effect of one or more levied process requirements on one or more performance requirements is negative rather than positive, the decision maker will need to justify the additional performance risk that is incurred through the risk acceptance process.

4.10.3 Accept the Risks of Selected Alternative and Document Decision Rationale

The final step in the RIDM process is for the decision-maker to document the rationale for the selected alternative in the Risk Informed Selection Report (RISR). A key part of the rationale to be documented is the identification of the risks associated with the selected alternative, and their "acceptance" as a result of the deliberation and selection process. "Acceptance of the risks" means in this context that the existence of such risks is explicitly acknowledged by the decision maker, who documents the corresponding information and transmits it to the organization(s) to which the activity or project that implements the actual realization of the selected alternative is assigned. It becomes then the responsibility of such organization(s) to decide as part of their CRM processes, applied while the activity or project born out of the selected alternative moves through its implementation and realization stages, how the risks identified by the RIDM process are addressed in time. It is therefore within such activity execution-stage processes that risks identified by RIDM may be actually "accepted as is" or acted upon with appropriate risk control measures.

Additional information on formulating and documenting the decision rationale can be found in Appendix G, Content Guide for the Risk-Informed Selection Report.

4.11 Graded and Special Focus RIDM Applications

This section provides guidance for tailoring the application of RIDM, according to criteria that depend on the objectives of its application and on the level of depth and fidelity that the programmatic and technical contexts of the application may require.

The first main subject covered in the following is the differences in the application of RIDM processes, when RIDM is used for “Activity Rebaselining” or for risk control analysis of alternatives in “Activity Execution,” with respect to the full reference definitions and illustrations provided in Sections 4.5 through 4.10 for Activity-Planning RIDM. In the following this is referred to as “RIDM specialization by type.” The RIDM specialization by type typically involves either a reduced scope of certain RIDM steps or sub-steps, or the elimination of some which may not be applicable in a given context. Section 4.11.1 covers these subjects.

A second main subject addressed in the section is the use of a graded approach to RIDM application, as it may be adopted in correspondence of a classification of the activity for which RIDM is applied, according to categories similar to those defined for robotic missions and instruments by NPR 8705.4. This is referred to as “RIDM graded approach by activity class.” The implementation of a graded approach according to activity class is addressed in Section 4.11.2 for Activity-Planning RIDM and Activity-Rebaseline RIDM, and in Section 5.3.1 for Activity-Execution RIDM (which as mentioned earlier in Section 4.7.2 is conducted in parallel with CRM and therefore subject to the same graded approach principles as CRM).

4.11.1 RIDM Specialization by Type

Table 4-IV provides a synthetic illustration of the differences in character that typically exist in the execution of Activity-Rebaseline RIDM and Activity-Execution RIDM, with respect to the full scope reference steps and sub-steps of an Activity-Planning RIDM. Key table explanations that are specific to either the Activity-Rebaseline or Activity-Execution RIDM types are provided below in Sections 4.11.1.1 and 4.11.1.2.

4.11.1.1 RIDM Specialization for Activity-Rebaseline

As introduced in Chapter 2, Activity-Rebaseline RIDM is conditionally called in the course of an activity execution by the Activity Decision Authority (e.g., in the case of a flight project, by the Project Manager), with the objective of modifying activity or system performance targets and, if necessary, performance and/or process requirements in order to satisfy the associated risk tolerance levels. The rebaselining is invoked if the CRM process applied in activity execution determines that the established activity requirements cannot be satisfied at the risk tolerance levels initially deemed acceptable, and if, to correct this situation, the Activity Decision Authority deems a modification of performance targets and requirements to be necessary. Activity contingencies requiring the activation of an Activity-Rebaseline RIDM are therefore “exception conditions” that generally are not expected to occur beyond the early stages of an activity execution.

Consistently with the rationale that triggers the possible initiation of an Activity-Rebaseline RIDM, the corresponding definition of steps illustrated in Table 4-IV assumes that this type of RIDM is carried out to re-evaluate and redefine risk tolerance levels and associated performance targets – as a key input to the redefinition and establishment, by the Activity Decision Authority and other responsible stakeholders, of new activity requirements that can be met at acceptable levels of risk. It must be noted that the Activity-Rebaseline RIDM definition adopted in this discussion assumes that the rebaselining objective of performance target, risk tolerance, and possibly requirement redefinition is applied to the specific activity solution alternative that was selected at Activity-Planning. In this definition, Activity-Rebaseline RIDM does not involve a

broader scope where completely new activity solution alternatives are introduced, or alternatives that were not selected at Activity-Planning are re-introduced and re-evaluated.

The above assumptions define a focused scope for Activity-Rebaseline RIDM. In particular, if a rebaselining RIDM process were not to remain limited to the redefinition of risk levels, performance targets, and requirements within a specific, previously selected alternative, then the associated context would be one where a full “activity replan” would actually be pursued, and initiated after completion of the RIDM processes. In such a case the RIDM execution would for the most part follow the blueprint of an Activity-Planning RIDM, with accordingly expanded scope and full set of attending steps.

In the context defined by the stated assumptions and conditions, Activity-Rebaseline RIDM is not scoped for the selection of an activity solution and plan alternative, but is instead limited to the determination of an appropriate balance between risk levels and performance targets for the alternative whose selection was informed by Activity-Planning RIDM. The risk vs. performance-target trade analysis is carried out in those performance dimensions for which the risk tolerance criteria and the performance and process requirements combinations established at activity-planning are not satisfied by the pre-rebaseline activity definition and plans. The step scope and content specializations suggested in Table 4-IV reflect this, by indicating that a typical Activity-Rebaseline RIDM process is mostly constituted by the execution of steps that correspond to Steps 3 through 6, i.e., Parts 2 and 3 of an Activity-Planning RIDM process. Execution of Part 1 is for the most part not needed, as any necessary information and data is already available per execution of the original Activity-Planning RIDM. Parts 2 and 3 are themselves executed with a significant reduction in scope, as the RIDM AoA does not concern activity alternatives, but only risk vs. performance targets and requirement trade studies within an existing selected alternative.

4.11.1.2 RIDM Specialization for Activity-Execution

Chapter 2 introduced Activity-Execution RIDM as an AoA-based process conducted as part of the CRM Plan step, for the specific purpose of selecting an “optimal” risk response among those that have been identified as possibly applicable in order to maintain activity performance risk below the established risk tolerance levels. As a necessary clarification, it is noted that this does not imply that an Activity-Execution RIDM must be invoked and carried out every time an Activity Decision Authority needs to make a choice among alternative risk responses. The decision of whether an Activity-Execution RIDM is needed to assure an appropriate and resource-effective selection and implementation of risk response rests with the activity decision authority, based also on the recommendation of the analysts and experts that are technically cognizant on the nature of the risk(s) to be addressed and on the range of means by which risk control can be accomplished.

The above definition of purpose characterizes Activity-Execution RIDM as being a RIDM execution with a specialized focus and scope. In Activity-Execution RIDM the “alternatives” being evaluated for deliberation and selection are not, as in Activity-Planning RIDM, general definitions of activity solutions and plans to be applied towards the achievement of stated activity objectives; instead, the alternatives being considered are the potentially applicable activity or system design solutions that may be applied to control activity and/or system performance risk, and keep such risk below the established tolerance levels. As in Activity-Rebaseline RIDM, the activity solution

and plans remain unchanged with respect to the output of the originally executed Activity-Planning RIDM. Moreover, performance requirements and risk tolerance definitions remain also as originally established and set at the activity pre-execution stages.

The risk response selection purpose of Activity-Execution RIDM is reflected across the board in all of its steps: the three RIDM reference Parts and six Steps thereof remain applicable as conceptual definitions of process execution; however, as Table 4-IV indicates and defines, their scope and analytical content changes significantly with respect to the corresponding elements of a “reference” Activity-Planning RIDM.

4.11.2 RIDM Graded Approach by Activity Class

The rationale for a “graded approach” stems from the recognition that breadth and depth of application of a process and its analytical elements is conditioned by how critical its outcomes are for the success of the activity or project within which it is applied, by how critical the activity or project itself is to the executing organization, and by the amount of resources reasonably available for the process execution. These factors are generally not independent, and their perceived intersection has been used to produce guidance for the tailoring and the graded implementation of various types of assurance processes and standards.

In proposing non-binding guidance for a graded approach to RIDM implementation, this handbook utilizes a classification scheme for the activities and projects to which RIDM may be applied. The adopted scheme is a generalization of the risk tolerance classes defined in NPR 8705.4, so that the corresponding definitions can be applied to any type of Agency organized activity, and not just to robotic missions and instruments.

The generalized, but essentially “NPR 8705.4 analogous,” Activity Classes are – Activity Class A+, Activity Class B, Activity Class C, and Activity Class D.¹⁶ The class assignment of any specific activity or project is then used to define RIDM step-tailoring and analytical-grading guidance. This guidance is organized and presented in detailed tabular format in Table 4-V.

The Table 4-V content is mostly self-explanatory, in that it reflects the general principle of adjusting the rigor and depth of RIDM application according to the criticality of the underlying activity and its availability of resource. These factors are generally strongly correlated and, as for the risk tolerance classes defined in NPR 8705.4, they decrease from Activity Class A+ to D.

As mentioned earlier in the beginning of Section 4.11, the graded approach criteria illustrated by Table 4-V can be applied in superposition to any limitation of RIDM process scope and steps dictated by the RIDM type that is to be executed. For example, in an Activity-Rebaseline RIDM, the execution of RIDM Part 1 steps is already very limited by the nature of the RIDM execution, thus the graded-approach criteria will generally be more focused on the tailoring of Part 2 and 3 steps according to the class assignment of the activity being rebaselined.

¹⁶ The “plus” (+) in Activity Class A+ indicates the inclusion of crewed missions.

Table 4-IV. Specialization of RIDM Steps by Activity Type

RIDM COMPONENT (NPR 8000.4C)			ACTIVITY PLANNING	ACTIVITY REBASELINE	ACTIVITY EXECUTION
PART	STEP	SUB-STEP			
Part 1 Identification of Alternatives <i>(Section 4.4.1)</i>	Step 1 Identify Objectives and Performance Measures <i>(Section 4.5)</i>	Sub-step 1A Identify Objectives	➤ In scope.	➤ Overall activity objectives are as identified at Activity-Planning. Rebaseline focus is on objectives impacted by risks that have triggered rebaselining.	➤ Objectives are pre-identified by the RIDM-invoking CRM process, as objectives of risk control for risks affecting specific performance measures.
		Sub-step 1B Identify Performance Measures	➤ In scope.	➤ Focus is on performance measures impacted by risks that have triggered rebaselining.	➤ Like objectives, performance measures of concern are pre-identified by the RIDM-invoking CRM process.
	Step 2 Identify Decision Alternatives <i>(Section 4.6)</i>	Sub-step 2A Compile an Initial Set of Alternatives	➤ In scope.	➤ Not in scope. Identification of new alternative design or solution concepts is not part of Activity-Rebaseline RIDM.	➤ Executed to identify risk response alternatives that appear to be potentially applicable.
		Sub-step 2B Identify Viable Decision Alternatives by Use of Trade Trees	➤ In scope.	➤ Not in scope, for same reason as Sub-step 2A.	➤ Executed to identify risk response alternatives that appear to be most viable.

Table 4-IV. Specialization of RIDM Steps by Activity Type (continued)

RIDM COMPONENT (NPR 8000.4C)			ACTIVITY PLANNING	ACTIVITY REBASELINE	ACTIVITY EXECUTION
PART	STEP	SUB-STEP			
Part 2 Analysis of Alternatives (Section 4.4.2)	Step 3 Conduct Integrated Risk Analysis of Each Alternative (Section 4.7)	Sub-step 3A Set the Analytical Framework	➤ In scope.	➤ Set risk-analytical framework that covers the performance targets deemed to necessitate rebaselining.	➤ Set risk-analytical framework that covers the risks for which risk responses are to be selected.
		Sub-step 3B Choose the Analysis Methodologies	➤ In scope.	➤ Choose methodologies suitable to modeling risk in performance measures for which performance targets rebaselining is sought.	➤ Choose methodologies suitable to modeling risk in performance measures affected by risk and controls to be selected.
		Sub-step 3C Conduct the Risk Analysis	➤ In scope.	➤ Focused on risk in performance measures for which performance targets rebaselining is sought.	➤ Focused on risk in performance measures for which risk responses are to be selected in AoA.
	Step 4 Develop the Technical Basis for Deliberation (Section 4.8)		➤ In scope.	➤ Focused on trade space for performance measures for which performance targets rebaselining is sought.	➤ Focused on trade space for risk responses that are to be selected in AoA.

Table 4-IV. Specialization of RIDM Steps by Activity Type (continued)

RIDM COMPONENT (NPR 8000.4C)			ACTIVITY PLANNING	ACTIVITY REBASELINE	ACTIVITY EXECUTION
PART	STEP	SUB-STEP			
Part 3 Risk Informed Alternative Selection (Section 4.4.3)	Step 5 Deliberate (Section 4.9)	Sub-step 5A: Convene a Deliberation Forum	➤ In scope.	➤ In scope, for convening forum of stakeholders and experts on performance targets being rebaselined.	➤ In scope, for convening forum of experts on performance risks being evaluated in controls AoA.
		Sub-step 5B: Define Risk Posture via Performance Markers and Risk Tolerances	➤ In scope.	➤ In scope, for risk tolerances and performance targets being rebaselined.	➤ Not in scope, RIDM uses performance markers and risk posture definitions set at Activity-Planning.
		Sub-step 5C: Identify Contending Alternatives	➤ In scope.	➤ Not in scope, for same reason as for Sub-step 2A.	➤ In scope, for identifying contending risk responses.
		Sub-step 5D: Communicate the Contending Alternatives to the Decision Maker	➤ In scope.	➤ Not in scope, for same reason as for Sub-step 2A.	➤ In scope, for communicating contending risk responses.
	Step 6 Select an Alternative and Accept the Associated Risk (Section 4.10)	Sub-step 6A: Select a Decision Alternative	➤ In scope.	➤ In scope, for acceptance of rebaselined risk.	➤ In scope, for selection of risk response.
		Sub-step 6B: Finalize the Performance Targets and Assist the Decision Authority's Deliberation on Requirements	➤ In scope.	➤ In scope, for finalization of rebaselined performance targets and requirements.	➤ Not in scope. Performance targets and requirements are same as at Activity-Planning.
		Sub-step 6C: Accept the Risks of Selected Alternative and Document Decision Rationale	➤ In scope.	➤ In scope, for acceptance of rebaselined risk tolerances and profiles.	➤ In scope, for acceptance of risk profiles, as projected after implementation of the selected risk response.

Table 4-V. RIDM Graded Approach by Activity Class

RIDM COMPONENT			CLASS A+	CLASS B	CLASS C	CLASS D
PART	STEP	SUB-STEP				
Part 1 Identifi- cation of Altern- atives	Step 1 Identify Objectives and Perfor- mance Measures	<u>Identify Objectives</u>	Consider and give appropriate weight to mission-objective expectations by all stakeholders inside and outside NASA. Decompose objectives into as many sub-objectives as necessary to describe the mission complexity.	Same as for Class A+, but limit the decomposition of mission-objectives into more detailed sub-objectives only to cases where important aspects of the mission would not be adequately defined.	Give priority consideration to mission-objective expectations by mission direct users, unless doing so leads to violation of externally-mandated constraints (e.g., in cost and schedule).	Consider only mission-objective expectations by mission direct users, unless doing so leads to violation of externally-mandated constraints (e.g., in cost and schedule). Focus on safety and major equipment damage caused by interactions with interfacing systems of higher value.
		<u>Identify Perfor- mance Measures</u>	Identify objectively-defined and quantitative performance measures to determine level of achievement of mission objectives and sub-objectives. Limit the use of ad-hoc defined scales or ratings as surrogate performance measures to situation where no other means of assessment are possible. Use as many dimensions of performance as it appears necessary to fully assess objectives and sub-objectives.	Identify objectively-defined and quantitative performance measures as for Class A+ to determine level of achievement of mission objectives. Iterate over the identification of objectives (preceding sub-step) in order to minimize the number of performance measures needed to assess the mission.	Identify objectively-defined performance measures that can be readily modeled and quantified (i.e., quantitative modeling and assessment does not require a high level of effort and organizational resources). Use constructed scales and discrete categories of performance for all objectives that cannot be readily modeled and measured in objectively quantifiable terms.	Identify performance measures that can be easily quantified. Use discretized and/or ad-hoc scales to assess performance in non-easily quantifiable performance dimensions.

	Step 2 Identify Decision Alternatives	<u>Compile an Initial Set of Alternatives</u>	Identify a comprehensive set of alternatives.	Same as for Class A+.	Identify alternatives that appear to satisfy mission or activity resource constraints.	Identify alternatives that appear to satisfy mission or activity resource constraints.
		<u>Identify Viable Decision Alternatives by Use of Trade Trees</u>	Apply trade tree process to identify alternatives that will be analyzed.	Same as for Class A+.	Set limit in number of alternatives that can be evaluated within resources. Apply judgment to eliminate alternatives that appear most costly / resource-intensive.	Same as for Class C, but possibly with even lower number of alternatives as "target."
Part 2 Analysis of Alternatives	Step 3 Conduct Integrated Risk Analysis of Each Alternative	<u>Set the Analytical Framework</u>	Set framework that covers Class A+ mission or activity performance dimensions.	Set framework that covers Class B mission or activity performance dimensions.	Set framework that covers Class C mission or activity performance dimensions.	Set framework that covers Class D mission or activity performance dimensions.
		<u>Choose the Analysis Methodologies</u>	Apply fully quantitative analytical methodologies in all analyses where it is possible.	Same as for Class A+, but apply simplified quantification where this is believed by experts to result in relatively minor loss of accuracy.	Alongside rigorous methodologies, consider the use of simplified, semi-qualitative, and expert-judgment-based analytical methodologies, where this is believed by experts to result in only minor or moderate loss of accuracy.	Preferentially use simplified, semi-qualitative, and expert-judgment based analytical methodologies, unless this is believed by experts to result in significant loss of accuracy.
		<u>Conduct the Risk Analysis</u>	Conduct analysis according to selection of methodology in preceding sub-step (Class A+ = maximum analytical effort).	Conduct analysis according to selection of methodology in preceding sub-step (Class B = significant analytical effort, but less than Class A+).	Conduct analysis according to selection of methodology in preceding sub-step (Class C = moderate analytical effort).	Conduct analysis according to selection of methodology in preceding sub-step (Class C = limited analytical effort).

	Step 4 Develop the Technical Basis for Delibera- tion	<i>n/a</i>	Document the analysis results consistently with the type and complexity of the analyses selected and applied in preceding sub-steps (Class A+ = maximum level of detail in analytical results documentation).	Document the analysis results consistently with the type and complexity of the analyses selected and applied in preceding sub-steps (Class B = significant level of detail in analytical results documentation, but less than Class A+).	Document the analysis results consistently with the type and complexity of the analyses selected and applied in preceding sub-steps (Class C = moderate level of detail in analytical results documentation).	Document the analysis results consistently with the type and complexity of the analyses selected and applied in preceding sub-steps (Class D = limited level of detail in analytical results documentation).
Part 3 Risk Informed Alterna- tive Selection	Step 5 Deliberate	<i><u>Convene a Delibera- tion Forum</u></i>	Identify SME's for deliberation forum as necessary to cover Class A+ performance dimensions (Class A+ = several SMEs generally needed).	Identify SME's for deliberation forum as necessary to cover Class B performance dimensions (Class B = several SMEs needed, but generally less than for Class A+)	Identify SME's for deliberation forum as necessary to cover Class C performance dimensions, giving preference to experts that can deliberate on multiple dimensions of performance (Class C = limited number of SMEs needed).	Identify SME's for deliberation forum as necessary to cover Class D performance dimensions, giving preference to experts that can deliberate on multiple dimensions of performance (Class D = very limited number of SMEs needed).
		<i><u>Define Risk Posture via Perfor- mance Markers and Risk Tolerances</u></i>	Class A+ = generally very low risk tolerance levels set in comparison of alternatives.	Class B = generally low risk tolerance levels set in comparison of alternatives.	Class C = generally medium risk tolerance levels set in comparison of alternatives.	Class D = generally medium or high risk tolerance levels set in comparison of alternatives.
		<i><u>Identify Contending Alternatives</u></i>	Apply iterative deliberation processes to downselect and rank alternatives.	Same as for Class A+ but keeping iterations to a minimum.	Downselect alternatives without iterations, unless explicitly requested by decision maker.	Same as for Class C.

		<u>Communi- cate the Contending Alternatives to the Decision Maker</u>	Present contending alternatives with full documentation of the processes applied to compare and rank performance (Class A+ = high level of detail in comparative results documentation allows review of deliberation process).	Present contending alternatives with documentation sufficient for insight in the processes applied to compare and rank performance (Class B = level of detail in comparative results documentation sufficient for insight into deliberation process).	Present contending alternatives with essential identification of the processes applied to compare and rank performance (Class C = comparative results documentation sufficient for identification of key elements of deliberation process).	Present contending alternatives without identification of the processes applied to compare and rank performance except where deemed important (Class D = comparative documentation limited to deliberation results).
	Step 6 Select an Alternative and Accept the Associated Risk	<u>Select a Decision Alternative</u>	Select alternative after as many iterations with deliberation forum as deemed necessary.	Select alternative. Use iteration with deliberation forum only if absolutely necessary.	Select alternative without iterations with deliberation forum.	Same as for Class C.
		<u>Finalize the Performance Targets and Assist the Decision Authority's Deliberation on Require- ments</u>	Class A+ = generally many performance targets to be finalized. Provide assistance on requirements as requested.	Class B = several performance targets to be finalized (generally less than Class A+). Provide assistance on requirements as requested.	Class C = moderate number of performance targets to be finalized (generally much less than Class A+ or B). Provide assistance on requirements as requested.	Class D = few performance targets to be finalized (generally less than Class C). Provide assistance on requirements as requested.
		<u>Accept the Risks of Selected Alternative and Document Decision Rationale</u>	Document in report-format accepted risks and decision rationale. Decision rationale documentation includes discussion of deliberation forum inputs and comparisons with discarded alternatives.	Document in report-format accepted risks and decision rationale. Decision rationale documentation includes synthetic description of deliberation forum inputs.	Document accepted risks and decision rationale in synthetic form (e.g., memo or briefing-format).	Same as for Class C.

4.12 References for Chapter 4

1. NASA Policy Directive, NPD 1000.0C, NASA Governance and Strategic Management Handbook. January 2020.
2. NASA Special Publication, NASA/SP-2014-612, NASA System Safety Handbook, Volume 2: System Safety Concepts, Guidelines, and Implementation Examples. November 2014.
3. Clemen, R., Making Hard Decisions, Duxbury Press. 1996.
4. Keeney, R., and Raiffa, H., Decisions with Multiple Objectives: Preferences and Value Tradeoffs, Cambridge University Press, 1993.
5. Hammond, J., Keeney, R., and Raiffa, H., Even Swaps: A Rational Method for Making Trade-offs, Harvard Business Review. March – April 1998.
6. U.S. Forest Service, Pacific Southwest Research Station, Comparative Risk Assessment Framework and Tools (CRAFT), Version 1.0, 2005.
7. NASA Aerospace Safety Advisory Panel. —Aerospace Safety Advisory Panel Annual Report for 2009, Washington, D.C., 2010.
8. Keeney, R., Value-Focused Thinking: A Path to Creative Decisionmaking, Harvard University Press, 1992.
9. NASA, Constellation Program Implementation of Human-Rating Requirements, Tracking Number 2009-01-02a. 2010.
10. Keeney, R., and McDaniels, T., A Framework to Guide Thinking and Analysis Regarding Climate Change Policies, Risk Analysis 21, No. 6, pp. 989-1000, Society for Risk Analysis, 2001.
11. NASA Special Publication, NASA/SP-2016-6105 Rev2, NASA Systems Engineering Handbook. Issued February 2017, Updated January 2020.
12. NASA Technical Memorandum, NASA-TM-2005-214062, Exploration Systems Architecture Study -- Final Report. 2005.
13. Maggio, G., Torres, A., Keisner, A., and Bowman, T., Streamlined Process for Assessment of Conceptual Exploration Architectures for Informed Design (SPACE-AID), AIAA Meeting Paper 2005-2587. Presented January 2005, Published Online December 2012.
14. NASA Cost Estimating Handbook, Version 4.0. Issued February 2015, Updated December 2020.
15. NASA Special Publication, NASA/SP-2011-3421 Probabilistic Risk Assessment Procedures Guide for NASA Managers and Practitioners, Second Edition. December 2011.
16. Morgan, M., and Henrion, M., Uncertainty: A Guide to Dealing with Uncertainty in Quantitative Risk and Policy Analysis, Cambridge Press, 1990.

17. Apostolakis, G., The Distinction between Aleatory and Epistemic Uncertainties is Important: An Example from the Inclusion of Aging Effects into Probabilistic Safety Assessment, Probabilistic Safety Analysis Conference (PSA'99), Washington, DC. 1999.
18. NASA Special Publication. NASA/SP-2009-569, Bayesian Inference for NASA Probabilistic Risk and Reliability Analysis. 2009.
19. U.S. Department of Defense (DoD) News Briefing Given by Donald Rumsfeld on February 12, 2002.
20. Mosleh, A., Siu, N., Smidts, C., and Lui, C., "Model Uncertainty: Its Characterization and Quantification," University of Maryland, 1993.
21. NASA Standard. NASA-STD-7009A w/ Change 1, Standard for Models and Simulations. December 2016.
22. Groen, F., and Vesely, B., Treatment of Uncertainties in the Comparison of Design Option Safety Attributes, PSAM 10, Seattle, WA. 2010.
23. National Research Council, Understanding Risk – Informing Decisions in a Democratic Society, The National Academies Press. 1996.

5 Continuous Risk Management

CRM is concerned with meeting all the organizational objectives that have been set out in the RMP. Because CRM generally takes place in the context of explicitly-stated performance measures, the risk that it manages is the potential for performance shortfalls which may be realized in the future.

The CRM process consists of the five cyclical steps, *Identify*, *Analyze*, *Plan*, *Track*, and *Control*, supported by comprehensive *Communicate* and *Document* functions. These steps and functions are illustrated in Figure 5-1. In practice, these steps operate in parallel, such that at any given time there may be individual risk scenarios being reported into the risk database; other individual risk scenarios being incorporated into the risk model; risk response plans being developed to reduce performance risk to tolerable levels; and implemented risk responses being tracked and controlled as needed to ensure their desired effects.



Figure 5-1. The CRM Process

CRM initiates its risk management processes by identifying specific scenarios or issues (which are not necessarily tied to well-defined scenarios) that are perceived as presenting a risk to the achievement of one or more organizational objectives. As discussed earlier in Section 3.2.2, these risk-significant scenarios and issues are referred to as *individual risk scenarios*, and collectively constitute the set of undesirable potential happenings that put the achievement of the activity objectives and requirements at risk of not being fulfilled. Each objective or requirement then has an associated risk that is the aggregation of the risk impacts from the set of individual risk scenarios that threaten the objective or requirement. The aggregate risk is quantified by applying risk analysis, using a scenario-based risk model that is informed from the risk model of the selected alternative developed during the Activity-Planning RIDM process. This risk model is augmented and refined as needed throughout implementation as individual risk scenarios are identified and incorporated into the model, and as detail is added commensurate with the maturity of the design or solution. In early stages of an activity, the risk model tends to be qualitative and based on expert judgment, whereas in later stages it tends to become mainly quantitative and based on data. In this

way, the risk model serves continuously as a tool for understanding the cumulative impacts that individual risk scenarios have on the stated objectives and requirements, and evolves to a degree that is consistent with the risk classification and level of maturity of the activity that is being analyzed.

CRM also manages opportunity, alongside managing risk, by identifying and managing desirable actions that, if implemented, have the potential for benefit (see Section 0). The potential benefit may arise either because the opportunity has the potential to reduce performance risk for a given activity (i.e., program, project, or other initiative), or because it has the potential to otherwise add benefit to the organization's portfolio of activities.

Once the organizational objectives have been specified, an alternative concept has been selected using the Activity-Planning RIDM process, and formal requirements have been developed for it as part of the technical requirements definition process [1], the risk associated with its implementation is managed using the CRM process during activity execution.

Consistently with the principle of application of a graded approach to risk management, throughout Chapter 5 the default manner of presentation for each step of the CRM process will be first to provide a table highlighting the main features of the topic being addressed individually for Activity Classes A+ through D, and then to follow up with more specific guidance on the characteristics and implementation of CRM that are appropriate to Activity Class A+. Since the guidance that pertains to Activity Classes B through D generally consists of culling the guidance for Activity Class A+, the combination of the introductory table and the subsequent guidance for Activity Class A+ is deemed sufficient to provide the coverage needed on the topic, including how to apply a graded approach for CRM which is aligned with the Activity classification that has been established.

5.1 Initialization of CRM

After a concept has been selected for implementation using the Activity-Planning RIDM process, and organizational objectives and formal requirements have been defined for it by the management of the acquiring organization, the CRM process is initiated to provide a framework for ongoing risk management in a manner that is appropriately standardized across the affected organizational units.

5.1.1 Development of the Risk Management Plan

As described in NPR 8000.4C [2], the development of a Risk Management Plan (RMP) occurs at each organizational level and is a responsibility of the management at that level. The RMP is part of the basic program/project or activity management planning and documentation, and should be cross-referenced and integrated with other parts of such documentation. It is therefore not a task pursued strictly within either RIDM or CRM. Each organizational entity producing an RMP is considered to be a *Provider* to another organizational entity considered to be the *Acquirer*. The *Acquirer's* organizational management reviews and accepts the *Provider's* RMP following a negotiating process that results, metaphorically speaking at least, in a "handshake." During the process of preparing and approving the RMP, those who have executed the RIDM process and those who will be executing the CRM process act as subject matter experts, providing the needed guidance.

An upfront phase of planning resulting in the RMP is necessary to assure the development of a robust risk management process and common understanding of the risks to be addressed. In addition to detailing how each CRM step will be carried out, the RMP should serve as the means to identify and define the key coordination and technical provisions that are to be implemented in the course of the CRM process.

Not all initiatives or activities require a detailed RMP, but some form of RMP should be prepared for all initiatives and activities during the planning stage. For large programs, projects, institutional initiatives, and enterprise initiatives, the key elements of the risk management framework that are developed and captured in the RMP include:

- Identification of stakeholders, such as risk review boards, to participate in deliberations regarding the response to risks
- Documentation of a complete set of organizational objectives, formulated first at the top organizational level responsible for the initiative or activity in question, and then allocated down to the lower levels.
- Establishment of performance markers (e.g., PMK-R and PMK-G values), associated risk tolerance levels (e.g., RTL-R and RTL-G values), and elevation protocols (the specific conditions under which a risk management decision must be elevated through management to the next higher level)
- Establishment of risk burn-down schedules for each objective and requirement, with respect to the risks identified during Activity-Planning RIDM (to be developed further in Section 5.1.5).
- For each objective and requirement, documentation, or indication by reference, of whether its associated risks (including the aggregate risk) are to be assessed quantitatively or qualitatively including a rationale for cases where it is only feasible to assess the risk qualitatively
- Establishment of risk communication protocols between management levels, including the frequency and content of reporting, as well as identification of entities that will receive risk tracking data from the unit's risk management activity
- Delineation of the processes for coordinating risk management activities and sharing risk information with other affected organizational units.

The RMP for smaller sized initiatives and activities should address only those elements cited in the above bullets that are relevant.

5.1.2 Inputs to CRM from the Activity-Planning RIDM Process

Many of the products of the Activity-Planning RIDM process carry over to CRM. These products include:

- Identification of objectives and associated performance measures – As part of its initialization, the RIDM process proffers a clear definition of the organization's objectives, consistent with the stakeholders' values and priorities. Management decomposes the top-

level objectives to lower-level objectives and allocates them to the supporting organizational units. For each objective at each level, the RIDM process identifies one or more quantifiable performance measures that provide the basis for assessing the degree to which the objective is being satisfied.

- The risk analysis of the selected alternative – The risk analysis that was developed during RIDM for the selected alternative is maintained throughout the CRM process. It is the risk model that provides the core risk analysis capability for assessing risk to organizational objectives and identifying risk drivers.
- Performance marker set-points (e.g., constraints and/or targets), and associated risk tolerances – These values are developed to characterize the risk posture initially expressed by the stakeholders and later refined by organizational management (see Section 0).
- The risk list – The risk list generated during Activity-Planning RIDM identifies the major scenarios (both defined and implied) contributing to the selected alternative's risk to objectives that are known at the time.
- The Risk-Informed Selection Report (RISR) – The RISR documents the RIDM process, including the decision rationale, as discussed in Appendix G.

The risk analysis performed during Activity-Planning RIDM will likely have addressed only those performance measures considered discriminators among the alternatives, and so will likely have to be supplemented during CRM. Moreover, because the initial risk list is based on the RIDM risk analysis, it is likely to contain only the major, top-level, initially evident risks and may therefore be incomplete, especially with respect to the non-discriminator performance measures.

As soon as feasible, the CRM process will need to complete the Activity-Planning RIDM risk analysis for the non-discriminator performance measures and expand and update the initial risk list to include any new risks from the completed risk analysis. There will also tend to be a transition from parameter-based modeling to scenario-based modeling wherever feasible.

The risk models developed during RIDM are sometimes developed at different levels of the organizational hierarchy, but because Activity-Planning RIDM is conducted during concept development, the RIDM risk analyses can be also done before the organizational hierarchy is established. This is not true of CRM risk models, because CRM is conducted at every level of the hierarchy. Therefore, CRM risk models rely on the upward flow of risk information from the models at lower levels of the organizational hierarchy, as well as on the consistency of models and assumptions among all units whose results are ultimately aggregated into performance risk at some shared higher level. Consequently, model sharing and data reporting protocols must be established as needed to support the distributed nature of risk management within the hierarchical structure of NASA. Section 5.7.1 and much of Part 2 will address the communication protocols necessary to ensure consistency and availability of data among units that are bound by a common set of objectives.

5.1.3 Inputs to CRM from Local Organizational Management

5.1.3.1 *Flow-Down of Organizational Objectives, Mandated Requirements, Performance Measures, and Associated Risk Tolerances*

Other inputs to the CRM process additional to those obtained directly from Activity-Planning RIDM include the organization-specific objectives for the selected alternative approach and the risk tolerances associated with them. The flowed-down objectives and requirements and associated performance measures and risk tolerances are the result of a negotiated decomposition and allocation process that flows downward through the NASA hierarchy, as discussed for example in the NASA Systems Engineering Handbook [1] with respect to programs and projects. It is expected that the performance marker values (e.g., PMK-R and PMK-G values) and associated risk tolerance levels (e.g., RTL-R and RTL-G values) will be informed by Activity-Planning RIDM's performance measure pdfs obtained from a risk analysis of the selected alternative and by other factors deemed pertinent by the management authorities. However, the totality of the requirements, including the parameters upon which requirements are levied, are not expected to be identical to or limited to the performance measures used by Activity-Planning RIDM. The management authorities that specify requirements are not constrained to base their requirements on the RIDM analysis, only to consider the risk information it produces in its deliberations.

5.1.4 Establishing the Performance Measures To Be Considered, Performance Markers, and Associated Risk Tolerance levels

Graded Approach Guidance			
Activity Class A+	Activity Class B	Activity Class C	Activity Class D
One or more quantitative performance measures are identified for each organizational objective as needed to gauge the level of achievement for each. All objectives and requirements specified in the risk management plan are considered, ensuring comprehensive coverage across programmatic, engineering, institutional, and enterprise activity domains and across safety, technical, security, cost, and schedule execution domains.	Same as for Activity Class A+, except reduced coverage is acceptable based on experience and expert judgement to deemphasize areas known historically to be not significant for achieving safety, mission success, or programmatic objectives and requirements.	Same as for Activity Class A+, except coverage focuses on readily evaluated performance measures that historically are significant to safety, mission success, or programmatic objectives and requirements.	Same as for Activity Class A+, except coverage focuses on safety and major equipment damage caused by interactions with interfacing systems of higher value.

Performance measures and associated risk tolerances are first identified during the Activity-Planning RIDM process as a means for comparing alternative high-level concepts, usually prior to

any significant amount of system design or solution development. The RIDM performance measures are identified within the framework of judging the degree to which the managing organization's objectives are being satisfied, wherein the objectives are decomposed down to a level that is consistent with distinguishing between the concepts. Compared to Activity-Planning RIDM, CRM works with a set of objectives that may be considerably more evolved in concert with the higher level of design definition. CRM is by its nature distributed among both high and low levels of the organization, as well as high and low levels of partnering organizations, so the performance measures identified during CRM are specific to the organizational entity whether it is at a high or low level in the organizational hierarchy, and not just those at higher levels.

Accordingly, the initialization of CRM includes a reevaluation of the performance measures identified during Activity-Planning RIDM and the possible development of an expanded set that addresses the objectives that have been allocated to each organizational entity. That said, the process of identifying applicable performance measures follows the path laid out in Sections 4.2 and 4.5, and the process of developing applicable risk tolerances proceeds along the path put forth in Section 4.9.2.

5.1.4.1 Update the Risk Models to Address All Performance Measures

Graded Approach Guidance			
Activity Class A+	Activity Class B	Activity Class C	Activity Class D
Analysis models during CRM start from Activity-Planning RIDM analysis models but are expanded in both breadth and depth to include design and/or solution details developed during the life cycle of the program/project or activity as well as test results and risk/opportunity findings. Fully integrated probabilistic modeling of performance measures are developed for all mission execution domains (safety, technical, security, cost, schedule).	Same as for Activity Class A+, except methods may be tailored to using more simplified models when justified that concentrate on obtaining point-estimate probabilities with uncertainty bounds rather than full distributions.	Same as for Activity Class A+, except methods may be tailored to more simplified models that concentrate on obtaining point-estimate probabilities with uncertainty bounds rather than full distributions. Emphasis on analysis of individual risk scenarios with a more qualitative assessment of aggregate performance risk.	Use of qualitative models and simplified quantitative models that provide reasonably conservative point estimates is sufficient to show no harm to other assets.

As discussed in Section 5.1.2, the risk models developed during the RIDM process in the early concept development stage of a program/project or other activity will have addressed only the performance measures considered to be discriminators among the various alternatives concepts being considered. In CRM, these models are updated to include the non-discriminator performance measures, and then maintained on a continuous basis by incorporating any pertinent new information that is generated during the design, development, and implementation of the activity.

Generally, this will include design details as they become available, test results, and new risks and opportunities that reveal themselves in different life-cycle stages.

5.1.5 Risk Burn-Down Schedules

It is common for risk to be reduced over the life cycle of an activity as risk drivers are identified, risk controls are implemented, and uncertainties are reduced. The analyzed risk at the beginning of an activity might not adhere to the established risk posture, but might be considered “on track” to meeting it if risks are actively being identified and controlled by the CRM process.

As an example, the risk of not meeting the required launch date for a space flight may be relatively high at the beginning of the project if the project requires the development of new technology. The RMP will specify that this risk should diminish to within acceptable levels by the Critical Design Review (CDR). Similarly, the risk of not being able to provide an effective IT security system within a specified budget and schedule may be high at first because of uncertainties that could result in cost and schedule overruns as the initiative progresses. In general, the RMP should include a risk burn-down schedule that reflects the anticipated success of the planned RM activities. The burning down of risk in accordance with the schedule then becomes a commitment that the activity is assessed against at the relevant milestone reviews such as LCRs, and should be incorporated into the success criteria of the reviews. Correspondingly, the activity budget should include a cost and schedule reserve consistent with the anticipated RM activities, and the adequacy of this reserve to address any remaining known or U/U risks should be likewise evaluated at milestone reviews.

As the activity evolves over time, mitigations are implemented; and as risk concerns are retired and the state of knowledge about the performance measures improves, uncertainty should decrease, with an attendant lowering of risk and increase in confidence, to the point where the risk is within the established risk posture. This is illustrated notionally in Figure 5-2, which shows risk being reduced from “unacceptable” (red) at Milestone 1, to “marginal” (yellow) at Milestone 2, to “acceptable” (green) at Milestone 3. When this decrease is within the risk burn-down schedule, the activity can be considered to be “on track” to adhering to the risk posture. When risk reduction is not within the burn-down schedule, doubt is raised as to whether or not the activity is “on track.”




Figure 5-2. Decreasing Risk over Time for an Activity that is “On Track” to Being Within the Established Risk Posture

Because risk responses often shift risk from one domain to another, risk burn-down schedules should be developed holistically for all performance measures, and the activity should be evaluated holistically against the full set of risk burn-down schedules. In particular, safety and technical risks are often controlled at the expense of cost and schedule risk. When the establishment of cost and schedule reserves is risk-informed and the realized costs and schedule impacts of RM activities are consistent with the established reserves, then cost and schedule risk is expected to decrease over time. However, if the cost and schedule impacts of RM activities exceed the anticipated impacts, then cost and schedule risk can increase dramatically. This mismatch between reserves and impacts is essentially a U/U issue – either risks arose that weren't known or anticipated during budget planning, or the impacts of managing the known or anticipated risks were underappreciated.

Figure 5-3 notionally shows a risk burn-down schedule for the risk to a hypothetical performance measure representing an organizational objective. The profile is based on the initial assessed performance risk and the time at which that risk is expected to meet the established risk posture. The specifics of the burn-down schedule depend on the details of the activity flow and the risk management plan.




Figure 5-3. Risk Burn-Down Schedule for a Hypothetical Organizational Objective

5.1.6 [The Risk Database](#)

NPR 8000.4 [2] requires risk dispositions and risk acceptance decisions and their rationales to be documented in the relevant organizational unit's risk database. Risk databases of individual risk scenarios and opportunities, leading indicators, performance parameters, performance measures, risk tolerances, etc., are powerful tools, not only for configuration management and institutional memory, but as integral parts of each organizational unit's ongoing risk management activities. In particular, the integration of separate organizational units' risk databases into an integrated risk database helps organizations to identify risks and opportunities they may have otherwise overlooked but were identified in other units; identify risk response alternatives and examine the rationales other units have used for risk response selection; allocate risk postures consistent with allocations for similar activities; and identify cross-cutting risks that are best managed cooperatively or by a higher-level organizational unit. It also helps to facilitate the process of

rolling up individual risk scenario and opportunity scenarios from lower levels to an aggregate view of the overall likelihood of success of meeting the activity's objectives and requirements.

The integration of standardized taxonomies into an integrated, inter-organizational risk database facilitates the identification and management of cross-cutting risks. For example, multiple organizational units might separately register individual risk scenarios related to a particular supply chain issue, or to the use of a particular piece of hardware, or to some phenomenon such as corrosion. The ability to search and filter risks on an Agency-wide basis using common taxonomic characterization of the risk enables a "birds-eye view" of individual risk scenarios across the Agency and the ability to craft systemic responses to systemic risks, either cooperatively by the directly affected units or from above by the appropriate superordinate unit.

The use of taxonomies is discussed further in Appendix I.

5.2 CRM Step 1: Identify

The sub-steps of the Step 1 *Identify* step are as follows:

- Identify individual risk scenarios, opportunities, and leading indicators
- Develop risk and opportunity statements
- Validate the risk and opportunity statements
- Develop accompanying risk and opportunity narratives

5.2.1 Identify Individual Risk Scenarios, Opportunities, and Leading Indicators

Graded Approach Guidance			
Activity Class A+	Activity Class B	Activity Class C	Activity Class D
Involves a comprehensive identification of individual risk scenarios, opportunities, and leading indicators. Includes coverage across programmatic, engineering, institutional, and enterprise activity domains and across safety, technical, physical security, cybersecurity, cost, and schedule domains.	Same as for Activity Class A+, except reduced coverage is acceptable based on experience and expert judgement to deemphasize areas known historically to be not significant for achieving safety, mission success, programmatic requirements, or other requirements important to the activity being pursued.	Same as for Activity Class A+, except coverage focuses on historically significant risks and leading indicators that affect safety, mission success, programmatic requirements, or other requirements important to the activity being pursued, and on newly discovered opportunities that meet a minimum level of potential benefit.	Same as for Activity Class A+, except coverage focuses on safety and major equipment damage caused by interactions with interfacing systems of higher value.

Individual risk scenarios, opportunities, and leading indicators can be identified at any point in the life cycle of a program/project or activity, and indeed it is a fundamental principle of CRM that the process of identification is ongoing. The primary sources of risk, opportunity, and leading indicator identification are expected to be:

- **The Initial Risk Analysis** – The initial risk analysis has its origin in the risk analysis of the selected alternative developed during the Activity-Planning RIDM process. During RIDM, significant uncertainties with the potential to adversely affect performance are identified, and a risk list is generated identifying the major scenarios contributing to the risk of that alternative. When CRM is initialized, the RIDM risk list is updated to reflect risk with respect to organizational objectives.
- **Taxonomy-Facilitated Brainstorming** – Brainstorming is a common method for identifying potential risks and opportunities. Various types of brainstorming techniques such as Checklist, What-if Analysis, Failure Modes and Effects Analysis (FMEA), and Hazard and Operability Studies (HAZOP) have been used for decades in the process industries to identify potential process upsets. Although numerous techniques are potentially applicable to the *Identify* step depending on the nature of the activity, the risk, opportunity, and leading indicator taxonomies can be used as powerful tools for brainstorming. Specifically, each unique combination of taxa from the taxonomies can be used in what-if fashion to structure the brainstorming sessions and stimulate thought.
- **Conditions Arising during Implementation** – As implementation proceeds, conditions can emerge that signal the presence of risk or new opportunity. There are a variety of potential sources of risk- or opportunity-indicating conditions including:
 - Data from systems engineering – As discussed in the NASA Systems Engineering Handbook [1], key success criteria and performance parameters are monitored during implementation by comparing the current actual achievement of the parameters with the values that were anticipated for the current time and projected for future dates. These data are used to confirm progress and identify deficiencies that might jeopardize meeting an objective or requirement, including cost and schedule constraints. When a parameter value falls outside the expected range around the anticipated value, it signals a need for evaluation and corrective action.
 - Tracking data for implemented risk responses – As discussed later in Section 5.5, the risk responses of *watch*, *research*, and *mitigate* entail the specification of data that will be tracked to monitor implementation of the response and assess its effectiveness in addressing the risk to organizational objectives. Similar to the case for data from systems engineering, when tracking data fall outside expectations, it signals the need for evaluation and corrective action. If such action can be accomplished within the framework of the risk response, it is not necessary to identify a new individual risk scenario in addition to the individual risk scenario(s) already underlying the current risk response. However, if the data are such that a new risk response is warranted, then the development of a new individual risk scenario is advisable.

- Inter-organizational communications – In each step of the CRM process, risk and opportunity information is communicated among the various risk management organizations within the NASA hierarchy, according to the degree to which they are all working towards the accomplishment of high-level objectives and requirements, and/or are vulnerable to similar conditions and departure events (i.e., so-called “cross-cutting risks”). GIDEP alerts are a good example. This information may indicate the presence of risk or opportunity within a given organizational unit that until then had been unidentified.
- Risks or opportunities elevated from lower levels in the organizational hierarchy – If a unit in the NASA hierarchy is unable to adequately manage its performance risk, or avail itself of a new opportunity, it may elevate the management of its risk or opportunity to the unit at the next higher level of the NASA hierarchy. When this is the case, the situation is identified as an individual risk scenario or opportunity at the higher level if it has not already been included as such. The CRM process can then be applied in order to assess the lower level unit’s performance risk in terms of its impact on the higher level unit’s performance risk.
- External sources of risk – The design and development of a solution is seldom accomplished in isolation from external considerations, such as the price and availability of parts and/or raw materials; the price, availability, and skill sets of human capital; or the availability of test facilities and other support functions. Project or activity planning necessarily involves assumptions about such considerations, which, as time goes on, may prove not to be the case. The *Identify* step captures such situations as individual risk scenarios when external conditions change in ways that adversely affect performance risk.
- **Rebaselining of Requirements** – It is possible that the unit at the next higher level of the NASA hierarchy will need to revise their derived requirements using the Activity-Rebaseline RIDM process as part of a risk mitigation effort at its level. When this is the case, objectives that flow down from the higher level to the current level are rebaselined in a negotiated fashion, leading to a modified set of objectives against which performance risk is assessed. The rebaselining may involve an adjustment process, wherein certain requirements are modified to make them more applicable and practicable, or alternatively an outright waiving of requirements that are unnecessary or counterproductive. Once approval is obtained, the rebaselining of the objectives will typically require modification of the local organization’s Risk Management Plan (and possibly the overall program/project or activity plan), which in turn will require an iteration on the identification of individual risk scenarios and opportunities.

5.2.2 Develop Risk and Opportunity Statements

Graded Approach Guidance			
Activity Class A+	Activity Class B	Activity Class C	Activity Class D
Separate statement formats are used for each of four statement types: (1) risks with defined scenarios, (2) risks with implied scenarios, (3) opportunities within the plan, and (4) opportunities outside the plan. Within these formats, each covered individual risk scenario or opportunity has a customized risk or opportunity statement.	Same as for Activity Class A+.	Same as for Activity Class A+.	Same as for Activity Class A+.

A *risk statement* can have either of two formats, depending on whether the scenario is well-defined or rather implied from a set of leading indicators. If the scenario is defined (i.e., there is a well-defined sequence of events leading to the end state), the risk statement has the following format:

“Given that **[Condition]**, there is a possibility of **[Departure]** affecting **[Asset]**, which can result in **[Consequence]**, adversely affecting the achievement of **[Affected Objective(s)]**. ”

In addition to the statement itself, there should be a notation about whether the risk is actionable (i.e., whether something can be done to prevent, or reduce the likelihood of the Departure and/or severity of the Consequence).

It is the job of the risk identifier, working as needed with risk management personnel, to develop verbiage for the condition, departure, consequence, and objective/requirement components of the risk statement.

- **Condition** – The Condition is a single phrase that describes a current key fact-based situation or environment that is causing concern, doubt, anxiety, or uneasiness. The fact-based aspect of the Condition helps to ground the individual risk scenario in reality, in order to prevent the risk database from becoming a repository for purely speculative concerns. The Condition represents evidence in support of the concern that can be independently evaluated by risk management personnel and which may be of value in determining an appropriate risk management response during the CRM *Plan* step.
- **Departure** – The Departure describes a possible change from the (agency, program, project, initiative, or activity) baseline plan. It is an undesired event that is made credible or more likely as a result of the Condition. Unlike the Condition, the Departure is a statement about what might occur at a future time. It is the uncertainty in the occurrence or non-occurrence of the Departure that is the initially identified source of risk.
- **Asset** – The Asset is an element of the organizational unit portfolio (OUP) (analogous to a WBS). It represents the primary resource that is affected by the individual risk.

- **Consequence** – The Consequence is a single phrase that describes the undesired state or condition of the Asset that could be produced by the realization of the individual risk scenario. The Consequence should not take into account any anticipated risk responses. In cases where a number of undesired states are possible, depending on existing mitigating or exacerbating factors, the stated Consequence should describe the state or condition that is most representative of the concern.
- **Affected Objective(s)** – The Affected Objectives(s) is/are the top-level objective(s) that is/are most threatened by the individual risk scenario. The identification of the Affected Objective(s) should not take into account any anticipated risk responses. The set of potentially affected objectives is the set of objectives for which PMs and PMKs are defined, and include mandated requirements and programmatic constraints. In cases where a number of top-level objectives are threatened, the stated Affected Objective(s) should identify those for which a shortfall is most representative of the concern.

RISK STATEMENT EXAMPLES FOR DEFINED SCENARIOS

- *Example 1 (Program/Project Context):* Given that **[Condition]** component Y of exploratory system X is highly susceptible to solar flares, there is a possibility of **[Departure]** a solar flare causing failure of **[Asset]** component Y, which can result in **[Consequence]** the quantity and quality of data received and transmitted by system X being compromised, adversely affecting the achievement of **[Affected Objective(s)]** significant gains in understanding current astrophysical mysteries
- *Example 2 (Institutional Context):* Given that **[Condition]** Center X's 5-year plan calls for doubling the number of its qualified specialists in field Z in order to support the Agency's new strategic plan, there is a possibility of **[Departure]** national economic conditions, national demographic changes, and competition from other employers changing in such a way that **[Asset]** Center X is not able to attract and maintain enough qualified graduates in field Z, which can result in **[Consequence]** shortages of personnel in field Z, adversely affecting the achievement of **[Affected Objective]** Center X's 5-year plan
- *Example 3 (Enterprise Context):* Given that **[Condition]** in a large percentage of the IT infrastructure, there is no known detection or audit mechanism available to determine if we are being attacked, or were attacked, although so far there is no indication from any of our personnel or Government sponsors that the operation of the IT system has been compromised or that sensitive information has been disclosed to unauthorized entities, there is a possibility of **[Departure]** an unauthorized entity such as a domestic or foreign hacker gaining access to **[Asset]** our system, **[Consequence]** degrading or terminating the operation of the system and gaining access to protected information, with the intrusion not being detected until the attack is over, affecting **[Affected Objective(s)]** the protection of sensitive technical data and NASA's reputation as a center of technical excellence.

If the scenario is implied but not defined (i.e., there is an established correlation between a leading indicator and failures to meet an objective without reference to a particular scenario), the risk statement has the following format:

*“Given that **[Leading Indicator]** has been determined to correlate significantly with unknown and/or underappreciated (U/U) risks in the present context and that **[Condition]** has caused this*

*leading indicator to exceed its “watch” value, there is a potential for **[Departure]** to cause the leading indicator to exceed its “respond” value, thereby affecting the achievement of **[Affected Objective]**. ”*

RISK STATEMENT EXAMPLE FOR AN IMPLIED SCENARIO

- *Example 4 (Program/Project Context): Given that **[Leading Indicator]** design complexity¹⁷ has been determined to correlate significantly with unknown and/or underappreciated (U/U) risks in the present context and that **[Condition]** new mission requirements combined with a limited budget has caused this leading indicator to exceed its “watch” value, there is a potential for **[Departure]** present underestimations or future requirement perturbations to cause the leading indicator to exceed its “respond” value, thereby affecting the achievement of **[Affected Objective]** launching System X before closure of the launch window*

It is of fundamental importance to the CRM process that risk statements be crafted without regard to potential mitigations or other risk responses that may suggest themselves to the risk identifier. The risk statement should not presume anything that is not in the current baseline project plan, other than the Condition, which has its basis in fact. In particular, the Consequence should presume that no risk response has been implemented that would shift the consequence from, say, a potential over-mass condition to a cost overrun and/or schedule slippage resulting from an anticipated redesign. It is recognized that the resulting risk statements can be considered artificial in that the issue might not be allowed to persist without a risk management response of any kind, but the point of the *Identify* step is specifically to capture the concern, not to presume the manner in which it will be addressed. This is not to say that the risk identifier should be silent on the topic of potential risk responses. On the contrary, such input is strongly encouraged, but should be included in the narrative description section of the individual risk scenario, as will be discussed in Section 5.2.4.

Like a risk statement, an opportunity statement can have either of two formats. The distinction is whether the opportunity concerns objectives that are already within the current plan, and simply increases the likelihood of being able to satisfy them, or whether it introduces a new objective that goes beyond the current plan but beneficially serves the organization’s overarching mission. If the opportunity lies within the current plan and serves one or more of its objectives and/or requirements, the opportunity statement has the following format:

*“Given that **[Condition]**, there is a potential for **[Action]**, which could result in **[Proximate Benefit]** and positively affect the achievement of **[Organizational Objective]**. ”*

¹⁷ Design complexity is often indicated in the literature as a leading indicator of risk, however is in itself a relative concept, driven in a given engineering system or project by such factors as the number of functional requirements, interfaces, moving parts, etc. Such factors may determine “complexity” in absolute terms, but also in relative relation to the resources available for the realization and implementation of a system design.

OPPORTUNITY STATEMENT EXAMPLE FOR AN OPPORTUNITY THAT ENHANCES THE ACHIEVEMENT OF EXISTING OBJECTIVES AND REQUIREMENTS

- Example 1 (Program/Project Context): Given that **[Condition]** (1) Company X has expressed its intention to build and field a space station capable of accommodating human occupants in far regions of space, including cislunar orbit: (2) it has also indicated its willingness to accept the risk of training its own astronauts and sending them on its own transport system to assemble the space station and subsequently maintain it, and (3) experience shows that commercial companies are typically able to conduct spacefaring missions at significant cost savings compared to the cost for NASA to do it using prime contractors, there is a potential for **[Action]** NASA commissioning Company X to build, assemble, maintain, and inhabit a space station in cislunar orbit, which could result in **[Proximate Benefit]** significant savings in cost and schedule, as well as benefits with regard to legal liability in the event of an accident, and positively affect the achievement of **[Organizational Objective]** building, assembling, maintaining, and inhabiting a space station in cislunar orbit

If instead the opportunity goes beyond the current plan and introduces new objectives that serves the organization's overarching mission, the opportunity statement has the following format:

"Given **[Condition]**, there is a potential for **[New Objective]** with corresponding **[Action]**, which could result in **[Proximate Benefit]** and positively affect the achievement of **[Mission Goal]**."

OPPORTUNITY STATEMENT EXAMPLE FOR AN OPPORTUNITY THAT ENHANCES ACHIEVEMENT OF THE ORGANIZATION'S OVERARCHING MISSION

- Example 2 (Enterprise Context): Given that **[Condition]** new technology developments in the area of controlled fusion have made it theoretically possible to increase the propulsive power and operating range of unmanned space systems by an order of magnitude, there is a potential for **[New Objective]** developing, demonstrating, and fielding a propulsion system that uses controlled fusion, which could result in **[Proximate Benefit]** opening up possibilities for deep exploration of our galaxy that didn't exist before, and positively affect the achievement of **[Mission Goal]** leading an innovative and sustainable program of exploration to increase understanding of the universe and our place in it

5.2.3 Validate the Risk and Opportunity Statements

The following seven questions can be used to guide the writing of an individual risk scenario or opportunity to ensure that it is valid. If the answer to any of the questions is "no" or "unknown," the risk or opportunity should not be considered valid and the author may wish to go back and modify it appropriately (possibly with the help of the appropriate risk management personnel) or abandon the effort.

1. Does the individual risk scenario or opportunity impact at least one agency/program/project/activity objective or mandated requirement, or contribute to a mission goal, and can that impact be objectively measured, described, and characterized?
2. Does the individual risk scenario or opportunity statement adequately communicate the possible sequence of events leading from the Condition or Leading Indicator, through the Departure or Action, to the Consequence or Proximate Benefit; and is the Consequence or Proximate Benefit expressed in terms of its effect on one or more Organizational Objectives or Mandated Requirements, or its contribution to a Mission Goal?
3. Is the individual risk scenario or opportunity based on relevant documentation or individual/group knowledge?
4. Does the individual risk scenario or opportunity involve a change from the program/project, initiative, or activity baseline plan for which an adequate contingency plan does not exist?

Note: If it involves a change that causes the existing contingency plan to be inadequate, the failure of that contingency plan should be addressed in the Departure portion of the risk or opportunity statement.

5. Is the Condition and/or Leading Indicator factually true and supported by objective evidence?
6. Is the Departure or Action credible (possible)?
7. Is the Consequence written without regard to potential mitigations?

After completion by the author and entry into the risk and opportunity database, the individual risk scenario or opportunity is reviewed by risk management personnel using the same seven validity test questions. If the answer to any is “no” or “unknown,” the individual risk scenario or opportunity is not considered valid and the author should be queried for his/her intent so that it can be either modified or rejected, with rationale.

5.2.4 Develop Risk and Opportunity Narratives

While the risk or opportunity statement provides a concise description of the individual risk scenario or opportunity, this information is not necessarily sufficient to capture all the information that the identifier has to convey, nor is it necessarily sufficient to describe the concern in enough detail that risk management personnel can understand it and respond effectively to it, particularly after the passage of time. In order that enough context is recorded so that the individual risk scenario or opportunity can stand on its own and be understood by someone not otherwise familiar with the issue, a narrative description field is provided. The narrative description is format-free and elaborates on key circumstances surrounding the individual risk scenario or opportunity, including:

- Contributing factors
- Uncertainties
- The range of possible consequences or benefits

- Related issues such as what, where, when, how, and why.

The narrative description is also a place where the risk or opportunity identifier can suggest or recommend potential mitigations, actions, or other responses that they feel is most appropriate. It is usually the case that the identifier is an engineer with significant subject matter expertise in the affected asset, and it is important to capture that expertise, not only concerning the nature of the issue, but also its remedy. When a risk response opportunity action is recommended, the identifier should also record the rationale for the recommendation, preferably including an assessment of the expected risk shifting (e.g., from a safety risk to a cost risk) that would result.

In cases where the risk or opportunity is not actionable, the narrative should explain the conditions that make it so and the prospects for those conditions to change.

5.2.5 [Categorize the Risk or Opportunity using Risk, Opportunity, and Leading Indicator Taxonomies](#)

Once a risk or opportunity has been formalized in a valid risk or opportunity statement and an associated narrative, its components should be categorized using the relevant risk, opportunity, and leading indicator taxonomies incorporated into the organizational unit's risk database system. This enables the new risk or opportunity to be identified as cross-cutting if that is indeed the case. The use of taxonomies is discussed further in Appendix I.

5.3 CRM Step 2: Analyze

The sub-steps of the Step 2 *Analyze* step are as follows:

- Develop a risk scenario diagram (RSD) (or other equivalent accident sequence logic diagram) for the identified risk
- Probabilistic Approach
 - Analyze the Likelihoods of the Events in the RSD
 - Analyze the Performance of Each End State in the RSD
 - Address Epistemic Uncertainty (As Needed)
 - Classify the Individual Risk Scenario
 - Integrate the Individual Risk Scenario into the Risk Model and Analyze Aggregate Risks (As Needed)
 - Classify the Aggregate Risks
 - Determine the Risk Drivers
 - Display and Communicate Risk
 - Analyze and display analogous results for opportunities using the spider chart format
- Heuristic Screening Approach as a Precursor to the Probabilistic Approach
 - Estimate individual risk scenario likelihoods and consequences from historical data and expert judgment

- Aggregate individual risk scenario likelihood and consequence rankings to risk rankings for base-level organizational objectives and mandated requirements
- Propagate the aggregate risk rankings for base-level organizational objectives and mandated requirements to aggregate risk rankings for top-level organizational objectives and mandated requirements
- Determine the risk drivers
- Analyze opportunities
- Display and communicate aggregate risks and risk drivers
- Use heuristic results as input to a graded-approach probabilistic analysis

5.3.1 Implementing a Graded Approach to Analysis

This section provides general guidance on grading the various aspects of analysis based on Activity Class and other factors, such as activity phase or stage of development. The guidance is intended to apply to all categories of risk management and all types of risk, including:

- All risk management activity domains (program/project, institutional, enterprise)
- All mission/activity execution domains (safety, technical, security, cost, schedule, liability protection, public education, etc.)
- All stages along the mission/activity timeline (concept development, design of the solution, implementation of the solution, conduct of the operation, decommissioning of the operation)
- All different shades and gradations of risk (individual risk scenarios, leading indicators of unknown and/or underappreciated (U/U) risk,¹⁸ aggregate risk to organizational objectives and/or mandated requirements)

As relates to the characteristics of analysis scope, inclusiveness, and completeness, the determinants of the appropriate level of analysis are a function of Activity Class as defined in Section 4.11.2., which is closely related to Risk Tolerance Class as defined in NPR 8705.4. This is addressed within this section, as well as in all other applicable sections of Chapter 5, by means of graded approach guidance tables.

5.3.1.1 *Tailoring the Analysis Methodology to the Program/Project Life Cycle*

Figure 5-4 indicates the types of analysis that are generally appropriate, as a function of the life-cycle phase of a program or project, for cost, technical, and safety estimation. The seven phases of a program/project are defined in the NASA Systems Engineering Handbook as follows [1]:

- Pre-Phase A: Concept Studies

¹⁸ A guiding principle is to ensure that the consideration of leading indicators for U/U risks is conducted honestly and without a bias toward optimism, so as to avoid the severe underestimation of costs and task durations that according to Government watchdog agencies have plagued some of the Agency's largest and most complex programs [3], [4].

- Phase A: Concept and Technology Development
- Phase B: Preliminary Design and Technology Completion
- Phase C: Final Design and Fabrication
- Phase D: System Assembly, Integration and Test, and Launch
- Phase E: Operation and Sustainment
- Phase F: Closeout (Not Included in Figure 5-4)




Figure 5-4. Analysis Methodology Guidance Chart

A brief description of the analysis types is provided below. More detailed information on methods can be found in discipline-specific guidance, e.g., [5], [6], and [7]. Discussion of uncertainty can be found in Sections 4.7.3.1 and 4.7.3.2.

5.3.1.2 Estimating Methodologies for Technical Performance Measures

- **First-Order Estimating Methodology** - First-order estimates involve the use of closed-form or simple differential equations which can be solved given appropriate bounding conditions

and/or a desired outcome without the need for control-volume based computational methods. The equations may be standard physics equations of state or empirically-derived relationships from operation of similar systems or components.

- *Detailed Simulation Estimating Methodology* - Estimates using a detailed simulation require the construction of a model that represents the physical states of interest in a virtual manner using control-volume based computational methods or methods of a similar nature. These simulations typically require systems and conditions to be modeled to a high-level of fidelity and the use of “meshes” or network diagrams to represent the system, its environment (either internal, external, or both), and/or processes acting on the system or environment. Examples are computational fluid dynamics (CFD) and finite-element modeling.
- *Testing Methodology* - Testing can encompass the use of table-top experiments all the way up to full-scale prototypes operated under real-world conditions. The objective of the test is to measure how the system or its constituent components may perform within actual mission conditions. Testing could be used for assessing the expected performance of competing concepts or for evaluating that the system or components will meet flight specifications.
- *Operating Experience Methodology* - Once the system is deployed data gathered during operation can be analyzed to provide empirically accurate representations of how the system will respond to different conditions and how it will operate throughout its lifetime. This information can serve as the basis for applicable changes, such as software uploads or procedural changes, that may improve the overall performance of the system. Testing and detailed simulation may be combined with operating experience to extrapolate from known operating conditions.

5.3.1.3 Safety, Reliability, & Operations Estimating Methodologies

- *Similarity Estimating Methodology* - Similarity estimates are performed on the basis of comparison and extrapolation to like items or efforts. Reliability and operational data from one past program that is technically representative of the program to be estimated serves as the basis of estimate. Reliability and operational data are then subjectively adjusted upward or downward, depending upon whether the subject system is believed to be more or less complex than the analogous program.
- *First-Order Parametric Estimation* - Estimates created using a parametric approach are based on historical data and mathematical expressions relating safety, reliability, and/or operational estimates as the dependent variable to selected, independent, driving variables through either regression analysis or first-order technical equations (e.g., higher pressures increase the likelihood of tank rupture). Generally, an estimator selects parametric estimating when the system and its concept of operation are at the conceptual stage. The implicit assumption of parametric estimating is that the same factors that shaped the safety, reliability, and operability in the past will affect the system/components being assessed.

- *Detailed Logic Modeling Estimation* - Detailed logic modeling estimation involves “top-down” developed but “bottom-up” quantified scenario-based or discrete-event logic models that segregate the system or processes to be evaluated into discrete segments that are then quantified and mathematically integrated through Boolean logic to produce the top-level safety, reliability, or operational estimate. Detailed technical simulation and/or testing, as well as operational data, can be used to assist in developing pdfs for quantification of the model. Typical methods for developing such models may include the use of fault trees, influence diagrams, and/or event trees.
- *Statistical Methods* - Statistical methods can be applied to data collected during system/component testing or from system operation during an actual mission. This is useful for characterizing the demonstrated safety, reliability, or operability of the system. In addition, patterns in the data may be modeled in a way that accounts for randomness and uncertainty in the observations, and then serve as the basis for design or procedural changes that may improve the overall safety, reliability, or operability of the system. These methods are useful for answering yes/no questions about the data (hypothesis testing), describing associations within the data (correlation), modeling relationships within the data (regression), extrapolation, interpolation, or simply for data mining activities.

As mentioned in Section 4.7.2, the amount of rigor exercised in modeling and simulation may not be the same for all risk scenarios. In general, the level of rigor should increase with the importance of the scenario being evaluated.

5.3.2 [Develop a Risk Scenario Diagram](#)

Graded Approach Guidance			
Activity Class A+	Activity Class B	Activity Class C	Activity Class D
Comprehensive risk scenario diagrams (RSDs) are developed including all potentially pertinent initial conditions, events, end states, and impacted requirements.	Same as for Activity Class A+, except weed out pathways in the RSDs and events in the pathways that are known with reasonable confidence to be insignificant.	Simple RSDs including scenarios that have been important historically and new scenarios discovered during the activity life cycle that meet a minimum level of significance.	RSDs are not required, so long as scenarios that can affect other assets are sufficiently understood.

A risk scenario diagram is a flowchart with paths that start from the departure event of the risk and lead to different end states. Along each path, pivotal events are identified as either occurring or not occurring. Each end state is expressed in terms of the resulting levels of performance across the set of performance measures established for the objectives and mandated requirements of the activity in question. RSDs were first discussed in Section 3.2.2. Figure 5-5 reproduces Figure 3-3 from that section.




Figure 5-5. Schematic of a Risk Scenario Diagram (RSD)

RSDs are primarily intended to be communications aids to facilitate inter-organizational discussions of the risks. They do not necessarily translate exactly to the more detailed risk analysis modeling, but they do serve as a tool for helping ensure that the models cover all the important elements of the risk. At this stage of CRM, the intent of the RSDs is to identify the possible pathways that can lead to positive or negative effects on the successful achievement of organizational objectives and mandated requirements, but not yet to quantify them. Importantly, they do not yet contain any reference to mitigation or research options that are not part of the present plan.

5.3.3 Analysis of Individual Risk Scenarios

NPR 8000.4 [2] stresses:

“When possible, quantitative characterizations of performance and corresponding risk levels are preferable, as they more directly enable risk vs. benefit ‘analysis of alternative’ (AoA) evaluations that constitute the foundation of the RIDM support to decision making.”

Correspondingly, this handbook stresses the importance and value of managing risk quantitatively, using a “probabilistic” approach to risk analysis, aggregation, and decision-making. Nevertheless, given the practical limits within which any risk management activity must take place, there is a potential need for more qualitative or “heuristic” approaches when more rigorous probabilistic methods are impractical due to limits of time, resources, or data, recognizing that the results of such methods may differ from those of more rigorous methods. Therefore, the analysis of individual risk scenarios presented in this section is partitioned into two subsections, one presenting a probabilistic approach and another presenting a heuristic approach. In any case, a graded approach should be followed in the selection and application of risk analysis methods and level of detail, in order to optimize the robustness of risk management decision-making within

programmatic constraints. The relative pros and cons of applying more or less rigor to risk assessment and risk management is discussed in more detail in Section 2.2.5.

5.3.3.1 Probabilistic Approach

To analyze the individual risk scenarios that have been identified in the Risk Scenario Diagrams and determine the aggregate risk to the organizational objectives and formal requirements that have been identified in the RMP, two types of analysis approaches are discussed in Chapter 5, probabilistic and heuristic. Probabilistic approaches are based on a more formal development of risk scenario progression and a typically (though not necessarily) quantitative treatment of event likelihoods and end state performance. Simpler heuristic approaches rely, instead, on the use of historical experience, empirical data, logical aggregation models, and expert judgment to determine ranges within which the individual and aggregate risks may be expected to lie. For most activities that are large and complex, probabilistic approaches is recommended, possibly augmented with heuristic approaches in areas of lesser risk significance, or when time is of the essence.

Section 5.3.3.1 and its subsections provide guidelines regarding the use of probabilistic approaches, whereas analogous guidelines regarding the use of heuristic approaches are provided in Section 5.3.3.3 and its subsections.

5.3.3.1.1 Analyze the Likelihoods of the Events in the RSD

Graded Approach Guidance			
Activity Class A+	Activity Class B	Activity Class C	Activity Class D
For each event in the detailed RSDs, the probability of occurrence of the event is analyzed, using data, analyses, and expert opinion. Analyses can take the form of fault tree analyses, phenomenological modeling, and simulation. Epistemically uncertain performance parameters are earmarked for uncertainty analysis.	Same as for Activity Class A+, except clearly insignificant pathways are weeded out, the scope of data, analyses, and expert opinion used in the evaluation may be reduced, and epistemic uncertainty may be neglected in scenarios that are of lesser importance.	Same as for Activity Class B, except the RSDs are more simplified and epistemic uncertainty may be neglected.	Likelihoods are evaluated qualitatively and only for events in scenarios that could affect other assets.

The likelihoods of the events in the RSD depend on the nature of the event. They represent the significant aleatory uncertainties associated with the activity that have the potential to affect activity performance. There is a wide range of possible causes for these events, including human error, equipment failure, market fluctuation, supply chain disruption, cyberattack, natural phenomena, unexpected litigation, and more. Correspondingly, there is a wide range of approaches to the analysis of event likelihood.

In the engineering realm, fault tree analysis (FTA) and phenomenological modeling are among the most common. Fault tree analysis involves the use of Boolean logic modeling to decompose event likelihood into its possible causes for which data is available. For example, a loss of power event might require both the loss of main power and the failure to provide backup power. Each of these contributing events could be further decomposed into their causes until a set of events is identified having likelihoods that can be substantiated by data. Phenomenological modeling involves the modeling of aleatory behavior to produce, in analysis, variations that can be expected in the actual activity. For example, the likelihood that a piece of foam will shed from a cryogenic stage and cause damage can be modeled in terms of variations in the number, size, and times of occurrence of the shedding events, along with the trajectory of each foam piece, including impact location. A more benign example might be the missing of a launch window, which might be due to a schedule slippage, which in turn might be due to insufficient staffing of a verification activity.

In developing likelihoods for the events in an RSD, the following precedence of sources is recommended:

- Use project-specific data as a first choice.
- Use relevant data from other projects as a second choice.
- Use indirect or surrogate data as a third choice.
- Resolve conflicting results using expert judgment elicitation.

Event likelihoods may be expressed quantitatively or qualitatively, depending on the nature of the information that is available for each event. Qualitative descriptors for the likelihood of an event in a five-choice ranking system might be as follows: “very unlikely”, “unlikely”, “even-odds”, “likely”, and “very likely” with respect to the chance that the event will occur during the life cycle of the activity.

5.3.3.1.2 Analyze the Performance of Each End State in the RSD

Graded Approach Guidance			
Activity Class A+	Activity Class B	Activity Class C	Activity Class D
End state performance is analyzed in detail for every performance measure and every end state, using quantitative performance models. Analyses account for the physical effects of each event in the RSD, using phenomenological modeling, simulation, data, and expert opinion. Epistemically uncertain performance parameters are earmarked for uncertainty analysis.	Same as for Activity Class A+, except that clearly insignificant pathways are weeded out, the scope of data, analyses, and expert opinion used in the evaluation may be reduced, and epistemic uncertainty may be neglected in scenarios that are of lesser importance.	Same as for Activity Class B, except the RSDs are more simplified, the analysis may focus only on a limited set of key performance measures, and epistemic uncertainty may be neglected.	End state performance is evaluated qualitatively and only for scenarios and performance measures that could affect other assets.

The end states of the RSD should be characterized in terms of performance, for every performance measure defined for the activity. This is the characterization that matters, and that ties the analysis to the activity's objectives. Typically, in the normal course activity planning and preparation, performance-related modeling is conducted to show that the activity, if executed as planned, will produce the intended result. These models may be technical, such as design studies, test plans, and technology development plans, or they may be programmatic, such as budget or earned value analyses. In any case, these models can often be leveraged in the service of the risk model, not only by using them in their nominal form to analyze nominal performance, but also by modifying them as needed to account for the effects of the off-nominal events in the RSD.

Figure 5-6 illustrates the analysis of end-state performance for an individual risk scenario. The activity begins nominally, but branches to include the RSD for the individual risk scenario in question at the point where the departure event can occur. This leads to a spectrum of possible end states (including the nominal end state), each of which has some probability of occurrence and set of performance values across the defined performance measures. In the figure, the values of the performance measures are indicated by the placement of the vertical bars on the abscissas, and likelihood of each end state is indicated by the height of the bars.




Figure 5-6. Analysis of an Individual Risk Scenario, Neglecting Epistemic Uncertainty

The bottom-right of Figure 5-6 shows the aggregation of the performance results for each end state into a single performance distribution for each performance measure. It also shows that the risk of shortfall with respect to a given performance marker is the area of the distribution on the shortfall side of the marker. What is not shown in the figure, but which is important to understand, is that there may be a number of performance markers associated with each performance measure. This

was discussed in Section 3.3.1, which presented the example of two performance markers for a single performance measure, namely a performance requirement and a performance goal.

5.3.3.1.3 Address Epistemic Uncertainty (As Needed)

Graded Approach Guidance			
Activity Class A+	Activity Class B	Activity Class C	Activity Class D
Probability distributions are developed for all epistemically uncertain performance parameters. Monte Carlo analysis is conducted, taking performance parameter dependencies into account, resulting in the development of distributions of the risk of not meeting each performance marker of every performance measure.	Same as for Activity Class A+, except that epistemic uncertainty may be neglected in scenarios that are of lesser importance, the uncertainty distributions may be defined qualitatively, and the evaluation of uncertainty may focus on a limited set of key performance measures.	Major epistemic uncertainties are qualitatively identified.	Epistemic uncertainty may be neglected.

The evaluation of epistemic uncertainty entails the development of uncertainty distributions for those performance parameters whose true values are not known well enough to infer from them a definitive characterization of the risks to the objectives. As such, they represent a lack of knowledge about the activity, rather than variability within the activity. Epistemic uncertainty is typically associated with novelty or rarity, and the lack of data associated with each.

Epistemic uncertainty distributions may be determined quantitatively or qualitatively, depending on the nature of the information that is available for each event. Quantitative distributions are usually derived by updating noninformative priors with available data, or by determining the functional form of the distribution based on theoretical considerations and using expert opinion to decide upon the values of the distribution parameters. Qualitative distributions are derived on discrete scales. Figure 5-7 illustrates three ways of representing epistemic uncertainty, each representing a different combination of quantitative and qualitative scales.

Epistemic uncertainty is typically incorporated into risk analysis by embedding the (aleatory) risk model within a Monte Carlo shell that performs a large number of iterations, sampling from the epistemic uncertainty distributions with each iteration, to produce a correspondingly large number of risk results, each of which represents a distinct set of assumptions about what the true values of the performance parameters are (i.e., a distinct “model of the world”). Care should be taken to account for dependencies among the performance parameter values. Monte Carlo analysis is illustrated in Figure 5-8, which shows how epistemic uncertainty leads to distributions of risk results rather than point values that emerge from the aleatory risk model alone (see Figure 5-6).




Figure 5-7. Quantitative and Qualitative Representations of Epistemic Uncertainty

5.3.3.1.4 Classify the Individual Risk Scenario

Graded Approach Guidance			
Activity Class A+	Activity Class B	Activity Class C	Activity Class D
For each performance measure, each individual risk scenario is classified as Green/“Acceptable,” Yellow/“Marginal,” or Red/“Unacceptable” based on its relationship to the performance markers and individual-risk RTLs defined for the performance measure. Epistemically uncertain risks are resolved through the establishment of an acceptable level of confidence, or by using mean risks. The classification of the individual risk scenario as a whole is the worst-case classification across the affected performance measures.	Same as for Activity Class A+, except that epistemic uncertainty may be evaluated for only a limited set of key performance measures.	Same as for Activity Class B, except that classification may be performed for only a limited set of key performance measures, and major epistemic uncertainties may be handled qualitatively by increasing the classification level based on engineering judgement.	Classification is qualitative based on the principle of “Do no harm.”

Monte Carlo Shell




Figure 5-8. Analysis of an Individual Risk Scenario, Including Epistemic Uncertainty

As shown in Figure 5-6, when epistemic uncertainty can be neglected, the risk associated with each performance marker is a point value. For a performance measure with multiple performance markers associated with it, the classification of an individual risk scenario follows the discussion presented in Section 3.3.2, with the important caveat that RTLs for individual risk scenarios are sub-allocated from the “parent” RTLs for aggregate risk, as discussed in Section 3.3.4. In any event, regardless of the number of performance markers and associated RTLs defined for a given performance measure, the risk to the associated objective should resolve into one of the three risk acceptability classes: Green/“Acceptable,” Yellow/“Marginal,” or Red/“Unacceptable.”

When epistemic uncertainty is significant, the risk associated with each performance marker takes the form of an uncertainty distribution, as illustrated in Figure 5-8. In this case, it may not be possible to say with certainty whether or not the risk is within the RTL, because there may be distribution mass on both sides of it. The situation is illustrated in Figure 5-9, where the most that can be said about the risk is that there is some level of *confidence* that the risk is within the risk tolerance defined by the RTL.




Figure 5-9. Confidence that the Risk of Not Meeting a PM X Performance Marker Due to Individual Risk Scenario q is Within the Individual Risk Scenario RTL

Two approaches for managing this complication are common. One is to take the mean value of each distribution of the risk of not meeting a performance marker. Another, arguably preferable, approach is for the Activity Decision Authority to declare a level of confidence that is needed in order to accept the results, and consider the risk to be the point risk value at that confidence. Different performance measures can warrant different levels of confidence. In general, the Activity Decision Authority can be expected to require a high level of confidence in claims about safety risk, whereas a lower confidence in other types of claims might be acceptable. Both of these approaches (mean value, confidence) reduce the risk distributions to point values, enabling the classification of the individual risk scenario along the same lines as above, where epistemic uncertainty is neglected.

Once the individual risk scenario has been separately classified with respect to each of the performance measures it threatens, the individual risk scenario as a whole can be classified as the worst-case performance-measure-specific classification.

5.3.3.1.5 *Integrate the Individual Risk Scenario into the Risk Model and Analyze Aggregate Risks*

Graded Approach Guidance			
Activity Class A+	Activity Class B	Activity Class C	Activity Class D
The individual risk scenario is integrated into the existing risk model, accounting for competing risks, synergistic or cumulative risk effects, departure event sequencing, and other complexities associated with vulnerability to multiple individual risk scenarios. Performance parameter correlations are accounted for across the risk model. Monte Carlo analysis is performed on the integrated set of individual risk scenarios. U/U risk is recharacterized based on the current state of knowledge.	Same as for Activity Class A+, except that epistemic uncertainty may be evaluated for only a limited set of key performance measures.	Integration of RSDs into the risk model may be limited to only those major individual risk scenarios having synergistic or cumulative effects. Major epistemic uncertainties are qualitatively identified.	N/A.

The integration of an individual risk scenario into an existing risk model is illustrated conceptually in Figure 5-10, which shows a risk model that has been updated from having just one individual risk scenario (Individual Risk Scenario 1) to having two individual risk scenarios (Individual Risk Scenario 1 and Individual Risk Scenario 2). Like the risk model of Figure 5-8, which Figure 5-10 expands upon, Figure 5-10 is conceptual only and does not illustrate complexities such as the possibility of both individual risk scenarios being realized, or of Individual Risk Scenario 1 being prevented due to realization of Individual Risk Scenario 2. Guidance on handling such complexities can be found in [7].

Figure 5-10, illustrates the need to reevaluate epistemic uncertainty for the updated risk model as a whole, recognizing that each iteration of the Monte Carlo routine represents a distinct “model of the world” that must be internally consistent with respect to its performance parameters, across the full set of RSDs in the model. U/U risk, which is typically assessed as a multiplier on the known risk, should also be evaluated with each iteration. Figure 5-10 also shows that risk distributions can be developed for known risk as well as for total risk (known and U/U).

5.3.3.1.6 Classify the Aggregate Risks

Graded Approach Guidance			
Activity Class A+	Activity Class B	Activity Class C	Activity Class D
The aggregate risk to each performance measure is classified as Green/“Acceptable,” Yellow/“Marginal,” or Red/“Unacceptable” based on its relationship to the performance markers and RTLs defined for it. Epistemically uncertain risk is resolved through the establishment of an acceptable level of confidence, or by using mean risk.	Same as for Activity Class A+, except that epistemic uncertainty may be evaluated for only a limited set of key performance measures.	Same as for Activity Class A+, except that classification may be performed for only a limited set of key performance measures, and major epistemic uncertainties may be handled qualitatively by increasing the classification level based on engineering judgement.	Classification is qualitative based on the principle of “Do no harm.”

The procedure for classifying aggregate risks is similar to that for classifying an individual risk scenario (Section 5.3.3.1.4). The main differences are:

- RTLs for aggregate risk are used instead of those for individual risk scenarios.
- The aggregate *total risk* to each performance measure due to both known and U/U risks is classified. (The aggregate *known risk* to each performance measure can also be classified.)
- The risk to the activity as a whole can be classified as the worst-case performance-measure-specific classification.

5.3.3.1.7 Determine the Risk Drivers

A risk driver is a significant source of risk to one or more organizational objectives or mandated requirements. Operationally, a risk driver can be a single performance parameter, a single event in a risk scenario, a single leading indicator, a set of performance parameters collectively, a set of events collectively, or a set of leading indicators collectively that, when varied over their range of uncertainty, causes the aggregate risk to change from acceptable or marginal to a higher level of unacceptability. Specific examples might include:

- Random mechanical failure of one or more specified components
- An externally or internally initiated cyberattack
- Cross-cutting supply chain problems
- Unexpected economic conditions
- Excessive growth in subsystem interface complexity due to a necessary redesign

- Increase in time pressures during a project due to changing requirements
- Lack of management commitment to risk management, cross-cutting several projects
- Uncertainties in the human reliability models
- Uncertainties in the load capacity, extent of deformation, or thermal properties of certain materials due to environmental conditions
- Combinations of scenarios, leading indicators, model shortcomings, or parameter uncertainties that are individually of low concern but collectively cause the aggregate risk to change from acceptable to not acceptable

Risk drivers focus risk management attention on those potentially controllable situations that present the greatest opportunity for risk reduction. Often, risk drivers affect more than one individual risk scenario and cut across more than one organizational unit. Risk drivers are identified during the *Analyze* step of CRM and are used during the *Plan* step to devise effective risk response options. Risk drivers are determined after the integrated performance models have been created and executed by performing sensitivity analyses within the calculation of performance measures.

The identification of risk drivers may sometimes need to be performed in two steps or more. The first step looks at each parameter, event, or leading indicator individually to determine whether it is a driver by itself. If no drivers are identified by this process, then combinations of parameters, events and/or leading indicators are considered.

5.3.3.1.8 Display and Communicate Risk

A recommended means for displaying and comparing aggregate risks to different objectives is through the use of aggregate risk spider charts, an example of which is shown in Figure 5-11. The radial dimension on the spider chart (which is also commonly referred to as a radar chart or web chart) shows the risk of not meeting a top-level objective in terms of risk acceptability classification. The spider chart format is useful for displaying results for all the top-level objectives on one chart.

Aggregate risk spider charts should not be confused with performance spider charts, for which the radial dimension reflects performance measure values rather than risk values.

For a given top-level objective, the “top risks” (i.e., top individual risk scenarios) that threaten the objective can also be communicated using a spider chart. In this case, the chart itself is objective-specific, and the spokes of the chart show the risk acceptability classifications of the top risks. An example of a “top risk” chart is shown in Figure 5-12. The figure clearly shows that individual risk scenarios 1, 2, and 7 are unacceptable and should be prioritized for further risk reduction.




Figure 5-11. Aggregate Risk Spider Chart for an Activity




Figure 5-12. Top Risks Spider Chart for an Objective

5.3.3.2 Analyze and Display Analogous Results for Opportunities Using the Spider Chart Format

Evaluation of opportunity actions also involves the assessment of aggregate risks. In general, an opportunity action is considered to be desirable if the following is true:

- 1) The opportunity action reduces the aggregate risk of not accomplishing one or more of the organization's existing objectives from unacceptable (red) to acceptable (green) or at least marginal (yellow).

- 2) Alternatively, the opportunity action makes it possible to achieve a new objective that would previously have been considered to be impossible (red) but is now considered to be achievable with acceptable (green) aggregate risk or at least marginal (yellow) aggregate risk.
- 3) The opportunity action does not cause the aggregate risk of failing to accomplish any of the organization's other existing objectives to increase from acceptable (green) or marginal (yellow) to unacceptable (red).
- 4) The likelihood of the opportunity action being non-implementable is reasonably low (green or yellow), unless the magnitude of the benefit in item 1) or 2), above, is sufficiently high to justify a higher risk of the opportunity action not being implementable.

The overall results of the analysis of an opportunity may be displayed using a spider chart format like that in Figure 5-11. The evaluation leading to the results for the aggregate risk accounts for the possibility that the implementation plan might succeed or might fail, leading to different conditional results with probabilities P_S and $P_F = 1 - P_S$, respectively. Building on Figure 5-11, Figure 5-13 below presents a set of conceptual results for the aggregate risk before and after implementing a plan for a hypothetical strategic opportunity.




Figure 5-13. Aggregate Risk Spider Chart of the Effects of Seizing Opportunity Y

5.3.3.3 Heuristic Approach

Given the practical limits within which any risk management activity must take place, there is a potential need for qualitative or “heuristic” approaches to the analysis of individual risk scenarios when more rigorous probabilistic methods are impractical due to limits of time, resources, or data, recognizing that the results of such methods may differ from those of more rigorous methods. Whereas probabilistic approaches are based on a more formal development of risk scenario progression, heuristic approaches rely instead on the use of historical experience, empirical data,

logical aggregation models, and expert judgment to determine ranges within which the individual and aggregate risks may be expected to lie.

In addition, there could be cases where some of the performance measures are not amenable to mathematical modeling, or when uncertainties for some performance measures are unusually large, so that the use of expert judgment based on historical experience is more reliable than the use of mathematical models based on data that is specific to the project or activity being analyzed.

5.3.3.3.1 Analyze and Classify the Individual Risk Scenario

Graded Approach Guidance			
Activity Class A+	Activity Class B	Activity Class C	Activity Class D
RSDs are developed for each individual risk scenario. Each path through each RSD is separately analyzed for its overall likelihood of occurrence and its effects on performance, using historical data and expert judgement. A risk matrix approach is used to classify each path as “Acceptable,” “Marginal,” or “Unacceptable.”	Same as for Activity Class A+, but concentrating on just RSD pathways that are relevant to Graded Approach Activity Class B.	Same as for Activity Class A+, but concentrating on just RSD pathways that are relevant to Graded Approach Activity Class C.	Same as for Activity Class A+, but concentrating on just RSD pathways that are relevant to Graded Approach Activity Class D.

Before the determination of risk scenario likelihoods and consequences for any given organizational entity, it is assumed that RSDs have been developed for each identified individual risk scenario, indicating the possible ways the scenario might progress from departure event to ultimate consequences in terms of shortfalls with respect to the top-level objectives. Section 3.2.2.1 discusses the development of RSDs.

For each path through each RSD, the overall likelihood of occurrence of the path and the effects on performance should it occur are analyzed using historical data and expert judgement. The level of rigor needed will ultimately be determined by the experts’ ability to classify (and defend) the acceptability of the path as “Acceptable,” “Marginal,” or “Unacceptable,” as discussed below, so some iteration between the analysis of the path and the classification of the path might be required.

Once the likelihoods and consequences of each path through the RSD have been analyzed, the acceptability of the individual risk scenario as a whole is classified using risk matrices such as those in Section 3.3.6.1. Two important features of these matrices are:

- The risk acceptability classes are anchored to IRTLs, which are anchored to the RTLs that constitute the organizational unit’s risk posture.
- The classifications in the matrices (i.e., the matrix element colors) are conservative with respect to the PMKs and IRTLs. For example, there are Red/“Unacceptable” matrix elements to the left of PMK-R in Figure 3-17. This conservatism is intended to account for the additional uncertainty inherent in the heuristic approach relative to the probabilistic approach.

Because an individual risk scenario has the potential to threaten multiple objectives, the heuristic method uses as many risk matrices as there are threatened objectives.

The procedure for classifying the acceptability of the individual risk scenario is as follows:

- For each risk matrix, map each path through the RSD to a colored region of the matrix, based on its analyzed likelihood and consequences. If the rigor of the estimated likelihoods or consequences is insufficient to support a defensible mapping to a specific color, then the analysis should be iterated with increased analytical rigor.
- If any path maps to Red/“Unacceptable,” then the individual risk scenario is classified as Red/“Unacceptable” with respect to that objective.
- Otherwise, if there are multiple Yellow/“Marginal” mappings in a matrix, then the individual risk scenario is classified as Red/“Unacceptable” with respect to the objective unless a defensible argument can be made that the cumulative effect of the multiple Yellow/“Marginal” paths is still Yellow/“Marginal,” in which case the individual risk scenario is classified as Yellow/“Marginal” with respect to the objective. The nature of the argument will depend on the nature of the paths. For example, if the multiple Yellow/“Marginal” paths all produce similar consequences, they can be combined into a single path whose likelihood of occurrence is the sum of the individual path likelihoods. If this combined path still maps to Yellow/“Marginal,” then the individual risk scenario can be classified as Yellow/“Marginal.” For paths leading to dissimilar consequences the argument will necessarily be more complex.
- Otherwise, if there is only one Yellow/“Marginal” mapping in a matrix, then the individual risk scenario is classified as Yellow/“Marginal” with respect to the objective.
- Finally, if there are multiple Green/“Acceptable” mappings in a matrix, then the individual risk scenario is classified as Yellow/“Marginal” with respect to the objective unless a defensible argument can be made that the cumulative effect of the multiple Green/“Acceptable” paths is still Green/“Acceptable,” in which case the individual risk scenario is classified as Green/“Acceptable” with respect to the objective.

The result of this procedure is that each individual risk scenario receives an acceptability classification for each of the organizational unit’s top-level objectives. This allows spider charts of the type illustrated in Figure 5-12 to be developed.

5.3.3.3.2 Classify the Aggregate Risks

The procedure for classifying the acceptability of the aggregate risk to each objective using a heuristic approach parallels the procedure for classifying the acceptability of an individual risk scenario, with the following modifications:

- The IRTLs in the risk matrices are replaced by the corresponding RTLs. For example, the Probability matrix element boundaries of Figure 3-17 are changed to (from top to bottom) 3 × RTL-G, RTL-G, RTL-R, and RTL-R/3.

- For each risk matrix, map each path through *every* RSD to a matrix element, based on its analyzed likelihood and consequences.
- The resulting risk acceptability classifications pertain to the objective associated with the risk matrix, rather than to the individual risk scenarios.
- Any arguments made regarding the cumulative effect of the multiple Yellow/“Marginal” paths or multiple Green/“Acceptable” paths should account for the possibility of multiple individual risks being realized, potentially resulting in the kind of issues addressed in Section 3.2.3. For example, two different individual risk scenarios might each have the potential to impact cost, resulting in a potential cost impact equal to the sum of the cost impacts of the individual risk scenarios considered separately, were both individual risks realized. This particular issue (cumulative effects) is not a concern when analyzing the cumulative effects of the pathways in a single RSD, because by definition only one path of an RSD can be realized.

The result of this procedure is that each top-level objective receives a risk acceptability classification, accounting for all individual risk scenarios. This allows spider charts of the type illustrated in Figure 5-11 to be developed.

The heuristic approach to analysis is illustrated in Figure 5-14.

5.3.3.3.3 *Determine the Risk Drivers*

Although the identification of risk drivers may be more qualitative when using the heuristic approach rather than for the probabilistic approach, the overall process is generally the same. See Section 5.3.3.1.7 for details.

5.3.4 Analyze Opportunities

As discussed in Section 2.1, risk management in an opportunity-seeking culture involves finding a balance between seeking opportunity and accommodating risk within an acceptable risk posture. Two types of opportunity were discussed, one of which is tactical in nature and the other strategic:

- 1) Tactical opportunities provide a potential for reducing the risk to one or more of an entity’s existing objectives. For example, an organizational unit that has begun execution on a project may be presented with an unexpected chance to involve a partner organization that has specialized expertise not currently present in the organization unit, resulting in the task being performed more expeditiously and at lower cost.
- 2) Strategic opportunities provide an opening for the organization to promote new objectives that advance the organization’s overall mission. For example, the emergence of a new technology may open up possibilities for exploring further into the universe than was previously possible.

Both types of opportunity require an action to be performed in order for the opportunity to be realized: e.g., establishment of a contractual working relationship with the partnering organization or maturation and implementation of the new technology.




Figure 5-14. Heuristic Approach to Risk Analysis

For tactical opportunities, estimation of the magnitude of the benefit of an opportunity action requires consideration of both its positive effects and its negative effects. For example, an action to involve a partnering organization with specialized skills and experience may *reduce* the risk of exceeding the schedule and cost performance markers (its positive effect). However, it may also *increase* the risk from cyberattack by ceding control of the cybersecurity function to an organization outside NASA's direct control (its negative effect). For this example, the positive effect can be evaluated by estimating the degree to which the action results in a reduction in one or more individual risk scenarios that threaten to produce schedule and cost overruns, whereas the negative effect can be evaluated by estimating the risk magnitude associated with any introduced risks that threaten cybersecurity.

Like individual risk scenarios, opportunities, both strategic and tactical, can be cross-cutting in that they may represent a potential benefit beyond the organization within which they are first identified. As such, opportunities should be communicated along the same pathways and in the same forums as individual risk scenarios in order to maximize the recognition of their potential benefits and take advantage of the potential economies of scale associated with systemic implementation.

As discussed in the beginning of Section 5.3.3, NPR 8000.4 [2] stresses:

“When possible, quantitative characterizations of performance and corresponding risk levels are preferable, as they more directly enable risk vs. benefit ‘analysis of alternative’ (AoA) evaluations that constitute the foundation of the RIDM support to decision making.”

Correspondingly, this handbook recommends using a “probabilistic” approach to opportunity management. Nevertheless, given the practical limits within which any opportunity analysis activity must take place, there is a potential need for more qualitative or “heuristic” approaches when more rigorous probabilistic methods are impractical due to limits of time, resources, or data, recognizing that the results of such methods may differ from those of more rigorous methods. Therefore, the analysis of opportunities presented in this section is partitioned into two subsections, one presenting a probabilistic approach and another presenting a heuristic approach. The relative pros and cons of applying more or less rigor to risk assessment and risk management is discussed in more detail in Section 2.2.5.

5.3.4.1 Probabilistic Approach

In a probabilistic approach, the analysis of a tactical opportunity takes the form of Activity-Execution RIDM where there are just two alternatives: the “no action” alternative in which the risk profile remains as is; and the “opportunity seizing” alternative where the seizure of the opportunity is integrated into the (potential) baseline activity. For the “opportunity seizing” alternative, risk identification and analysis is then conducted altered baseline, most likely heavily leveraging the existing risk analysis of the “no action” alternative but modifying the analysis of existing risks as needed, and possibly adding new ones arising from the alterations.

This results in two integrated risk models, each along the lines of Figure 5-10: one for the “no action” alternative and one for the “opportunity seizing alternative.” The communication of analysis results, the deliberations, and the selection of a preferred alternative (i.e., whether or not to seize the opportunity) proceed along the lines of the existing RIDM process. A probabilistic example of Activity-Execution RIDM (in the context of risk response) is provided in Section 4.5.1 of Part 2 of this handbook.

5.3.4.2 Heuristic Approach

For the heuristic approach, the rationale behind the evaluation of tactical opportunities as a balancing between positive and negative effects is shown in Figure 5-15, using a risk and opportunity matrix format for illustration purposes. In the example shown in the left-hand chart, an opportunity has been identified that if successfully implemented will reduce an existing individual risk scenario from “yellow” (risk matrix element 20) to “green” (risk matrix element

10). However, it also introduces a new individual risk scenario, which is analyzed as “green” (risk matrix element 7). The net benefit of the opportunity, if successfully seized, is expressed as a rank, from 1 to 5, using the opportunity benefit ranking scheme of Table 5-I. Because the example opportunity reduces the existing individual risk scenario from a 20 to a 10 and the introduced risk is a 7, the net benefit of the opportunity satisfies the dual criteria in Table 5-I for a risk rank of 5, or “Very High.”

Separately, the likelihood of successful implementation of the opportunity, given a decision to seize it, is estimated and ranked according to the likelihood ranking scheme of Table 5-II. In the current example, the opportunity is ranked as “High,” with a numerical value of 4.

The overall benefit of the opportunity is expressed in terms of its placement on the opportunity matrix on the right side of Figure 5-15. The “Magnitude of Benefit” of the opportunity is 5, as determined from Table 5-I, and its “Likelihood of Successful Implementation” is 4, as determined from Table 5-II. Communicating opportunities using an opportunity matrix such as that in Figure 5-15 provides an analog to the communication of individual risk scenarios using a risk matrix.




Figure 5-15. Example Illustration of the Use of an Opportunity Matrix

Table 5-I. Example Net Benefit Ranking for Tactical Opportunities

Magnitude of Benefit of Analyzed Opportunity		
Rank	Decrease in Significance Rank for Targeted individual Risk Scenarios Given Success of Implementation	Maximum Significance Rank for Introduced Risks
5 Very High	≥ 10	7
4 High	7 – 9	9
3 Moderate	4 – 6	11
2 Low	1 – 3	13
1 Very Low	≤ 0	N/A

Table 5-II. Example Likelihood Ranking for Tactical Opportunities

Likelihood of Successful Implementation of Analyzed Opportunity	
Rank	Likelihood of Successful Implementation
5 Very High	0.8 – 1.0
4 High	0.6 – 0.8
3 Moderate	0.4 – 0.6
2 Low	0.2 – 0.4
1 Very Low	0.0 – 0.2

5.4 CRM Step 3: Plan

The sub-steps of Step 3: *Plan* are as follows:

- Generate risk response options
- Generate risk response alternatives
- Perform risk analysis of mitigation alternatives
- Deliberate and select a risk response alternative

5.4.1 [Generate Risk Response Options](#)

Graded Approach Guidance			
Activity Class A+	Activity Class B	Activity Class C	Activity Class D
Decision to accept, mitigate, watch, research, elevate, or close individual risk scenarios based on placement on performance-based risk matrices. Comprehensive inclusion of all reasonable options when risk mitigation is required.	Same as for Activity Class A+, except the number of risk mitigation options considered may be limited to as two per performance measure, as long as all significant risks are considered.	Same as for Activity Class A+, except the number of risk mitigation options considered may be limited to as few as one per performance measure, as long as all risks historically significant to safety, mission success, or programmatic requirements are considered.	Selection of response options based on qualitative considerations is sufficient.

Risk response options are the individual responses from which candidate risk response alternatives are generated. Each risk response option pertains to one of the following risk disposition types specified in NPR 8000.4C: Accept, Mitigate, Watch, Research, Elevate, and Close.

Risk Response Dispositions

Accept – Applies to the totality of the performance risk. A decision to accept the performance risk means that no risk response is needed at that time.

Mitigate – Applies to risk drivers. A decision to mitigate a risk driver means that positive action will be taken to reduce its impact on performance risk.

Watch – Applies to risk drivers. A decision to watch a risk driver means that a watch plan is developed for one or more observables related to the risk driver, possibly including contingency plans that are to be executed contingent on specific values of the observables.

Research – Applies to risk drivers. A decision to research a risk driver means that a research plan is developed to investigate the risk driver, possibly including contingency plans that are to be executed contingent on specific findings.

Elevate – Applies to risk management decisions. A decision to elevate risk management decision making means that a unit's performance risk cannot be adequately managed by that unit.

Close – Applies to individual risk scenarios. A decision to close an individual risk scenario means that the risk drivers in the risk no longer exist or are no longer cost-effective to watch.

The following bullets provide guidance on when to apply each disposition:

- **Accept** – A risk response of *Accept* indicates that no risk management action needs be taken, given the current analyzed performance risk. This is typically because the aggregate risks associated with the organizational objectives are all within tolerable levels, reflecting an activity that is on track to accomplish its objectives within established risk tolerances. Within the risk model, this means that at the time of the risk analysis, none of the identified risk drivers are of sufficient magnitude to create intolerable risk to objectives or requirements.

However, a risk response of *Accept* does not mean that no risk management action relating to the existing risk drivers will be needed in the future. As the activity proceeds, additional conditions and departures may be identified that compound the effects of existing risk drivers in a manner that produces intolerable risk. In such cases, risk drivers that previously did not warrant a risk response might now be the most attractive targets for reducing requirement risk.

As the “no action” risk response option, the *Accept* option does not combine with other options when generating candidate risk response alternatives. An option of *Accept* applies to the entirety of the activity's risk posture and is superseded by any other risk response.

A risk response of *Accept* must be documented by the organizational unit, including the assumptions and conditions on which it is based.

- **Mitigate** – A risk response option within *Mitigate* is the taking of positive action to address the activity's risk. This is typically because the aggregate risk to one or more organizational objectives is outside tolerable bounds. However, it is important to allow for the possibility that mitigation may also be employed simply because an opportunity exists to reduce

performance risk even when it is within bounds. Mitigation options typically address one or more risk drivers, and are focused on improving the performance risk where it is most in need of improvement, without producing too large a collateral increase of performance risk in other areas (e.g., cost and schedule). Because mitigation options can address one or more risk drivers, and because a single risk driver can be present in a number of individual risk scenarios, a single mitigation option can potentially be responsive to a substantial number of individual risk scenarios.

Mitigation can be classified into two broad categories: departure prevention and consequence reduction¹⁹. Departure prevention refers to those risk response options that, if successfully deployed, prevent or reduce the likelihood of the departure event, and therefore the likelihood of the performance measure shortfalls associated with that event. Consequence reduction refers to those risk response options that, if successfully deployed, reduce the severity of the consequence produced by the departure, and therefore, of the magnitude of the associated performance measure shortfalls. Departure prevention and consequence reduction are illustrated graphically in Figure 5-16.

Most often, successful mitigation can be accomplished without having to change the derived objectives or mandated requirements, or the risk tolerances, that have been allocated to the organizational unit that is seeking to obtain mitigation of a risk. If the controls needed to accomplish this mitigation are not obvious and a variety of options exist, then Activity-Execution RIDM would be initiated (will be discussed further in Section 5.4.3). As explained in Section 2.2.5, this form of RIDM does not entail any rebaselining of the existing objectives or requirements or of the risk tolerances assigned to these objectives and requirements.




Figure 5-16. The “Mitigate” Risk Response Disposition

In some cases, a mitigation option can affect achievement of the derived objectives and requirements that flow down to organizational units at the next lower level of the NASA

¹⁹ In other contexts, the term “mitigation” refers only to consequence reduction, and is distinct from the term “prevention,” which refers to departure prevention. However, as discussed here and in NPR 8000.4, the “Mitigate” risk response disposition encompasses both departure prevention and consequence reduction.

hierarchy. When this is the case, implementation of the mitigation option includes the negotiation of a rebaselined set of derived objectives and non-objectives-based requirements among the affected units. In situations where risk drivers have been elevated from lower levels in the organizational hierarchy, it is not unexpected that mitigation would involve changes to the objectives and requirements, given that the rationale for elevation is the inability of the lower level to manage the risk to the very same objectives and requirements within their scope of authority and capability. However, it may also be the case that an organizational unit, when managing their own internally identified risk, identifies an attractive mitigation option that entails rebaselining of derived objectives and requirements. In either situation, all affected organizations should participate in defining the mitigation option(s), and Activity-Rebaseline RIDM should be initiated once approval is obtained from the activity decision authority.

In more extreme cases it might be advantageous for an organizational unit to consider mitigation options that go beyond the scope of the design solution chosen during Activity-Planning RIDM. Candidate alternatives would typically include the contending alternatives from the original RIDM activity, but might also include previously discounted alternatives that are now attractive, or other alternatives not previously considered but which, due to changed conditions, are now attractive. In these cases, re-execution of Activity-Planning RIDM generally produces entirely new sets of derived objectives and/or requirements flowing down to lower-level units in the NASA hierarchy, and the costs associated with such a major shift must be factored into the risk analysis of alternatives.

In all cases, mitigation plans are documented, including the appropriate parameters that will be tracked to determine the effectiveness of the mitigation.

- **Watch** – A risk response of *Watch* identifies one or more risk drivers that will be monitored according to a documented set of tracking requirements that include, at a minimum, the specific parameters to be watched and a monitoring schedule according to which the parameters will be observed. Additionally, depending on the circumstances, the watch plan may include contingency plans or other types of deferred decisions that will be invoked conditional on the results of the monitoring activity.

Watching entails the periodic updating of the risk analysis with current values of the watched parameters, according to the monitoring schedule. Because the activity's risk is analyzed using a single, integrated risk model, it may be efficacious to coordinate, as much as practicable, the monitoring schedules of the watched parameters, so that the risk analysis is updated and evaluated in a "batch" fashion.

- **Research** – A risk response option of *Research* applies to one or more risk drivers whose uncertainties are large enough that they interfere with robust risk management decision making. The *Research* option seeks to reduce uncertainty concerning some aspect of a risk driver by actively generating additional information about it. It entails the development of a research plan that identifies the subject to be researched, the specific parameters about which information is expected to be generated, and a research schedule including timeframes for results (and integration of the results into the risk model). Additionally, like

the *Watch* option, the Research option may include contingency plans or other types of deferred decisions that will be invoked depending on the results of the research.

- **Elevate** – A risk response of *Elevate* transfers the management of a performance risk to the organizational unit at the next higher level. Elevation occurs when no satisfactory combination of *Mitigate*, *Watch*, and *Research* options can be found that return the risk to tolerable levels. The Elevate option recognizes that the inability to manage performance risk at one level of the NASA hierarchy directly impacts the performance risk at the next higher level. Elevation is a pathway into the *Identify* step of the organizational unit at the higher level.

It is expected that the *Elevate* option will typically be combined with a *Watch* option that monitors the status of any risk drivers that may be associated with the unmanageable performance risk. In addition, the Elevate option may be combined with other options that address some fraction of the risk, though not enough to bring it to within tolerable levels. An *Elevate* option also entails coordination of the risk modeling activities of the organizational unit that is elevating the management of risk, and the unit to which the management of risk has been elevated. It is expected that the analysis of risk response alternatives will involve a coordinated risk analysis effort by both levels. As a practical matter, when elevating, the elevating unit should propose alternatives that it considers attractive but which exceeds the scope of its authority to implement.

A risk response of *Elevate* should only be made in response to an inability of the organizational unit to effectively manage performance risk at its level in the NASA hierarchy. As such, *Elevate* should not be proposed as an initial option. Instead, it should be reserved for situations in which the available risk response options have been analyzed and shown to be inadequate.

- **Close** – A risk response of *Close* applies to individual risk scenarios whose risk drivers no longer exist or are no longer cost-effective to watch. This can occur when their probability has been reduced below a defined level of insignificance; the consequence potential has been reduced below a defined level of insignificance; or the event has occurred, thus becoming a problem rather than a risk management issue (and is tracked as such). Closing an individual risk scenario indicates not only that it is currently not a significant contributor to performance risk, but that there is no expectation that it will be a significant contributor to performance risk in the future.

5.4.2 Generate One or More Risk Response Alternatives

Experience has shown that risk responses can be multidimensional, involving a number of discrete responses that act together to reduce performance risk. Risk responses may consist of a number of individually defined risk response options, each of which is of a particular risk disposition type specified in NPR 8000.4C. For example, the best response to an individual risk scenario that has large uncertainty may involve a combination of *Research* and *Mitigate*. In addition, since it is often the case that mitigation is the response of choice for more than one individual risk scenario, the alternative of choice may involve a combination of mitigation options integrated in a way that produces an overall synergistic effect. Thus, while the process of generating a set of candidate risk

response alternatives consists of first generating a set of candidate risk response options, the alternatives to be analyzed will generally involve one or more combinations of these options integrated collaboratively together. Figure 5-17 illustrates this schematically.




Figure 5-17. Relationship between Risk Response Options and Risk Response Alternatives

Theoretically, for N risk response options it is possible to define 2^N candidate response alternatives (including the no-action alternative). However, in practice, it will usually be possible to accept a single alternative if the cost is not too high and/or if the best solution is obvious, thereby bypassing the need for Activity-Execution RIDM. If that is not the case, the candidate risk response alternatives can most often be constrained to a reasonable number by downselecting attractive alternatives that:

- Address the performance risk of multiple organizational units
- Address all (or most) of the organizational objectives whose risk, as measured by its associated performance measures, is outside tolerable levels
- Introduce less risk in other performance areas (e.g., cost, schedule) in order to achieve the intended risk reduction.

5.4.3 Perform Risk Analysis of Mitigation Alternatives

Graded Approach Guidance			
Activity Class A+	Activity Class B	Activity Class C	Activity Class D
The models for analysis of mitigation alternatives start from the risk models developed in the <i>Analyze</i> step for Activity Class A+ and include modifications so as to be able to account for the mitigation alternatives being considered.	Same as for Activity Class A+, except start from the Activity Class B risk models and mitigation alternatives.	Same as for Activity Class A+, except start from the Activity Class C risk models and mitigation alternatives.	Same as for Activity Class A+, except start from the Activity Class D risk models and mitigation alternatives.

For risk response alternatives that include mitigation, a risk analysis of the alternatives is conducted to determine the residual risk for each performance measure and to ensure that those residual risks comply with the associated risk tolerances.

Under many circumstances, to manage day-to-day risks there is no need for a formal RIDM component within the CRM *Plan* step. That is the case if either of the following two situations pertains:

1. None of the individual risk scenarios are judged to be unacceptable on their own, and the magnitudes of the aggregate risks to the entity's organizational objectives and formal requirements (or more specifically the risks of not meeting the performance markers) are within their respective risk tolerances,
2. Any departures from Situation 1 can be corrected through simple and obvious fixes.

There may be occasions, however, where significant changes in existing conditions, including the possible emergence of new risks, cause risk tolerances for one or more of the performance measures to be exceeded, and additionally a correction of these departures requires a rigorous analysis of risk response alternatives. In those cases, there would be a need to initiate Activity-Execution RIDM within the CRM *Plan* step. When this occurs and new controls are formulated as a result, the design of the system or solution being developed to accomplish the activity is amended to include the new control(s), and the risk management plan is revised accordingly.

On still rarer occasions, it might be determined that conditions have changed so much, or that there are new risks that are so significant, that it is not possible to bring the aggregate risks of not meeting the designated performance marker (e.g., PMK-R or PMK-G) set-point values within their respective risk tolerances. If that is the case, a decision has to be made by the activity decision authority as to whether to approve and authorize a rebaselining of the activity. If the decision is affirmative, then Activity-Rebaseline RIDM is initiated to rebaseline performance requirements, performance goals, and/or risk tolerances, or waive mandated requirements. This decision is the prerogative of the activity decision authority. In general, Activity-Rebaseline RIDM involves more time and effort than Activity-Execution RIDM, but in most cases, it involves less time and effort than Activity-Planning RIDM.

Activity-Execution RIDM can be characterized as a scaled-down version of Activity-Planning RIDM, Part 2, Analysis of Alternatives (Section 4.2.2). The main difference stems from the fact that the alternatives being considered during Activity-Planning RIDM are typically very broad in scope and often involve fundamentally different concepts, whereas that is generally not the case during activity execution. During Activity-Planning RIDM, each alternative may require its own unique risk analysis starting more or less from scratch and requiring the development of uniquely different risk models. During activity execution, the alternatives being considered are typically narrower in scope, as they involve modifications to an already established concept with an already established design basis or solution approach. Such narrower alternatives can in essence be defined as "alternative risk-control solutions / measures" within that overall activity design basis and solution approach, and are considered and evaluated for the specific objective of mitigating a significant individual or aggregate risk in an optimal fashion, i.e., in a way that balances their risk reduction benefit with respect to the cost and schedule to be expended for their implementation.

Because of the narrower scope, Activity-Execution RIDM tends to rely more on sensitivity analyses starting from the integrated risk analysis model that has already been developed as part of the CRM *Analyze* step. Some of the submodels may have to be modified or even reconstituted to accommodate the mitigation approaches being considered, but the overall modeling framework stays intact.

As in Activity-Planning RIDM, the purpose of analyzing the risk response alternatives during activity execution is to support decision making. The goal is a robust decision, where the decision-maker is confident that the selected risk response alternative is actually the best one, given the state of knowledge at the time. This requires the risk analysis to be rigorous enough to discriminate between alternatives, especially for those performance measures that are determinative to the decision.

Figure 5-18 illustrates the invocation of Activity-Execution RIDM within the CRM *Plan* step, along with the possible rebaselining of the activity if needed.




Figure 5-18. Invocation of RIDM within the CRM Plan step

5.4.4 If Needed, Deliberate and Select a Mitigation Alternative

If it is determined that Activity-Execution RIDM is needed to downselect from a number of risk response alternatives, then it will be necessary to perform the deliberation sub-step. Deliberation and selection of a risk response generally proceeds along the same lines as during Activity-

Planning RIDM, Part 3, Risk-Informed Alternative Selection (Section 4.2.3). The activity execution version of deliberation and alternative selection tends to be scaled down from the Activity-Planning RIDM version, because although the decisions made during both activity start and activity execution involve tradeoffs over the same mission execution domains and over the same performance measures, the tradeoffs involve fewer parameters when considering mitigation alternatives than when considering entirely different concepts.

In addition to before-and-after risk values, other information captured during deliberation should be summarized and forwarded to the decision-maker, including:

- The Pros and Cons of Each Alternative – An itemized table of the pros and cons of each alternative is recommended for the contending alternatives, because it enables conflicting opinions to be documented, and captures elements of subjective value to the deliberators.
- Individual Risk Scenarios Introduced by Each Alternative – Mitigation alternatives can potentially generate new individual risk scenarios as the cost of addressing existing requirement risk. These should be identified and communicated to the decision-maker so that he or she understands the downside of each alternative.
- Cost-Benefit Tradeoff for Each Alternative – The cost of implementation/execution of the alternative, compared to the risk reduction benefit,
- It may be the case that no risk response alternative is available that reduces requirement risk to tolerable (or at least marginal) levels. In this case, elevation of the risk decision to the next level of the NASA organizational hierarchy is necessary. This situation would be documented, along with any other measures that are taken to at least partially address the intolerable requirement risk.

There may also be situations that endanger the activity but are outside the activity execution purview of accomplishing defined organizational objectives and non-objectives-based mandated requirements. Examples of these include poorly defined or missing requirements and requirements creep. In such cases it may be necessary to proceed to the Activity-Rebaseline RIDM activities that lead to recommendations for waiving or adjusting unneeded or conflicting formal requirements. The decision to rebaseline the requirements would be documented in the risk database and in a risk response document.

5.4.5 Implement the Risk Response

The key element of implementing the CRM risk response resides in setting up the roles and responsibilities to ensure that the response is carried out as intended in an effective and timely manner. Principally, this involves defining, empaneling, and obtaining commitments from the entities that are *responsible*, *accountable*, *consulted*, and *informed*:

- People who are *responsible* must complete the risk response task successfully, on time, and within budget. Those who have ownership of individual risk scenario items and who handle plan implementation directly (the risk owners) support the risk response process by providing periodic risk updates to the risk manager, briefing the risk board on current progress, and receiving / providing input at contractor RM activities.

- The person who is *accountable* (a single individual) must sign off or approve when each major subtask in the risk response is complete. That person must make sure that responsibilities are assigned for all related activities.
- People who are *consulted* (risk boards, risk managers, technical reviewers) are relied upon to give input before each major subtask is completed and signed-off on.
- Those who are *informed* (any person or group of people who have a stake in the outcomes) generally include people from other program or project organizations who are affected by the cross-cutting nature of the risks being responded to.

The implementation plan is reviewed and updated as the activity moves through each key decision point, with special emphasis on the following areas:

- Assessment level of detail and focus (risk tolerance changes as the program progresses)
- Review cycle/board frequency (as the activity moves toward completion, the frequency must often increase)
- Reporting formats and RM analysis tool capabilities
- Interaction with contractor and lower-level / higher-level risk processes

The existing risk management plan and the internal control plan are updated to incorporate the selected alternative. The update focuses on identifying, monitoring, and fine-tuning the key processes needed to ensure that the selected alternative is implemented as intended, and that the key assumptions made in the analysis of the alternative remain valid.

5.5 CRM Step 4: Track

Graded Approach Guidance			
Activity Class A+	Activity Class B	Activity Class C	Activity Class D
Tracking of individual risk scenarios, leading indicators, and performance measures in a timely manner. Tracking includes risks in all risk management domains (program/project, institutional, enterprise) and all mission execution domains (safety, technical, security, cost, and schedule), and concentrates on realization and operational stages of the life cycle.	Same as for Activity Class A+, consistent with the individual risk scenarios, leading indicators, and performance measures identified for Activity Class C.	Same as for Activity Class A+, consistent with the individual risk scenarios, leading indicators, and performance measures identified for Activity Class D.	Tracking of individual risk scenarios and leading indicators in a timely manner.

The objective of the *Track* step is twofold. It entails:

- Tracking the progress of the implementation of selected risk responses
- Tracking observables, related to performance measures and risk drivers, that are affected by the selected risk responses.

As such, the *Track* step ensures that data are generated to monitor not only the implementation status of risk response options, but also their effectiveness once implemented.

Tracking applies to the *Mitigate*, *Watch*, and *Research* risk response option types. The option types *Accept*, *Close*, and *Elevate* do not have tracking requirements associated with them. The nature of tracking is a function of the option type for which the tracking is being performed:

- **Mitigate** – Mitigation produces a modification to the baseline project plan that reflects the implementation of the selected mitigation option(s). As such, implementation is expected to be integrated into the project schedule of the responsible organizational unit, and progress tracked by project management processes within that unit. The progress should be communicated to the risk management functions of other organizational units so that all affected risk management functions have an awareness of the current status/configuration of the activity.

Mitigation also entails the scheduled monitoring of observables related to the effectiveness of the mitigation option(s). These observable quantities should also be communicated to the risk management functions of other affected organizational units. Monitoring enables risk management to assess the actual risk reduction relative to the forecasted risk reduction and the actual risk cost relative to the forecasted risk cost.

The actual risk reduction relative to the forecasted risk reduction – Mitigation options are implemented with the intent of reducing performance risk to the level forecasted by the risk analysis, conducted during the *Plan* step, of the selected risk response alternative. Observables selected for tracking should enable the actual performance risk reduction to be assessed and compared to that forecasted during *Plan*. These observables are expected to be directly related to the risk drivers that the mitigation options address.

The actual risk cost relative to the forecasted risk cost – Mitigation usually requires the acceptance of an increased level of requirement risk in some areas (e.g., cost and schedule). When this is the case, it is expected that these increases will be reported as new individual risk scenarios in accordance with the *Identify* step. Observables should also be selected that enable the monitoring of actual performance risk increase relative to that which is forecasted. These observables typically will not directly relate to the risk drivers that the mitigation options address; rather, they will tend to relate to low-risk areas of the activity where margin exists that can be sacrificed in the service of an improved overall performance risk posture.

- **Watch** – A decision to watch a risk driver entails the scheduled monitoring of observables related to that risk driver that can be used to assess the current performance risk and the contribution of the risk driver to that risk. Tracked parameters serve as early warning indicators so that further action can be taken. This enables timely execution of contingency plans or other types of deferred decisions that may be invoked conditional on the results of the monitoring activity. Tracked parameters should be communicated to the risk management functions of all affected organizational units.

In contrast to the *Mitigate* risk response option type, the *Watch* option type does not involve changes to the baseline project plan, and consequently does not involve the monitoring of implementation.

- **Research** – A decision to research a risk driver produces a research plan whose implementation should be tracked, and the scheduled monitoring of observables related to the research that, like *Watch*, can be used to assess the current performance risk and the contribution to that risk of the risk drivers associated with the research. Tracked parameters should be communicated to the risk management functions of all affected organizational units.

Tracking data can be used to construct requirement risk tracking charts that show how requirement risk increases and decreases over time as new individual risk scenarios are identified and responses are implemented. Figure 5-19 illustrates such a chart.




Figure 5-19. Requirement Risk Tracking Chart

5.6 CRM Step 5: Control

Graded Approach Guidance			
Activity Class A+	Activity Class B	Activity Class C	Activity Class D
Control both individual risk scenarios and performance measures in a timely manner. Control activities focus first on minor corrective actions that do not require management intervention, but allow for initiating a replanning of the response and reiteration of the CRM steps if necessary.	Same as for Activity Class A+, consistent with the individual risk scenarios and performance measures identified for Activity Class B.	Same as for Activity Class A+, consistent with the individual risk scenarios and performance measures identified for Activity Class C.	Timely control of individual risk scenarios, and timely responses to leading indicators.

The objective of the *Control* step is to evaluate the tracking data to determine whether or not risk responses are being implemented as planned, and if so, whether or not they are effecting the anticipated changes in targeted risk drivers and in the performance risk generally. Control includes an assessment of the need to take action to keep the relevant risk responses on track. These actions are kept within the control function unless it is clear that the objective of the risk response cannot be attained within the current plan. If that is the case, the *Plan* step is reinitiated and a new or modified risk response alternative is selected for implementation.

Because the *Control* step is focused on responding to the tracking data, it too is a function of risk response type:

- **Mitigate** – The role of risk management regarding control of implementation is primarily one of monitoring progress, evaluating the potential risk associated with departures from the implementation plan, implementing contingencies when needed, and making small changes in the plan when needed that do not require reinitiation of the *Plan* step. As mitigation options are implemented, it is the function of the *Control* step to evaluate the updated risk model in light of the tracked parameters and assess the degree to which they have successfully mitigated the effects of the risk driver(s) they address. If the assessed performance risk falls short of that forecasted during *Plan*, *Control* acts within the scope of the selected alternative to achieve, at least approximately, the intended result.
- **Watch** – In the case of *Watch* options, the *Control* step evaluates the watched parameters and, as appropriate, executes the contingency plans or other deferred decisions according to pre-established criteria.
- **Research** – Like *Mitigate*, *Research* involves the execution of a plan of action (in this case, the research plan) whose implementation is expected to be integrated into the activity plan. Therefore, the role of risk management regarding the control of research option implementation is analogous to that for mitigation. Also, like *Watch*, *Research* involves the execution of contingency plans or other deferred decisions based on an evaluation of the researched parameters relative to pre-established criteria.

To some extent, the CRM *Control* step may involve the implementation of new controls in order to address new risks and to ensure that the responses to them are implemented successfully. When that is the case, it will be necessary to ensure that the new controls are implemented seamlessly into the organization's existing internal control structure. Refer to Section 2.2.5 for a discussion of the interfaces between risk management and internal controls in the context of implementing controls that address responses to new risks.

5.7 Communicate and Document

Graded Approach Guidance			
Activity Class A+	Activity Class B	Activity Class C	Activity Class D
Establishment of communication protocols within the multi-organizational RM team including, at a minimum, regularly scheduled, weekly or bi-weekly cross-organizational meetings. Communication with organizational management on results, decisions, and associated rationale. Recommendations to organizational management on reformulation/reallocation of objectives, requirements, and risk tolerances if deemed advisable. Comprehensive documentation that includes the rationale behind all recommendations to management and all RM-related decisions reached within the RM team.	Same as for Activity Class A+.	Same as for Activity Class A+, but regularly scheduled meetings may be less frequent.	Regular communication between risk manager, participating entities, and organizational management. Documentation that includes the rationale behind all recommendations to management and all RM-related decisions made by the RM Manager.

Communication and documentation are central to CRM, and are integrated into each of the five CRM steps of *Identify*, *Analyze*, *Plan*, *Track*, and *Control*. Each of these steps involves the generation of information that must be properly documented and communicated to the appropriate personnel at the appropriate time, using appropriate standardized communication aids to assure that the intended meaning has been conveyed.

5.7.1 [Communication of Risk Information and Deliberations](#)

Throughout CRM, communication takes place among stakeholders involved in risk management, organizational management, and systems engineering to make sure that risks are effectively managed during implementation of risk responses. As discussed in previous subsections, communication can take place in a variety of forums, ranging from informal meetings, phone calls, and emails among personnel within an organizational unit, to technical interchange meetings involving personnel from numerous units in the NASA hierarchy and the authoring and dissemination of detailed reports. A graded approach is appropriate to determining the scale and formality of a given forum. In general, forums should facilitate dissemination of information to

the relevant affected parties, and provide ample opportunity for discussion and feedback to assure that issues are fully understood at a level that supports the decision making needs of all participants.

Inter-organizational communication is an integral part of CRM across the NASA organizational hierarchy. Different organizational units at different levels in the hierarchy must work together to ultimately achieve the top-level objectives that motivate their derived lower-level objectives and mandated requirements. Throughout the CRM process, communication takes place among these units to assure that:

- Every unit is aware of the individual risk scenarios that affect its performance risk.
- Individual risk scenarios are integrated into the risk analyses of the affected units in a consistent fashion (i.e., using consistent modeling assumptions).
- Individual risk scenarios are aggregated to performance risks correctly at each level (i.e., with each acquiring organizational entity taking into account the aggregation rationale and assumptions used by its providing organizational entities).
- Every unit's risk driver list is available to other units and is updated according to an established schedule.
- Every unit that is affected by a risk driver, or by the proposed responses to a risk driver, is adequately engaged in planning a response to it, including deliberation and selection of a response for implementation.
- Every unit is aware of the risk responses that affect its performance risk and/or its risk analysis.
- Elevation of risk management decisions is timely and unambiguous.
- Graphical illustrations pertaining to communication between organizational entities have been presented and described earlier in Section 5.3.3.1.8.

Standardized communication aids should be developed that support the information needs of the decisions they support. Examples are:

- Risk and opportunity statements, which include:
 - The conditions leading to the risk or opportunity
 - The departure event(s) or action(s) required to cause the risk or opportunity to emerge
 - The relevant leading indicators
 - The negative impact of the risk or positive benefit of the opportunity in terms of its effect on one or more objectives or requirements
 - The affected requirements or objectives
- Risk and opportunity narratives, which include:
 - Contributing factors

- Uncertainties
- The range of possible impacts or benefits
- Suggested or recommended responses
- Related issues such as what, where, when, how, and why
- Risk burn-down schedules for each performance measure
- Populated risk taxonomies showing the distribution of individual risk scenarios among the taxons of the specified taxonomies (e.g., see the taxonomies in Appendix I)
- RSDs that enumerate the spectrum of possible outcomes (and their likelihoods) resulting from an individual risk scenario's departure event
- Risk driver lists to support risk response planning
- Tables and charts of performance risk and risk tolerances for the contending risk response alternatives
- Risk tracking aids such as performance risk tracking charts that show the trajectories of each performance risk with respect to its risk burn-down profile.

Risk communication protocols should be negotiated among involved organizational units and documented in the RMP. This includes scheduled periodic reporting of risk information, such as to the unit at the next higher level, as well as protocols for risk reporting in response to triggers such as the exceedance of an elevation threshold.

5.7.2 [Documentation of Risk Information and Deliberations](#)

Principal decisions and associated rationale stemming from the communications described in the preceding subsection are documented as a matter of good practice. In addition, each organizational entity that participates in the CRM process prepares a risk response planning document during the *Plan* step and a risk response evaluation document during the *Track* and *Control* steps. The risk response planning document identifies the risk drivers that the entity has decided to act upon, the risk response options that the entity has considered, and the rationale for promoting one versus another. This enables multiple organizational units to work cooperatively to address the risk drivers that are cross-cutting, the impact that the risk drivers have on performance risk, and the capacity of proposed risk response options to reduce the risk across multiple units. The risk response evaluation document provides traceable evidence of how the risks are being addressed in real time and the degree to which the controls are or are not succeeding, from the perspective of the participating organizational entity.

The risk database can be used as the central repository of risk management documentation related to CRM. In order to fully support the process, the risk database must be relational, allowing for many-to-many linkages between individual risk scenarios, performance risks, risk drivers, and risk responses. It provides storage and archiving of the risk analysis results as they evolve over the course of the activity. It also provides storage and archiving of risk response planning, including the set of risk response alternatives, the risk analyses of the alternatives, the selected risk response, and the rationale for the selection.

5.8 References for Chapter 5

1. NASA Special Publication, NASA/SP-2016-6105 Rev2, NASA Systems Engineering Handbook. Issued February 2017, Updated January 2020.
2. NASA Procedural Requirements, NPR 8000.4C, Agency Risk Management Procedural Requirements. April 2022.
3. Government Accountability Office Report, GAO-23-106021, NASA: Assessment of Major Projects. May 2023.
4. SpaceNews Article, Cost and Schedule Overruns Continue to Grow for NASA Programs. May 2021.
5. NASA Cost Estimating Handbook, Version 4.0. Issued February 2015, Updated December 2020.
6. NASA-HDBK-7009A, NASA Handbook for Models and Simulations: An Implementation Guide for NASA-STD-7009A. May 2019.
7. NASA Special Publication, NASA/SP-2011-3421 Probabilistic Risk Assessment Procedures Guide for NASA Managers and Practitioners, Second Edition. December 2011.

6 Organizational and Managerial Aspects of Objectives-Driven Risk Management

The application of an integrated perspective in the identification and management of risk has been discussed in Chapter 2 as being one of the key principles and pillars of the NASA risk management framework. The present chapter identifies the basic organizational and managerial elements that must be considered, and are to be put in place, as a necessary pre-requisite and supporting structure for the practical implementation of an integrated risk management perspective.

In general, it is necessary to assume an organizational structure when discussing the details of a cross-organizational execution of risk management that includes both mission and institutional objectives. This chapter assumes the existence of the current NASA organizational structure when considering cross-organizational risk-related interactions, i.e., it recognizes an organization that consists of Agency executive-level management and its supporting councils that operate for the most part in an Enterprise Domain, a set of mission directorates responsible for coordinating mission-oriented and technical activities in the Program/Project Domain, a mission support directorate responsible for promoting and maintaining institutional efficiency through coordination and pooling of in-house resources and out-of-house acquisitions, and a number of centers and facilities responsible for supporting the programs and projects and maintaining core capabilities. It should be recognized, however, that the principles of the integrated, cross-organizational approach apply irrespective of the specifics of the organizational structure and would continue to apply if those specifics were to change in the future.

6.1 Principles of Organizationally Integrated Risk Management

The integrated implementation of risk management across NASA activity domains and organizational interfaces is a subject that has already been partially discussed in Chapter 2. In that context it was noted that the Agency organizational structure is designed to execute a broad spectrum of activities in pursuit of objectives that span the three principal interconnected domains listed below:

- Enterprise Domain, where strategic decisions are made and corresponding execution priorities and resources are established;
- Program/Project Domain, where the execution of activities and projects for development of technology and missions is assigned and carried out, reflecting the translation of Agency strategic objectives into operational and practical ones;
- Institutional Domain, where the human, physical, and technical support structure for the other two domains is developed and maintained.

An integrated view of risk across the above activity domains and the organizations that are responsible for the associated activities requires that, notwithstanding the fact that day-to-day risk identification and evaluation processes are initiated and conducted within the boundaries of individual organizational units, due consideration must be given also to risks of a cross-cutting nature. These are risks that cut across organizational boundaries because of the interconnections that typically exist when top-level objectives are allocated for execution across activity domains and to multiple organizational units.

The Agency's higher level strategic objectives are the unifying focuses of all derived activities, including the risk management activities set up to address the attending risks. Therefore, the understanding and support of such higher-level objectives by all lower-level organizations that are assigned suballocated portions of the associated execution activities are necessary pre-requisites for an effective integration of risk management throughout the Agency.

In practical terms, an integrated approach to risk management needs to be attuned to some specific objectives, prioritizing its focus on the identification of activity and organizational interface areas that without sufficient attention to risk integration priorities may easily become barriers to an effective flow of risk communication, decision, and action. In the following sections, the discussion will address aspects of risk management integration intended to overcome such barriers, via the application of appropriate processes and protocols. To keep a practical focus the discussion considers the following risk integration subjects that should be of main concern for any activity or project:

- Risk management integration across life-cycle stages
- Recognition of cross-cutting risks
- Risk management across organizational boundaries

6.2 Risk Management Integration across Life-cycle Stages

This section addresses risk management in a life-cycle context, such as is required by [1] for space flight programs and projects. Section 6.2.1 addresses the application of RIDM across different life-cycle stages. Section 6.2.2 addresses the evaluation of the risk management effort at LCRs and the development of assurance that the established risk posture is being adhered to.

6.2.1 Application of RIDM and CRM Across Life-cycle Stages

A first important aspect of necessary risk management integration arises when an activity or project progresses through the stages of its life cycle and its nature evolves accordingly. This aspect involves the application of appropriate risk management processes and tools at each life-cycle stage and the transfer of risk information and actions resulting from the application of these processes, across the decision gates at the times of transition from an activity stage to the next.

The two principal processes that complement each other within the NASA risk management framework are RIDM and CRM. While they have been discussed at length in all aspects of their implementation steps, their essential characteristics can be summarized here as follows:

- RIDM is a process that considers the risk profiles of alternative technological and/or activity execution solutions – including the selection of alternative acquisition strategies and types of Acquirer-Provider formal agreements and relationships – or the achievement of declared activity objectives. RIDM then applies an AoA (Analysis of Alternatives) approach to identify the solution that can be selected by decision makers as being optimal from the perspective of minimizing risk and maximizing benefits and opportunities.
- CRM is a “within-activity” process that, once a path and plan of activity and/or project execution has been decided and set in motion, systematically identifies risks that may

impact the achievement of the activity objectives and applies suitable measures and controls to prevent or mitigate such risks.

An integrated application of risk management across activity life-cycle stages involves the utilization of these two processes in coordinated fashion and consistently with their intrinsic characteristics, as predicated by NPR 8000.4 in specific paragraphs that address this subject, [2, Sections 1.2.2 to 1.2.4]. To achieve such a coordination, it is important to keep in mind the above distinct characteristics of the two complementary processes. Notably, the complementarity resides in the fact that RIDM can be applied at any level of decision-making support; e.g., RIDM can be used at different levels, to decide what type of project is best suited for the realization of a high-level strategic objective, which type of acquisition solution (e.g., internal development, co-development by Acquirer and Provider in specified roles, purchase of system or service) is best suited for that projects, or which actual space missions to select for execution, out of a theoretically possible set, or to select the type of design solution best suited for a given mission execution.

Thus, it is worthwhile noting that Activity-Planning RIDM is a process to be typically applied in the stage of setting up a project or activity and of deciding on its course. Since its application involves the identification and characterization of the principal risks that may potentially affect the pursuit of the activity objectives and the type of mission that is eventually selected for implementation, risk management integration requires that this information should be transferred directly into the CRM process that is set in motion as soon as the implementation and execution of the activity or project starts. This represents a first and key necessary element of continuity and complementarity between RIDM and CRM that needs to be deliberately pursued. That is, since pre-execution RIDM may typically be carried out before the organizational structure of a project or activity is actually set up, the risk information that it produces at this stage and that should be treated as the starting point of the execution-stage CRM-based risk management processes should be formally documented and transmitted across the pre-execution to execution gate of the activity life cycle.

A second and no-less important element of RIDM-CRM integration is directly related to what in Part 1 Chapter 2 has been referred to as “Activity-Execution RIDM.” This type of RIDM application is invoked from within a CRM process, whenever the control or mitigation of an identified risk may require the identification of an “optimal” risk control solution among a set of theoretically possible ones. Although in theory this type of risk-control AoA could be invoked in many situations, in practical terms it applies primarily to cases where an optimal selection of risk control or mitigation solutions addressing major risks cannot be easily identified without an in-depth consideration of the resulting benefits against the cost and resources needed for their implementations, or even the potential for other risks that their application could induce. These cases represent the second significant class of situations where a close integration of RIDM and CRM processes is warranted and necessary.

Figure 6-1 illustrates the aspects of RIDM-CRM processes integration that have been discussed above.

Life Cycle Progression




Figure 6-1. Life-cycle Integration of RIDM and CRM Processes

The figure articulates the above concepts one step further, in that it considers the application of RIDM and CRM in three conceptual stages of an activity life-cycle progression, which are labelled as “Activity & Activity Selection,” “Activity Execution Planning,” and “Activity Execution.” The first is the stage where an executive decision is made to identify a specific type of project or activity and the associated acquisition strategy and means that are deemed to be most appropriate and effective to realize an Agency high-level objective or set of objectives. The second is the stage where a project or activity has been identified and defined in terms of its principal execution objectives, but the practical course and timing of its execution is being elaborated and formally defined. The third is the activity true execution stage according to well defined plans and technical solutions. This characterization of an activity or project life cycle is conceptually general and therefore serves well the purpose of the present discussion. It is in no way intended to be an alternative to the formal systems engineering life-cycle phase definitions that are common use within programs and projects, and the respective stages and phases can be easily cross mapped.

The logic flow of RIDM-CRM execution and results integration illustrated by the figure shows that a first type of RIDM execution, in the Acquisition/Activity Selection stage, results in the selection and definition of an activity or project and associated means of acquisition, informed by the identification of an associated top-level risk profile, which is deemed more favorable than the risk profiles of the other activity or project alternatives considered in the RIDM AoA. Such a risk profile constitutes the key risk management information that needs to be transmitted across the time and logistic boundary between the Acquisition/Activity Selection and the Activity Execution Planning stages. A practical challenge for this aspect of risk management integration is that the management and technical staff in charge of an activity execution may in many situations differ from the personnel staffing the teams and councils that participate in the activity and acquisition-strategy selection and associated RIDM processes. Therefore, it is important that effective channels of communication and documentation of the decisions and supporting analyses elaborated in the latter be established so that the necessary continuity of information and risk perspective can be assured across those activity stages.

6.2.2 Evaluation of the Risk Management Effort at Life-cycle Reviews

The partitioning of programs and projects into life-cycle phases, each with one or more LCRs, was discussed in Section 2.2.6.2. The present section addresses the task of evaluating the risk management effort at LCRs by providing a general structure for the evaluation and the making of findings. It is assumed that:

- Success criteria have been defined for each LCR that collectively provide a sound and approved basis for determining whether or not the program/project is adhering to the established risk posture, given the state of the program/project at the time of the review. Valid, approved sets of LCR success criteria effectively reduce the question of adherence to the risk posture to one of meeting the success criteria.
- The evidence that will be used to substantiate meeting the success criteria has been specified and approved prior to the execution of the phase during which it is produced. Approval of the evidence specifications are based on a determination that they provide a valid basis for determining whether or not the LCR success criteria have been met.

LCRs are conducted by the program or project, with participation by an independent SRB at selected reviews [1]. Consistent with the above bullets, the evaluation of the program/project at each LCR, and the making of findings and recommendations, can be structured as follows:

- Is the evidence produced consistent with the approved evidence specifications? Ideally, the evaluators (e.g., the SRB) should know in advance what kind of evidence to expect, how to interpret it, and how to connect it to claims that the LCR success criteria are met. When unexpected forms of evidence are presented at an LCR, uncertainty is created in both the meaning of the evidence and its connection to the LCR success criteria, which erodes assurance that the risk posture is being adhered to. It also raises the question of whether the presented evidence has been selected after the fact in a manner that biases the assessment of the status of the program/project.
- Does the evidence indicate that the LCR success criteria have been met? Even if the evidence is consistent with the approved specifications, the question remains as to whether or not it substantiates the meeting of the LCR success criteria. Analytical results, test results, audit findings, budget data, schedule projections, and other forms of evidence might indicate that one or more success criteria are not met, which may indicate that the risk posture is not being adhered to. Conversely, evidence that is consistent with what was approved, and which clearly indicates that the LCR success criteria are met, provide a sound basis for concluding that the program/project is adhering to its risk posture.
- Are the program's/project's claims about whether or not the LCR success criteria are met consistent with the evaluator's assessment of the evidence? Independent evaluators should be skeptical towards arguments that rely on unexpected evidence, that interpret the evidence in unexpected ways, or that connect the evidence to the LCR success criteria in ways that differ from the basis for approving the evidence specifications initially.²⁰
- To what extent to deficiencies identified during the LCR affect the assessment of whether or not the program/project is adhering to the established risk posture? The answer to this question depends on the original argument made early in the program/project life cycle that establishes the validity of the LCR success criteria. Different criteria may be more or less determinative of adherence to the risk posture, and ability to remediate a deficiency depends on its nature and magnitude.

In practice, deficiencies are not unexpected, so one question for evaluators becomes whether the deficiencies are large enough to recommend corrective actions, either prior to or in parallel with proceeding to the next life-cycle phase.

6.3 Recognition and Handling of Cross-Cutting Risk

An integrated perspective on risk and on the approach to risk management requires a timely recognition of the cross-cutting nature of certain specific individual risk scenarios. Such a recognition constitutes the trigger for the management provisions and organizational protocols that

²⁰ Evaluators should be on the lookout for “special pleading,” e.g., claims of mitigating circumstances whose ability to negate adverse evidence is asked to be taken on faith.

should also be planned and implemented to address risk that for its effects or for the actions necessary for its handling affects multiple areas of the organization.

In practical terms, the determination of whether a given risk should be considered cross cutting can be based on the answer to the following questions:

1. Does the risk originate from within activities that are the execution responsibility of an organizational unit that is managerially distinct from the one that has identified it and deems to be affected by it?
2. Does the risk have the potential to significantly impact the objectives of a unit which is managerially distinct from the identifying organizational unit – i.e., it operates at a higher or lower level, or in a parallel but separate area of the organization?

If conditions 1 or 2 are met, the risk should be considered cross-cutting. An example of the first condition is a situation where a project deems to be vulnerable to a potential cyberattack because of insufficient firewall protections for which a supporting information technology (IT) organization is responsible. An example of the second is a situation where a program that includes several projects sees the possibility of budget shortfalls by which it may be forced to reduce the resources assigned to each and all of the included projects. Another example is a situation where significant delays in a technology development and validation program may affect the execution of a set of distinct projects in which the technology was to be implemented.

The actual recognition of the cross-cutting nature of a risk may occur at any level according to the general criteria identified above. However, in some cases it may be easier for higher level organizations to recognize that some risks affect several of the organizations that are operating at lower levels. It is important to maintain awareness that this type of recognition, however, is critically dependent on a free flow of information and communication about risk from the lower levels of the organization upwards, and vice versa. The need for improved risk communication was one of the key findings of the NASA RMTT report previously cited [3]. Besides the above, perhaps the most important distinction regarding the nature of a cross-cutting risk concerns whether such a risk can or cannot be dealt with, and controlled from, within the boundaries of the identifying organizational unit.

The possible conditions under which a risk scenario may be classified as being “cross-cutting” and the resulting options for its handling are depicted in Figure 6-2.

The figure underscores that that the determination of whether a risk scenario should be classified as cross-cutting may be arrived at in either of two ways. The first of these may occur when, in a review of risk scenarios reported to a higher-level organizational unit by its lower-level sub-units, commonalities in underlying causes or resulting effects of reported risks are identified, or a risk / risk scenario comes to attention as originating and/or impacting several of the reporting units, or even as having the potential for affecting other organizational units of whose activities the reviewing higher-level unit is aware. The second case occurs when the lower-level unit that has originally identified a risk scenario recognizes its potential cross-cutting nature and flags it to its parent higher-level unit, so that the latter can make a final determination. As the figure suggests,




Figure 6-2. Recognition and Classification of Cross-Cutting Risks

in either case it is the higher-level organizational unit's responsibility to decide whether risks being evaluated indeed have cross-cutting connotations.

Once a risk has been officially determined to be cross-cutting by a responsible higher-level organization, different degrees of risk management integration may be required to deal with it, depending on additional more specific associated characteristics. As the figure indicates, the cross-cutting risk scenario may, from the handling perspective, be classified into one of two categories, i.e.:

Type A Cross-Cutting Risks: In Case A, the risk management activity cannot be effectively executed by a single lower-level organizational unit and its integration extends into the coordination of the risk handling itself (i.e., risk control planning, tracking, etc.). The risk therefore is cross-cutting not only in the impacts that it may produce but also because its causes reside in the areas of responsibility of several organizational units and due to the coordination of actions that therefore its handling requires. Depending on the specific situation, this may involve the execution of a relevant portion of the risk handling process by the higher-level coordinating entity, and/or the allocation of risk control and tracking activities to the multiple organizational entities that are best suited to execute them. As an example, consider a risk constituted by the difficulty in recruiting qualified personnel in some specific technical discipline, in multiple NASA programs or centers that support several projects. The risk is clearly cross-cutting and cannot be addressed by one single organizational unit. All of the affected centers have to address it, possibly under the overall coordination of a headquarters office that provides special hiring incentives for recruiting and/or a balanced distribution of the needed personnel across the centers.

Type B Cross-Cutting Risks: In Case B, although its effects may impact several units and activities, the conditions that originate the cross-cutting risk of concern are confined in the area of responsibility of a specific lower-level unit and it is determined that such a unit has the capabilities and resources to address it. As an example, consider a situation by which the development / procurement of a new launch vehicle slated to serve several different spacecraft projects is affected by technical difficulties and potential serious delays. Although the risk is cross-cutting because of its possible multi-program impact, its resolution is likely to be possible if the organization responsible for the new launch vehicle development can successfully address the technical problems that are at its roots. In such a case, the handling of the risk will typically remain to be the responsibility of the lower-level entity, which may be either the unit which has initially identified it, or a unit that the higher-level reviewing entity judges to be the best suited for the handling tasks. In this case the integration aspects of the cross-cutting risk management can be limited to a timely notification given to the other potentially affected entities of the existence of the risk, and of the measures being taken for its proper handling. This should also be followed by information updates to such entities about the outcome of the undertaken risk handling actions, so that they may take remedial actions of their own in case the original handling plan of the responsible unit cannot be executed successfully in the time frames that are acceptable for the other affected units. Thus, in the example situation described above, if

the entity responsible for the development / procurement of the new launch vehicle cannot resolve in a timely fashion the technical issues affecting the procurement schedule, the projects that were originally slated to use the vehicle should be given notice in a timely fashion of the risk handling difficulties that may be occurring, so that they may seek the launch services of alternative vehicle to avoid negative impacts on their projects.

From the above it is evident that a realistic view of the flow-down of objectives from the highest level into the various organizational areas of execution responsibility and across the activity domains (Enterprise, Program/Project, and Institutional) is necessary for the determination of any cross-cutting nature of identified individual risk scenarios and of the associated risk management integration activities. Examples of “flow-down trees” of organizational objectives, and of their allocation to organizational units for practical execution, is provided in Part 2 Chapter 3. Alongside that type of practical instrument, another tool that may be helpful in the determination of the cross-cutting nature of certain individual risk scenarios is represented by the risk-identification taxonomies provided in Appendix I. A “risk taxonomy tree” facilitates the identification of credible risk scenarios and of the potential breadth of their direct impacts on a project, system, or mission. This enables a parallel evaluation, via the perspective provided by considering the “objectives tree” of the affected organizational entity (or entities) and the corresponding allocation of execution responsibilities, of the possible propagation of these impacts, and it makes more rigorous and reliable the determination of whether a risk of concern has or does not have cross-cutting characteristics.

6.4 Integration Across Domain and Organization Boundaries

The preceding section has addressed the question of what characteristics make a risk recognizable as being “cross-cutting.” Such a recognition is part of the overall subject of integration of risk management across organizational boundaries discussed in the present section, but was presented in its own section for emphasis, and because as mentioned there it constitutes the “trigger” for the application of integration provisions and protocols.

The difficulty of implementing an effectively integrated risk perspective has often originated in the past from a lack of consistent risk management planning in different areas and across the activity domains of the organization. Historically, risk management has been set-up and executed as a formally planned activity predominantly in the Program/Project Domain and within specific programs or projects. A corresponding practice, however, has not been firmly established in the Enterprise or Institutional domains, with the effect that risk management responsibilities have not been assigned and corresponding activities have not been planned in consistent fashion across the organizational structure. Under these conditions it has been difficult to establish and implement protocols for timely management of cross-cutting risks that exist because the activities and actions carried out by distinct operational units contribute at some level to important common objectives, even though such units may for the most part operate in separate domains.

Advocating for the creation of specific risk management roles or responsibilities for the execution of risk management activities is outside the scope of this handbook. However, the recognition and handling of risk is an across-the-board management responsibility, which remains true for cross-cutting risks as well. Because, by definition, cross-cutting risks have potential impacts across

activity domains and organizational boundaries, their timely identification and effective handling requires a **coordination of management actions across those boundaries**, regardless of whether such actions are classified as being part of a formal risk management function or not in any particular organizational units.

Consistently with the basic distinction between the two types of cross-cutting risk defined in the preceding section, the integration of risk management across organizational boundaries should concern the formulation and implementation of two basic types of management interaction protocols, i.e.:

- a. Cross-Organizational Risk Communication Protocols (CORCP)
- b. Cross-Organizational Risk Handling Protocols (CORHP)

CORCPs are protocols that should be established for definition of timely and appropriate lines of communication, for exchange of information relative to cross-cutting risks of any type affecting distinct organizational units. Such protocols should establish, as a minimum standard to be implemented across the Agency:

- A definition of the conditions for triggering the initiation of communications concerning cross-cutting risk among the affected organizational units, as well as for the updating of such communications at regular intervals or upon request, and for their eventual termination.
- The timeliness requirements relative to the above communications.
- For each type of organizational unit, the standard and “default” point-of-contact (POC) individuals or offices designated to originate or receive the risk communications.

CORHPs are protocols that should be defined, in addition to CORCPs, to assure that an adequate level of coordination among distinct organizational units is in place for the situations where the handling and control of a cross-cutting risk requires a combination of decisions and actions by those units. A CORHP should as a minimum establish:

- Criteria to determine, for the various possible situations of concern, which organizational unit should be the risk handling coordinator, i.e., should be the entity that would have the responsibility of:
 - a. Identifying the other organizational units called to actively contribute to the cross-cutting risk handling activities, and
 - b. Organizing and leading coordination meetings where the cross-cutting risk handling strategy and implementation timetable is discussed, decided upon, and set.
- Decision elevation criteria, for cases where a higher authority deliberation and decision is required to define, set, and implement a coordinated risk handling strategy.

Both types of protocols should to the extent possible be defined in the Risk Management Plans (RMPs) formulated by organizational units at all levels. A preliminary coordination should therefore be established to make sure that such protocols are defined in cross-compatible fashion in different RMPs, i.e., with no contradictory criteria defined in one RMP vs another.

Examples of integration of risk management activities for the handling of cross-cutting risks are provided in Part 2 Chapters 2 and 3.

6.5 Risk Management Execution Planning

This section discusses guidelines for preparing a Risk Management Plan (RMP) that (1) addresses the RM organizational processes and interactions that are needed to support the principles of risk leadership, as defined in NASA NPD 1000.0 [4], and expanded upon in Chapter 2 of this handbook; (2) conforms with the requirements for NASA RMPs presented in NASA NPR 8000.4 [2]; and (3) ensures that risk management is aligned with the success criteria defined in NASA NPR 7123.1 [5] for each Key Decision Point (KDP).

In past standard practice NASA RMPs have primarily focused on the following topics:

- Purpose, scope, and relevant documents
- Roles and responsibilities
- Avenues of communication
- Definition of likelihood and consequence categories
- Brief discussion of the traditional CRM steps (identification of risks, analysis, planning, tracking, control, communication, and documentation), as applicable to the organization issuing the RMP
- Preparation of risk registers and risk matrices of likelihood vs. consequence, oriented toward evaluation of individual risk scenarios

The current NPR 8000.4 requires that NASA RMPs be modified and expanded in scope to cover the following areas:

- Risk leadership principles; risk posture; risk tolerances; and risk acceptability criteria
- Identification of stakeholders
- Risk types
- Sources of risk
- RIDM and CRM approaches implementation and integration
- Applicable organizational objectives and requirements
- Identification of risks to meeting each key objective and requirement
- Applicable commitments for providing evidence that objectives and requirements will be met (e.g., testing)
- Level of quantification vs. qualitative treatment
- Risk aggregation from individual risk scenarios to the risk of not meeting organizational objectives and requirements

- Coordination with higher level RM plans and the Systems Engineering Management Plan (SEMP), as applicable
- Definition of categories for likelihood and consequence severity, consistent with the definition of performance objectives – e.g., as expressed by performance markers – and corresponding risk tolerance levels
- Protocols for assessment of risk levels via estimation of PM outcomes likelihood and/or probability distributions
- Special provisions for cybersecurity and physical security risks
- Treatment of uncertainty, including consideration of the potential magnitude of unknown and/or underappreciated (U/U) risk scenarios
- Risk elevation protocols
- Cross-organizational coordination protocols
- Risk communication protocols and display formats
- Documentation of management decisions
- Intervals for periodic reviews
- Management concurrence and signature

The guidance presented in this chapter covers the subject areas listed in the two sets of bullets provided above, as well as some areas implied by the above lists though not directly stated (e.g., the tie-in of RM to the success criteria specified in NPR 7123.1 [5] for each KDP, and the development of a case for risk acceptance, when this may be required at KDPs).

6.5.1 Contents of the RMP

The list of contents presented in this section is intended to be comprehensive. It is incumbent upon the responsible organizational unit to select and tailor the contents of each RMP for the specific context, scope and purpose of the risk management activity for which it is generated.

Introductory Information

- Identify the purpose and scope of the RMP.
- Identify the source documents, including NASA NPDs and NPRs, that are relevant to the RMP, and explain their relevance.
- Identify and describe the activity that is the subject of the RMP (“the subject activity”).
- Specify the key decision points (KDPs) for the subject activity, consistent with the identification of KDPs in the NASA NPR 7123.1D if applicable.
- Identify other activities that are relevant to the subject activity, and explain their relevance.

Identification and Characterization of Performance Objectives, Performance Markers and Associated Risk Tolerances

- Identify the top-level objectives for which performance measures are defined and on which the activities of the organizational unit are focused.
- Identify and describe the quantitative performance measures that will be used to evaluate the aggregate risk to each top-level objective, and explain why they are considered to be appropriate measures of performance relative to them.
- Identify the relevant performance markers to be used (e.g., PMK-R and PMK-G values) and indicate how these selections comport with the top-level objectives.
- Identify the risk tolerance levels for each performance marker (e.g., RTL-Rs and RTL-Gs) and indicate how these selections comport with the *Acquirer*/stakeholder risk posture.

Identification and Description of Roles and Responsibilities

- Identify which persons are *responsible* (i.e., the *risk manager*, who is in charge of organizing the risk management tasks for the subject activity, delegating responsibilities for their accomplishment, overseeing the implementation, reporting to the activity manager, and providing periodic updates to the *Acquirer*; and the *individual risk scenario owners*, who are responsible for ensuring that the individual risk scenarios under their purview are properly analyzed and responded to, communicating continuously with the analysis team including staff members and contractors, and reporting to the risk manager).
- Identify which persons are *accountable* (e.g., the designated administrator or member of management who has sign-off or approval authority when each major subtask in the risk response is complete, or when the risk management plan has to be updated to reflect decisions made at key decision points).
- Identify which persons are to be *consulted* (e.g., risk boards, technical reviewers, risk managers of other activities that cross-cut with the subject activity, and others that are relied upon to give input before each major subtask is completed and signed-off on).
- Identify which persons are to be *informed* (any other person or group of people who have a stake in the outcomes, including people from other program or project organizations who are affected by the cross-cutting nature of the risks being responded to).
- Explain the means by which these designated people have been empaneled and have formally accepted the commitments they are expected to honor.

Processes for Identifying Risks, Opportunities, and Leading Indicators

- Identify the categories of risk that will be included in the RM analyses and deliberations (e.g., spaceflight safety risks, spaceflight technical risks, physical security risks, cybersecurity risks, cost risks, schedule risks, staffing risks, training risks, maintenance risks, supply-chain risks, facility safety risks, facility availability risks, facility technical risks, organizational strategic risks, operations risks, compliance risks, acquisition risks, fraud risks, and reputational risks).

- Specify the source information that will be used to identify risk scenarios, potential opportunity scenarios, and leading indicators of unknown and/or underappreciated (U/U) risk.
- Describe how risks, opportunities, and leading indicators will be classified according to characteristic attributes such as departure events, development phase, affected assets, and affected objectives, using taxonomies or other classification schemes, and specify how these classifications will be used to help ensure that the identification of risk scenarios is comprehensive and as complete as reasonably achievable.
- Describe how risk and opportunity statements will be structured for individual risk scenarios, opportunity scenarios, and leading indicators of U/U risk, and specify what will be contained in each, including the activity objective(s) and/or non-objective-based mandated requirement(s) that are affected.
- Describe how the risk and opportunity statements will be validated, using the criteria in Section 5.2.3: Validate the Risk and Opportunity Statements
- Describe what information will be provided in the accompanying narratives for each risk scenario, opportunity scenario, and risk leading indicator.

Processes for Analyzing Individual Risk Scenarios and Their Flow-Up to Aggregated Risks

- Describe how the level of analysis rigor to be employed for each performance measure will be graded on a scale ranging from low to high based on factors such as Activity Class, the complexity and novelty of the activity, and the stage of activity implementation.
- Describe how the analysis models and data sources that will be used to analyze individual risk scenarios and their flow-up to aggregate risks will be identified and selected for each performance measure in a way that is consistent with the graded analysis approach.
- Describe how estimates will be made for the likelihoods and consequences of the individual risk scenarios using heuristic (experience- and judgment-based) methods; then describe how combinatorial logic and expert judgment will be used to estimate the aggregated risks to performance measures.
- Identify the analysis models that will be used for probabilistic analysis of individual and aggregate risks, and describe how the selection of models from among this set will be made based on the refined graded analysis approach.
- Describe how the models to be used will be verified and validated prior to use based on procedures for verification and validation provided in the NASA Handbook for Models and Simulations, NASA-HDBK-7009A.
- Describe how margins to account for potential unknown and/or underappreciated (U/U) risks will be determined, what leading indicators will be examined to support this determination, what data sources will be used to evaluate the leading indicators, and how the values of the leading indicators will ultimately be related (e.g., through correlations) to the potential magnitude of the U/U risk.

- Describe how the results of the analyses will be used to determine and prioritize the most important risk drivers.
- Describe and illustrate how the risk analysis results will be organized and displayed using, for example, risk spider charts or other informative display techniques.

Processes for Planning and Implementing Responses to Unacceptable Risks, Including the Use of RIDM during Activity Execution

- Describe the process for determining how to decide whether an individual or aggregate risk should be retired, watched, researched, mitigated, or elevated to a higher level.
- For risks that need to be mitigated, describe the process for deciding whether the mitigation can be planned and executed without requiring additional RIDM analysis, or whether the RIDM process needs to be initiated in the Activity-Execution mode.
- For mitigation decisions that need additional RIDM analysis (i.e., a formal Analysis of Alternatives), explain the process for petitioning management to approve the re-opening of RIDM in the Activity-Execution mode, the supporting documentation that would be provided by the activity manager to support that petition, and the sign-off process by which the management authority would approve or reject the petition.
- Describe the process for determining whether the *Acquirer's* requirements and/or risk tolerances need to be modified or rebaselined in order for the mitigation action to succeed, and if that is the case, explain the process for petitioning management to approve the re-opening of RIDM in the Activity-Rebaselining mode, and the supporting documentation and sign-off process.
- Describe the process to be followed if it is decided to implement a RIDM Analysis of Alternatives in the Activity-Execution mode, including how the mitigation alternatives will be identified, how appropriate data sources and analytical models will be identified, how unfruitful alternatives will be weeded out using an initial screening approach, how more promising alternatives will be analyzed using a more in-depth graded-analysis process, how the results will be displayed, and how deliberations will occur to provide recommendations to management concerning the selection of an alternative for implementation
- For risks that need to be elevated, describe the process for conducting communications between the organizational entity conducting the subject activity and the higher-level organizational entity to which the risk has been elevated, so as to ensure that the analysis of residual risks after planning and implementation by the higher-level entity has been completed is appropriately reflected in analyses performed at the subject activity level.

Processes for Tracking and Controlling Individual and Aggregate Risks

- Describe the processes for identifying, in a timely manner, any changes that have occurred in internal and external conditions that lead to changes in the conditions cited in the risk and opportunity statements.

- Similarly, describe the processes for determining, in a timely manner, whether any departure events cited in the risk and opportunity statements have actually occurred, causing them to become conditions rather than departures.
- Describe the processes by which the observed changes in conditions will be factored into the analysis of individual risk scenarios and aggregate risks, and the process for determining whether any resulting adverse effects on risk can be accommodated through straightforward and easy-to-implement controls, as opposed to more complex controls that require a detailed identification and analysis of new mitigation alternatives.
- Where it is determined that straightforward and easy-to-implement controls are sufficient, describe the processes through which these controls will be implemented and whether or not the RMP needs to be amended.
- Describe the processes for ensuring that risk responses are being implemented as planned, and if they are not, describe the processes for taking corrective actions.

Protocols for Communicating Cross-Cutting Risks, Ensuring that Related Analyses are Conducted Consistently across Entities, and Ensuring that Risk Responses are Decided-upon Interactively

- Describe the communication protocols that will be adhered to across the subject activity, between the subject activity and other activities for which there are cross-cutting interests, and with organizational management entities that possess decision making authority that affects the subject activity, including the subject content and frequency of inter-organizational meetings.
- Describe the protocols that will be adhered to in documenting areas of agreement and disagreement (particularly with regard to the handling of cross-cutting risks and the selection of risk mitigation options), plans for resolving disagreements, decisions on how to proceed thereafter, the rationale behind these decisions, and plans for elevating decisions to a higher organizational level when necessary.
- Describe the processes that will be followed to ensure that all cross-cutting risks, risk drivers, and risk response alternatives being considered are shared between entities, and that individual risk scenarios are integrated into the analyses of aggregate risks in a consistent way.

Criteria and Documentation of Risk Acceptance at Key Decision Points

- Identify the success criteria defined for each KDP and explain why satisfaction of the success criteria implies that the activity is adhering to the established risk posture or is on track to adhering to the established risk posture.
- Identify the evidence that will be produced to evaluate the activity against the success criteria at each KDP.
- Identify and establish appropriate lines of communication between the entity carrying out RM processes and the agency entities responsible for developing and implementing internal controls, so that the latter be informed by the former of the magnitudes of the risks

and the risk drivers that may need to be controlled, and in turn those entities be able to inform the RM executing entity about any tailoring of the current internal controls they plan to implement to address the identified risk drivers.

Procedures for Reviewing the Risk Management Plan and Updating It as the Activity Moves through Each Key Decision Point

- Describe the procedures for reviewing, updating, and approving changes to the RMP to reflect decisions reached about changing roles and responsibilities, implementing new risk mitigation alternatives, pursuing new opportunities, adjusting risk tolerances, adding new requirements, tailoring or waiving of existing requirements, and/or elevating risks to higher levels.

6.5.2 Cross-Organizational Integration of RM Plans

Risk management plans (RMPs) are to be prepared by any organizational unit assigned the execution of a set of activities or tasks in pursuit of identified and specified objectives – e.g., a program or project, but also a support organization executing a set of closely interrelated institutional tasks which are thus identifiable as a well-defined “institutional project.” The production of an RMP for a given unit is usually the responsibility of an “RMP authority” – usually a member of the unit’s management team although the actual task of producing the plan may be delegated to a technical expert or group of experts – but its approval and implementation should remain the ultimate responsibility of the program, project, or institutional activity overall manager.

Earlier sections of this chapter have introduced and discussed the subject of integration and coordination of risk management activities across organizational entities. To the extent possible and foreseeable, such a coordination should be planned and reflected in a corresponding coordination of the respective RMPs. Figure 6-3 shows a schematic representing typically desirable interfaces between RMP authorities for the coordination of their plans.

The figure underscores that in the course of an RMP production, the RMP authorities for entities that have interfacing or cross-cutting objectives, and therefore potentially interfacing risks, should coordinate with one another to assure that the respective RMPs identify the corresponding inter-organizational interfaces and define agreed-upon protocols to address them in a coordinated fashion. The coordination of RMPs between / among units that share potential risk interfaces is needed to ensure that the assumptions made in the definition of the plans are consistent with one another, and that roles, responsibilities, and cross-organization consultation activities are organized and assigned in a way that produces necessary coverage and maximum efficiency.

Consistently with NPR 8000.4, Section 3.2.2 Paragraph (i), the topics listed earlier in Section 6.5.1 should generally be included within each RMP. Their coverage and possible tailoring should in general be subjected to an inter-organizational review for coordination purposes, as described above.




Figure 6-3. Schematic of desirable coordination of RM plans

6.6 References for Chapter 6

1. NASA Procedural Requirements, NPR 7120.5F, NASA Space Flight Program and Project Management Requirements w/Change 3. August 2021.
2. NASA Procedural Requirements, NPR 8000.4C, Agency Risk Management Procedural Requirements. April 2022.
3. NASA Internal Report, Risk Management Tiger Team Report. September 07, 2023.
4. NASA Policy Directive, NPD 1000.0C, NASA Governance and Strategic Management Handbook. January 2020.
5. NASA Procedural Requirements, NPR 7123.1D, NASA Systems Engineering Processes and Requirements w/Change 1. July 2023.

7 Definitions

The reader should note that in this chapter of definitions, the numbered notes appear as endnotes in the end of the chapter instead of as footnotes.

Acquirer	A NASA organization that tasks another organization (either within NASA or external to NASA) to deliver a product (e.g., a system) or a service.
Aggregate Risk	The cumulative risk associated with a given goal, objective, or performance measure, accounting for all significant risk contributors thereof ¹ .
As Safe as Reasonably Practicable (ASARP)	Employing the safest means of achieving specified technical objectives within programmatic constraints (e.g., on cost and schedule) ² .
Continuous Risk Management	A systematic and iterative process that efficiently identifies, analyzes, plans, tracks, controls, communicates, and documents risks associated with implementation of designs, plans, and processes.
Cross-Cutting Risk	A risk that is generally applicable to multiple mission execution efforts, with attributes and impacts found in multiple levels of the organization or in multiple organizations within the same level.
Cybersecurity	Prevention of damage to, protection of, and restoration of computers, electronic communications systems, electronic communications services, wire communication, and electronic communication, including information contained therein, to ensure its availability, integrity, authentication, confidentiality and nonrepudiation.
Cybersecurity Risk	Cybersecurity risk can be thought of in two primary ways: Cybersecurity asset risk and cybersecurity compliance risk. Cybersecurity asset risk arises from threats to and vulnerabilities of cyber-assets - i.e., data, information, control systems, and information systems stored and/or implemented in digital form. Cybersecurity compliance risk arises from potential failures to comply with laws, regulations, or other requirements.
Decision Robustness	The character of a decision that is supported by sufficient technical evidence and characterization of uncertainties to determine that: a) the selected

decision alternative best reflects decision-maker preferences and values, consistently with the informing state of knowledge at the time of the decision, and b) can be deemed to be insensitive to credible modeling perturbations and realistically foreseeable new information.

Graded Approach

The application of risk management processes at a level of detail and rigor that adds value without unnecessary expenditure of resources. The resources and depth of analysis are commensurate with the stakes and the complexity of the decision situations being addressed³.

Imposed Constraint

A limit imposed by a higher decision authority on the allowable values of the performance measure with which it is associated⁴.

Individual Risk Scenario

A sequence of events or combination of such sequences, originated by an event or condition followed by other events or conditions, which are judged to be unique and defining with respect to how the consequences of concern are produced and impact one or more activity objectives.

Mission Objective

An explicitly established and stated desired outcome or product of a mission⁵.

Mission Success

A mission outcome in which all mission technical objectives have been met. Mission success can be whole, where all mission objectives are fully met, or partial, where some mission objectives are not met or are only partially met.

Mission Success Risk

The likelihoods that mission technical objectives will not be achieved.

Objectives-Driven Risk Management

An approach to risk management that explicitly focuses on ensuring that an activity's risk profile is within the established risk posture.

Opportunity

The possibility of an existing goal, objective, or desired outcome being met more efficaciously, or a new goal, objective, or desired outcome becoming feasible.

Organizational Unit	An organization, such as a program, project, Center, Mission Directorate, or Mission Support Office that is responsible for carrying out a particular activity.
Performance Measure	A metric used to measure the extent to which a system, process, or activity fulfills its intended objectives ⁶ .
Performance Parameter	A performance parameter is any quantifiable variable whose value is needed to execute the models that quantify the performance measures.
Performance Requirement	The value of a performance measure to be achieved by an organizational unit's service or product that has been agreed upon to satisfy the needs of the next higher organizational level ⁷ .
Physical Security	The objective and function of securing an organizational infrastructure, program, project, mission, or activity, and protecting their material, functional, control, information, and human assets from interferences of a physical and material nature that are capable of causing damage, or of resulting in unauthorized possession or control of such assets.
Program Objective	An explicitly established and stated desired outcome of a program. Program objectives typically fall into categories such as safety, technical, cost, and schedule.
Project Objective	An explicitly established and stated desired outcome of a project. Project objectives typically fall into categories such as technical, safety, cost, and schedule.
Provider	A NASA or contractor organization that is tasked by an accountable organization (i.e., the <i>Acquirer</i>) to produce a product (e.g., a system) or a service ⁸ .
Risk	The potential for shortfalls with respect to achieving explicitly established and stated objectives ⁹ .
Risk Attitude	The general inclination of a stakeholder or decision-maker to accept risk in pursuit of defined objectives.
Risk Driver	A significant source of risk to one or more organizational objectives or mandated requirements.

Risk-Informed Decision Making

A risk-informed decision-making process that uses a diverse set of performance measures (some of which are model-based risk metrics) along with other considerations within a deliberative process to inform decision making¹⁰.

Risk Leadership

One of the “NASA Senior Leadership Focus Areas” referred to in NPD 1000.0C, it is there described as the application by NASA of a risk culture that has the goal of “increasing ‘decision velocity’ within a proper risk posture.” It is implemented from the higher levels of management by communicating to the work force a clear and balanced understanding of risk and benefits, by defining and indicating appropriate technical standards, and by ensuring the workforce has the proper experience and commitment to collaboration [adapted from NASA NPD 1000.0C].

Risk Management

A coordinated flow of activities, included and closely integrated with all other management activities, to identify, evaluate, and address risk with appropriate actions, which combines RIDM and CRM in an integrated framework¹¹.

Risk Posture

A definition, expressed in qualitative or quantitative terms, of the level of acceptable risk to an activity’s top-level objectives¹².

Risk Profile

The ensemble of assessed risks to an activity’s top-level objectives¹³.

Risk Tolerance

The limit of acceptable likelihood of falling short of achieving an explicitly established and stated objective.

Risk Tolerance Level (RTL)

Risk Tolerance Level defines the level of risk declared to be acceptable with respect to the achievement of a performance requirement¹⁴.

Safety

In a risk-informed context, an overall condition that provides sufficient assurance that mishaps will not result from an activity, or, if they occur, that their consequences will be mitigated¹⁵.

Safety Risk

The likelihoods that the identified and declared mission safety objectives will not be achieved.

Technical Risk

The likelihoods that the identified and declared mission technical objectives will not be achieved.

Unknown and/or Underappreciated Risk

A risk that may elude an explicit operational characterization in the form of specific scenarios and events.

¹ For example, the total probability of loss of mission is an aggregate risk metric quantified as the probability of the union of all scenarios leading to loss of mission

² In practice, this entails prioritizing safety in decision-making throughout the program or project life cycle insofar as is practical. The ASARP objective may be separate and independent from any safety risk tolerances that may be levied on the mission to define thresholds of acceptable safety

³ For example, the level of rigor needed in risk analysis to demonstrate satisfaction of safety-related performance requirements depends on specific characteristics of the situation: how stringent the requirements are, how complex and diverse the hazards are, and how large the uncertainties are compared to operating margin, among other things. Both RIDM and CRM are formulated to allow for this flexibility in the depth and breadth of their analytical processes.

⁴ Imposed constraints are minimum performance requirements that are pre-defined and negotiated between NASA organizational units in order to define the task to be performed

⁵ Mission objectives include mission technical objectives, which relate to the purpose for which the mission is conducted (e.g., Collect 10 kg of lunar regolith and return it to Earth); mission safety objectives, which relate to the protection of relevant at-risk entities (e.g., Return crew safely to Earth, protect the public from reentry debris); as well as objectives in other mission execution domains such as cost and schedule. Mission objectives are defined at the mission level. Mission objectives are deterministic – they are either achieved or not achieved in any given instance of mission execution

⁶ Performance measures should in general relate to observable quantities. For example, engine performance parameters, cost metrics, and schedule are observable quantities. Although safety performance measures can be observed in principle, many of them have to be modeled. Partly because of this, in ranking decision alternatives, one may use a risk metric (e.g., probability of loss of crew) as a surrogate for a performance measure

⁷ In an Acquirer-Provider context, a performance requirement is the agreed upon level of performance to be achieved by a Provider's product that satisfies the needs of the Acquirer organization.

⁸ Synonymous to the term "Supplier" as used in NPD 1000.5

⁹ As applied to programs and projects, the objectives are translated into performance requirements, which may be related to mission execution domains (mission success, safety, physical and cybersecurity, cost, and schedule) or institutional support for mission execution. Risk may operationally be characterized as a set of triplets:

a. *The scenario(s)* leading to degraded performance with respect to one or more performance measures (e.g., scenarios leading to injury, fatality, destruction or compromise of key assets; scenarios leading to exceedance of mass limits; scenarios leading to cost overruns; scenarios leading to schedule slippage).

b. *The likelihood(s)* (qualitative or quantitative; unconditional or conditional) of those scenarios.

c. *The consequence(s)* (qualitative or quantitative severity of the performance degradation) that would result if those scenarios were to occur.

Uncertainties are included in the evaluation of likelihoods and identification of scenarios

¹⁰ A decision-making process relying primarily or exclusively on a narrow set of model-based risk metrics would be considered "risk-based"

¹¹ Risk management is done in order to foster proactive management of risk items, to inform better decision making through better use of risk information, and then to manage more effectively the implementation of risk-related activities and actions by focusing the CRM process on the baseline performance requirements and risk trades identified via the application of the RIDM process.

¹² An activity's risk posture expresses the agreed upon limits of risk an organization's leadership team is willing to accept in order to achieve one or more of its objectives. It is defined up front and in tandem with the development of objectives, consistently with risk leadership principles, and serves as the attitudinal framework for seeking a balance between the likelihood and benefit of achieving the objective(s), vs. the likelihood and severity of risks that may be introduced by the pursuit of achievement. Risk posture may change with time, in reflection of the evolution of leadership team attitudes or because of changes in priorities, but at any particular time, risk posture provides the de-facto basis for risk-informed decision making and continuous risk management.

¹³ Depending on the characterization of the activity's risk, its risk profile can consist of actuarially derived risk values, analytically constructed individual risk scenarios, and/or estimates of U/U risk.

¹⁴ An RTL is a quantitative definition of risk tolerance, usually expressed in terms of what probability of not meeting a performance target can be tolerated and accepted, and is the "other side of the coin" of stating the "confidence level" by which a performance requirement is to be met

¹⁵ NPR 8000.4C uses the term "safety" broadly to include human safety (public and workforce), environmental safety, and asset safety.

Appendix A Roadmap for Risk Management Handbook Utilization

Table 1-I in Chapter 1 presents a depiction of the common types of roles and objectives that an individual or team with risk management organizational or executional responsibility may have in the two principal top-level stages of a project or activity – i.e., *Definition/Planning* vs. *Execution*. This identification of risk management roles is then used in Table A-I to identify the topics covered in this handbook which may have different degrees of relevance and priority in relation to the risk management role and stage of application identified in Table 1-I.

A few clarifications are in order with respect to the above and a user's best utilization of the two tables (Table 1-I and the detailed roadmap in Appendix A). The first is that the roles defined in Table 1-I represent general definitions as represented in the table itself: they should not be construed as having any direct pre-defined correspondence with the personnel roles officially assigned within the NASA organizational and programmatic hierarchies. A second observation, also directed at a correct interpretation of the tables, is that in the context of the present discussion the terms "Definition/Planning" and "Execution" are used to identify the two major activity stages that are relevant in relation to the type of risk management processes that are to be executed within them: in the context of a formally structured project life cycle and the corresponding systems engineering definition of project phases, these two major stages would correspond, respectively, the former to a combination of the Pre-Phase A and Phase A portion of the project life cycle, and the latter to the remainder of all the following life-cycle phases.

In light of the above a user is invited to decide which among the activity-stage associated roles defined in Table 1-I they best identify with, then, based on that identification, to use the color-coded, stage and role dependent relevance classifications of handbook topics in Table A-I, they can decide on the order of priority by which the handbook topics may be consulted. With regard to this the color-coding and meaning of the Table A-I relevance classifications is as follows:

- Blue – High-Relevance Topic: a topic that provides a user with key background and/or technical information and skills necessary for an effective, NPR8000.4 compliant execution of his/her risk management responsibilities and functions in the activity stage of concern.
- Dark Green – Recommended Topic: a topic that provides a user with important managerial and/or technical skills relevant to and applicable in the execution of his/her risk management responsibilities and functions in the activity stage of concern.
- Lighter Green – Useful Topic: a topic that provides a user with information and-or technical skills relevant and applicable to the execution of his/her risk management responsibilities and functions in the activity stage of concern.
- Grey – Optional Topic: a topic that a user may decide to investigate as an optional background subject of interest.

Table A-I. Roadmap of Risk Management Handbook Utilization

		TOPIC	Risk leadership and risk posture	Risk management in the context of risk leadership	Risk management integration across organizations and application domains	RIDM and CRM processes integration	Graded analysis approaches for RIDM and CRM	Use of heuristic rules and criteria in RM processes
		PART 1 SECTIONS	2.1.4, 2.1.7, 2.2.2, 2.2.3, 2.2.4, 2.2.6, 3.3, 4.2.3, 4.4.3, 4.9, 6.2, 6.5	1.1, 1.4, 2.1.7, 2.2.2, 2.2.3, 2.2.4, 2.2.6, 3.3, App E	2.2.6, 4.2, 4.4, 4.7.2, Ch 6	2.2.6, 4.2, 4.4, 4.7.2, 5.4	1.4, 2.2.4, 2.2.5, 4.1.1, 4.2.2, 4.7.2, 4.9.4, 4.11, Ch 5 intro, 5.1 to 5.7, App C, App H, App I	3.3.4, 4.9.3, 5.3, 6.5.1, App D, App E
		PART 2 SECTIONS	3.2.1, 3.2.7		2.1, 2.3, 4.2	4.5	4.3, 4.4	3.2.4
PHASE	USER	FUNCTION						
All	Deliberation & Decision Board / Council	<i>Providing High Level Directions and Making Decisions</i>	High Relevance	High Relevance	Recommended	Useful	Useful	Useful
	Higher Level Leader / Manager	<i>Providing Top-Level RM Direction / Focus</i>	High Relevance	High Relevance	Recommended	Useful	Useful	Useful
Definition / Planning	Recommendation Board / Council	<i>Providing Oversight</i>	Recommended	High Relevance	High Relevance	High Relevance	Recommended	Recommended
	Program / Project / Activity Manager	<i>Organizing and Directing RM Tasks</i>	Recommended	High Relevance	High Relevance	High Relevance	Recommended	Recommended
		<i>Making and Accounting-for Risk Relevant Decisions</i>	Recommended	High Relevance	Recommended	High Relevance	Recommended	Recommended
	Program / Project / Activity RM Specialist	<i>Executing RM Processes and Tasks</i>	Recommended	Recommended	Recommended	High Relevance	High Relevance	High Relevance
	Program / Project / Activity Risk Analyst	<i>Executing Specific Risk Analysis Tasks</i>	Useful	Useful	Useful	Useful	High Relevance	High Relevance
Execution	Review & Recommendation Board / Council	<i>Reviewing and Providing Oversight</i>	Recommended	High Relevance	Recommended	High Relevance	Recommended	Recommended
	Program / Project / Activity Manager	<i>Organizing / Directing RM Tasks</i>	Recommended	High Relevance	Recommended	High Relevance	Recommended	Recommended
		<i>Making and Accounting-for Risk Relevant Decisions</i>	Recommended	High Relevance	Recommended	High Relevance	Recommended	Recommended
	Program / Project / Activity RM Specialist	<i>Executing RM Processes and Tasks</i>	Recommended	Recommended	Recommended	High Relevance	High Relevance	High Relevance
	Program / Project / Activity Risk Analyst	<i>Executing Specific Risk Analysis Tasks</i>	Useful	Useful	Useful	Useful	High Relevance	High Relevance

Table A-I. Roadmap of Risk Management Utilization (cont.)

		TOPIC	Flow-down of objectives and integration with externally mandated requirements	Performance measures and parameters for objectives and requirements	Assignment and interpretation of performance markers and associated risk tolerances	Identification and characterization of individual risk scenarios	Identification and characterization of opportunity
		PART 1 SECTIONS	2.1.3, 5.1.3, 6.3	2.1.5, 3.2.3, 3.3.1, 4.2, 4.4.2, 4.5, 4.7, 4.8, 4.9.2, 4.9.3, 4.9.4, 4.9.5, 4.11.2, 5.1.2, 5.1.3, 5.1.4, 5.1.5, 5.3.1, 5.3.3, 5.4.3, 5.5, 5.6, 6.5.1, App C, App F	3.3, 4.2.3, 4.2.4, 4.4.3, 4.9.2, 4.9.3, 4.10.2, 5.1.1, 5.1.4, 5.3.3, 6.5.1, App C, App E	1.4, 2.1.6, 3.2.2, 3.2.3, 3.2.6, 3.3.5, 3.3.6, Ch 5, 6.3, 6.5.1, App C, App D, App E, App I	2.1, 2.2.3, 5.2, 5.3.3, 5.3.4, 5.7.1, 6.5.1, App I
		PART 2 SECTIONS	2.1	3.2, 4.1, 4.3, 4.4, 4.5	3.2, 4.1, 4.3, 4.4	4.3, 4.4	
PHASE	USER	FUNCTION					
All	Deliberation & Decision Board / Council	Providing High Level Directions and Making Decisions	Useful	Useful	Optional	Optional	Optional
	Higher Level Leader / Manager	Providing Top-Level RM Direction / Focus	Useful	Useful	Optional	Optional	Optional
Definition / Planning	Recommendation Board / Council	Providing Oversight	Recommended	Recommended	Useful	Useful	Useful
	Program / Project / Activity Manager	Organizing and Directing RM Tasks	Recommended	Recommended	Optional	Useful	Useful
		Making and Accounting-for Risk Relevant Decisions	Recommended	Recommended	High Relevance	Useful	Useful
	Program / Project / Activity RM Specialist	Executing RM Processes and Tasks	Recommended	High Relevance	High Relevance	Recommended	Recommended
	Program / Project / Activity Risk Analyst	Executing Specific Risk Analysis Tasks	Useful	High Relevance	High Relevance	High Relevance	High Relevance
Execution	Review & Recommendation Board / Council	Reviewing and Providing Oversight	Recommended	Recommended	Recommended	Useful	Useful
	Program / Project / Activity Manager	Organizing / Directing RM Tasks	Recommended	Recommended	Optional	Useful	Useful
		Making and Accounting-for Risk Relevant Decisions	Recommended	Recommended	High Relevance	Useful	Useful
	Program / Project / Activity RM Specialist	Executing RM Processes and Tasks	Recommended	High Relevance	High Relevance	Recommended	Recommended
	Program / Project / Activity Risk Analyst	Executing Specific Risk Analysis Tasks	Useful	High Relevance	High Relevance	High Relevance	High Relevance

Table A-I. Roadmap of Risk Management Utilization (cont.)

		TOPIC	Aggregate risk to objectives and requirements	Identification of leading indicators of unknown and/or underappreciated (U/U) risks	Individual and aggregate risk acceptability classification	Identification and characterization of cross-cutting risks and risk drivers	Display of individual and aggregate risks in various formats
		PART 1 SECTIONS	1.4, 2.1.6, 2.2.2, 3.2.1, 3.2.2, 3.2.3, 5.3, 6.5.1, App D	1.4, 3.2.4, 4.7.3, 5.2, 5.5, 5.6, 6.5.1, App B, App H, App I	3.3.2, 3.3.5, 3.3.6, App C, App E	5.3, 5.4.1, 5.5, 5.6, 6.5.1	3.3.2, 3.3.5, 3.3.6, 4.8, 5.3.3, 5.3.4, 6.5.1, App C, App E
		PART 2 SECTIONS	4.4		3.2.6, 3.2.8, 4.3, 4.5		3.2, 4.3, 4.4, 4.5
PHASE	USER	FUNCTION					
All	Deliberation & Decision Board / Council	Providing High Level Directions and Making Decisions	Optional	Useful	Useful	Optional	Useful
	Higher Level Leader / Manager	Providing Top-Level RM Direction / Focus	Optional	Useful	Useful	Optional	Useful
Definition / Planning	Recommendation Board / Council	Providing Oversight	Recommended	Recommended	Recommended	Optional	Useful
	Program / Project / Activity Manager	Organizing and Directing RM Tasks	Recommended	Recommended	Useful	Optional	Useful
		Making and Accounting-for Risk Relevant Decisions	Recommended	Recommended	High Relevance	Recommended	Recommended
	Program / Project / Activity RM Specialist	Executing RM Processes and Tasks	High Relevance	High Relevance	High Relevance	Recommended	High Relevance
	Program / Project / Activity Risk Analyst	Executing Specific Risk Analysis Tasks	High Relevance	High Relevance	High Relevance	High Relevance	Recommended
Execution	Review & Recommendation Board / Council	Reviewing and Providing Oversight	Recommended	Recommended	Recommended	Optional	Useful
	Program / Project / Activity Manager	Organizing / Directing RM Tasks	Recommended	Recommended	Useful	Optional	Useful
		Making and Accounting-for Risk Relevant Decisions	Recommended	Recommended	High Relevance	Recommended	Recommended
	Program / Project / Activity RM Specialist	Executing RM Processes and Tasks	High Relevance	High Relevance	High Relevance	Recommended	High Relevance
	Program / Project / Activity Risk Analyst	Executing Specific Risk Analysis Tasks	High Relevance	High Relevance	High Relevance	High Relevance	Recommended

Table A-I. Roadmap of Risk Management Utilization (cont.)

		TOPIC	Special considerations for cybersecurity risk	Communications and interactions leading to risk acceptance decision making	Detailed guidance on the steps of RIDM	Detailed guidance on the steps of CRM	Contents of the Risk Management Plan
		PART 1 SECTIONS	1.4, 2.1.6, 3.2.2, 3.2.5, 4.2.1, 4.5.1, 4.7.2, 5.3.4, 6.5.1, App B, App C	2.2.3, 4.9.5, 4.10, 5.4.4, 5.7, 6.3, 6.4, 6.5.1, App C, App E	4.4 to 4.10	5.2 to 5.7	6.5.1
		PART 2 SECTIONS	4.1	3.2, 4.3, 4.5	3.2	4.1 to 4.5	
PHASE	USER	FUNCTION					
All	Deliberation & Decision Board / Council	Providing High Level Directions and Making Decisions	Useful	Optional	Optional	Optional	Optional
	Higher Level Leader / Manager	Providing Top-Level RM Direction / Focus	Useful	Optional	Optional	Optional	Optional
Definition / Planning	Recommendation Board / Council	Providing Oversight	Recommended	Useful	Recommended	Optional	Recommended
	Program / Project / Activity Manager	Organizing and Directing RM Tasks	Recommended	Useful	Recommended	Optional	Recommended
		Making and Accounting-for Risk Relevant Decisions	Recommended	Recommended	Useful	Optional	Recommended
	Program / Project / Activity RM Specialist	Executing RM Processes and Tasks	High Relevance	Recommended	High Relevance	Optional	Recommended
	Program / Project / Activity Risk Analyst	Executing Specific Risk Analysis Tasks	High Relevance	Optional	High Relevance	Optional	Recommended
Execution	Review & Recommendation Board / Council	Reviewing and Providing Oversight	Recommended	Recommended	Useful	High Relevance	Recommended
	Program / Project / Activity Manager	Organizing / Directing RM Tasks	Recommended	Recommended	Useful	High Relevance	Recommended
		Making and Accounting-for Risk Relevant Decisions	Recommended	High Relevance	Optional	Useful	Recommended
	Program / Project / Activity RM Specialist	Executing RM Processes and Tasks	High Relevance	Recommended	Useful	High Relevance	Recommended
	Program / Project / Activity Risk Analyst	Executing Specific Risk Analysis Tasks	High Relevance	Optional	Recommended	High Relevance	Optional

Appendix B Example of Leading Indicators

Work reported in [1] and [2], suggests that the following design, organizational, and programmatic factors are among the principal leading indicators of U/U risk:

- Amount of complexity, particularly involving the interfaces between different elements of the system. Technical systems more prone to U/U failure are complex, tightly coupled systems that may make the chain of events leading to a potential disaster difficult for operators to recognize in their true level of danger.
- Amount of scaling beyond the established domain of knowledge. U/U risks may occur either from incrementally scaling up a design to achieve higher performance or incrementally scaling down a design to save on cost or time, without providing adequate validation.
- Use of fundamentally new technology or fundamentally new application of an existing technology. The use of new technology in place of heritage technology may lead to an increase in U/U risks when other factors within this list are not well handled.
- Degree to which organizational priorities are focused toward safety and reliability. U/U risks occur more frequently when top management is not committed to safety as an organizational goal, when there is no or little margin in the availability of qualified personnel, and when established technical expertise and organizational learning are not sufficiently valued.
- Degree to which the management style is hierarchical in a unidirectional upward direction (i.e., lower and middle level managers are more preoccupied with satisfying the directives and needs of higher-level managers than with tending to the needs and support of the personnel in their charge). Two-way flows of information and discussion are essential in technological systems to maximize the sharing of decision-relevant information among all personnel regardless of position in the organizational hierarchy.
- Degree of oversight when responsibilities are distributed among various entities. Interfaces between different elements of the system provided by different suppliers require stringent oversight by the managing agency.
- Amount of pressure to meet schedule and budget constraints. In particular, time pressure beyond the level of comfort is a fundamental reason for high human error rates.
- Likelihood of major or game-changing external events that affect the agency's direction, such as changes in the Administration or geopolitical upheavals. Such events impact the stability of long-term strategic planning and of constraints such as International Traffic in Arms Regulations (ITAR), etc.

References [1] and [2] provide useful guidelines on how various combinations of the above leading indicators affect the relative magnitude of U/U risks compared to the magnitude of known risks.

In addition to the leading indicators cited above, types of leading indicators that in their manifestations in past projects appeared to be correlated with specific types of performance risks and issues are identified and discussed in guide published by the NASA Office of the Chief Engineer [3]. Moreover, categories of leading indicators associated with an organization's ability to maintain its mandated core competencies have also been discussed in the technical literature. For example, the National Academy of Public Administration (NAPA) in 2007 [4] recommended for NASA use the following indicators of the health of a center, and more specifically of the risk of not being able to maintain a robust workforce:

- Median age of workforce
- Number of uncovered full-time equivalents (FTE)
- Ratio of fresh-out hires to total hires
- Ratios of civil service persons to contractors and supervisors to staff
- Center-by-center use of workforce incentives such as flexible work schedule, bonuses, and subsidized student loan payments.
- Percentage of people participating in training over the past year
- Number of turnovers and absenteeism
- Overall productivity rating
- Employee perceptions/assessments of management, e.g. from 360-degree feedback and *Best Places to Work* survey
- Number and severity of disciplinary actions
- Number of unfair labor practices and Equal Employment Opportunity (EEO) complaints

Similar lists can be postulated for physical assets and instructional assets. For example, the following short list of attributes can be thought of as leading indicators of the health of a center with respect to the availability and capability of an organization's physical assets:

- Median age of facilities
- Maintenance history of facilities
- Scale factors for testing
- Unaddressed cybersecurity threats
- History of changes to policies and procedures

There has also been a variety of studies within both Government and the commercial sector directed at identifying leading indicators of cost and schedule problems, including a National Academy of Sciences study of NASA cost overruns in 2010, a systems-engineering leading indicator guide by INCOSE and others in 2010, and a Government Accountability Office (GAO) report on causes of cost issues on selected NASA program delivered to Congress in 2008. In addition to the indicators listed above, the following conditions were prominently mentioned in these studies as indicators of impending cost growth and schedule slippages:

- Excessive number of requirements, often competing against one another
- Failure to provide adequate funding and lead time for the development of new technologies

- Failure to address interfaces between interdependent systems until late in the program
- Software complexity
- Supplier financial difficulties
- Insufficient quality assurance

B.1 References for Appendix B

1. Frenner, C., et al., An Assessment of the Inherent Optimism in Early Conceptual Designs and its Effect on Cost and Schedule Growth, European Aerospace Cost Engineering Working Group, May 2008.
2. Benjamin, A., Dezfuli, H., and Everett, C., Developing Probabilistic Safety Performance Margins for Unknown and Underappreciated Risks, Journal of Reliability Engineering and System Safety, Vol. 145 (329-340), January 2016.
3. NASA Common Leading Indicators Detailed Reference Guide, NASA Office of the Chief Engineer, January 2021.
4. NAPA Report (Panel of the National Academy of Public Administration), NASA: Balancing a Multisector Workforce to Achieve a Healthy Organization. February 2007.

Appendix C Further Technical Considerations on Risk Evaluation and Decision Making

This section discusses several technical aspects of risk characterization, assessment and decision-making that are related to the subjects introduced in this chapter, either in a direct way that warrants further explanation, or because they may be encountered as alternative, but still valid ways, of accomplishing in NASA-relevant contexts the same RM processes and tasks with which NPR 8000.4 and this handbook are concerned.

C.1 Probability as Continuous Complementary Parameter for Binary Performance Measures

The characterization or quantification of risk represents an organization's attempt to identify and evaluate the uncertainty affecting performance in the achievement of declared objectives, so that informed decisions can be made regarding how to best pursue such objectives. Characterizing performance uncertainty may appear more conceptually straightforward when performance can be expressed as an outcome within a continuous range of possible outcomes, e.g., when the possible total cost of a project is estimated vis-à-vis all the factors that may eventually determine it. In many other cases, however, the performance outcome may be more intrinsically binary. For example, in a crewed mission context, crew safety is typically treated as a binary objective, and performance is defined in terms of whether the crew may survive or be lost.

For such cases of “binary performance,” it is common to identify, for purposes of risk evaluation, communication, and decision-making before the evidence of the actual binary outcome becomes available, a probability parameter that can be used as a continuous performance measure (PM) complementing for such purposes the binary representation of possible outcomes. Thus, in such cases the probability that the binary outcome be positive or negative may be utilized as the continuous PM of choice, over which requirements or goals can be set, just like for any other continuous performance parameter. As a concrete example of this practice, in programs involving crewed missions NASA has used the probability of loss of crew, P(LOC), as an evaluation parameter for a future hypothetical binary PM outcome “crew is safe / crew is lost.” Correspondingly, what initially may have been expressed as a deterministic objective, i.e., “the crew must be kept safe,” has been translated into a probability-based performance requirement, PMK-R, to be satisfied by the best available estimations of P(LOC). That is, the qualitative binary requirement “the crew must be kept safe” has been translated into a quantitative PM requirement of the form: “the probability that $P(\text{LOC}) > \text{PMK-R}$ must be less than some small value X.”

As mentioned earlier, the setting of probability thresholds on a probability parameter may be confusing, since it appears to in essence set a requirement on the “probability of a probability.” As also mentioned earlier however, when the probability of an outcome is being used as a complementary substitute for the binary performance outcome itself, it is useful, in order to avoid any confusion, to think of and treat such probability parameter as one would any other type of system design parameter.

C.2 Assessment of Risk in Relation to Performance Markers and Risk Tolerance Levels

This section further addresses the concepts of performance markers and associated performance risk tolerance levels. Besides providing the means for translating top-level declarations of risk posture into definitions of operational levels of risk tolerance, and practical instruments of risk management, these concepts come into prominent play during any risk-informed Analysis of Alternatives (AoA). Accordingly, the material described here will be particularly relevant to the in-depth discussion of RIDM and CRM processes presented, respectively, in Chapters 4 and 5.

The notion of a performance marker, and possible shortfalls or surpluses with respect to it, has been introduced as a means of highlighting the idea that each performance measure comes with the definition of specific values in its range that the decision maker specifies as being the expectation for an activity or project. If the direction of goodness is toward the negative, as in the case of project cost, or P(LOC) (probability of loss of crew), then values higher than the marker value are to be avoided and represent a performance shortfall, and possibly a violation of requirements; whereas values lower than the marker value are desirable and represent a performance surplus. The opposite is true if the direction of goodness is toward the positive, as in the case of mass to orbit capability of a launch vehicle, or useful life on orbit of a satellite.

In the discussion that follows, which addresses the consistent application of risk posture and risk tolerance to the performance expectations in relation to identified activity objectives and/or decision alternatives, two different types of *performance markers* are considered: *performance constraints* / *performance requirements* and *performance targets* / *performance goals*. We note in this regard that, although there may be a different flavor of meaning between them, for the purposes of the following discussion the terms “constraints” and “requirements” will be used interchangeably, as they both are generally intended to similarly represent strict thresholds of performance, the violation of which is viewed as being highly undesirable, if not altogether unacceptable. Correspondingly, the terms “targets” and “goals” will also be used interchangeably, as they indicate thresholds of performance for which greater flexibility exists. Thus, the difference between the two types of markers which remains relevant to the present discussion is that a performance constraint or requirement reflects pre-set conditions and, once defined and set, is assumed to be non-negotiable between an *Acquirer* and a *Provider*, whereas a performance target or goal defines levels of performance that reflect the aspirations and preferences of stakeholders and organization leaders, but are generally not set in rigid terms and remain to some degree negotiable between *Acquirer* and *Provider* in the course of an activity or project execution.

Figure C-1 illustrates the concept for the case of a performance measure PMX for which two markers are set, one representing an established *performance requirement* (or *constraint*), the other a *performance goal* for which some degree of discretionality exists, as reflected by the magnitude of the *requirement-to-goal margin* that its value may have with respect to the value of the requirement/constraint itself.




Figure C-1. Risk Levels Expressed by Performance Markers in Performance Measure pdf

For risk-informed deliberation purposes, any given marker can be set by the decision maker to correspond to some allowable percentile value or range of a corresponding performance measure pdf. The pdf percentile values associated with a marker corresponds to a **risk tolerance level** (RTL) that the decision maker wishes the deliberation process to adopt as the limit-criterion for the possibility (i.e., in quantitative terms, probability) that the performance measure PMX fall on the “wrong side” of that marker.

In Section 0, the concept of *risk tolerance* was associated with to the identification of two types of risk boundaries used in the deliberations that may be made by a decision-maker, i.e., an RTL-G was associated with a *performance goal or target*, whereas an RTL-R was associated with a *performance requirement or constraint*. These threshold values are also referred to, respectively, as “*risk acceptability / watch boundary*” and “*risk tolerance / response boundary*,” or as “*acceptable-to-marginal risk boundary*” and “*marginal-to-unacceptable risk boundary*,” when defined as the “*green-to-yellow*” and “*yellow-to-red*” limits between the color-coded risk regions of common forms of risk representation and display.

C.3 Use of Cumulative Distributions in the Setting of Performance Markers and Determination of Associated Risks

The relation between performance markers and risk levels is best illustrated and discussed by using the cumulative distribution function (CDF) representation of performance measure uncertainty, as shown in Figure C-2.




Figure C-2. Comparative Illustration of CDF and pdf of a Performance Measure

A CDF curve represents the mathematical integral of the probability distribution function (pdf) of a parameter or variable. Thus, given the CDF of a performance measure, PMX, and any given PMX value, call it V, on the abscissa axis, the corresponding CDF value intersected on the ordinate axis via the CDF curve is a probability P, such that:

$$P = \text{probability that } PMX < V ;$$

or in shorthand:

$$P = p (PMX < V)$$

Because a CDF expresses full probability values, it must be noted that such values are always in the range between 0 and 1, and in fact, being the integral of a pdf, which is always positive in value, a CDF is always a non-decreasing function of the underlying performance measure.

Applying the above to the two performance markers shown in Figure C-2 for the measure PMX, which represent, respectively, a Performance Requirement (PMK-R) and a Performance Goal (PMK-G), yields the following:

$$A = p (PMX < PMK-R) = [\text{risk of not meeting the Performance Requirement}]$$

$$A + B = p (PMX < PMK-G) = [\text{known risk of not meeting the Performance Goal}]$$

$$A + B + C = 1 = [\text{probability of the range of PMX possible values}]$$

Figure C-2 shows how the markers and probability values originally illustrated in Figure C-1 are translated from the pdf to the CDF form of representation. For the purposes of the present discussion the CDF representation of a PM uncertainty and risk distribution offers the practical

advantage of permitting a direct readout of known-risk levels associated with performance markers. Thus, referring again to Figure C-2, the value A, which represents the known constraint risk, or calculated probability of not meeting the performance requirement, can be read directly off the ordinate axis of the CDF curve representation of the PMX performance measure distribution function. The same is true for the value [A + B], which represents the known target risk, or calculated probability of not meeting the performance goal.

C.4 Mutual Constraints between Performance Markers and Associated Performance Risk Tolerances

The illustration in Figure C-2 and the above related discussion of the relation between performance marker and risk levels have important implications for understanding the constraints that the setting of risk tolerance limits generates in terms of the performance markers that may also be set, and vice versa.

Given the uncertainty distribution of a performance measure like the one notionally depicted in Figure C-2, the *risk tolerance level* (RTL) for that performance measure can be expressed by setting a risk threshold value, i.e., a maximum value of known performance risk that can be tolerated or accepted, which in practice means setting a maximum value for the probability that the performance marker of concern is not complied with. Given an assessed probability distribution for the performance measure of interest, which implies a representation of known risk thereof, setting a risk/probability threshold – i.e., limits of acceptable / non-acceptable risk, such as the RTL-R and RTL-G values previously discussed in Section 0 – also defines the maximum value of a correspondingly admissible **performance marker**. Vice versa, if a minimum value of performance is set in the form of a **performance marker**, this also defines an admissible range for a compatible performance **risk tolerance level** (RTL).

To illustrate the above, Figure C-3 shows the lower probability / risk region of a performance measure CDF. The two cases referred to above are depicted, respectively, in Figure C-3a and Figure C-3b.

Figure C-3a shows that, given a performance measure CDF as the one depicted, if a performance *risk tolerance level* (RTL) (i.e., RTL-G or RTL-R value per Section 3.3.1 terminology) is established, the range of *performance marker* (PMK) values that may be set compatibly with such a value is as shown in the figure: an admissible PMK cannot be set at a value greater than the intersect obtained by projecting the RTL value on the performance measure axis via the given CDF curve. If a PMK value were chosen and set outside the admissible range shown in Figure C-3a, it would be impossible for the system or mission design alternative represented by the depicted CDF to satisfy both the pre-established RTL and the chosen PMK value. Similarly, Figure C-3b shows the admissible range for setting an RTL, once a PMK value was set as shown.




Figure C-3. Relationship between Performance Risk Tolerance Levels and Performance Markers

C.5 Risk Representations for Cases of Inverted Direction of Performance Measure Goodness

As mentioned above, there frequently exist situations where better performance is expressed by lower values of the performance measure. Obvious examples are program cost and the time to reach a milestone. In the technical performance domain, an example of this type of situation is the resolution capability of an imaging instrument, when this is expressed as the minimum linear distance between two distinct objects that the instrument can “resolve.” Better performance is in this case expressed by smaller resolution distance. Another example is found in the safety domain, where the probability of loss of crew, $P(\text{LOC})$, is commonly used as a measure of performance: the lower the value of $P(\text{LOC})$, the better the safety of the system of interest is reasonably assessed to be.

When using the uncertainty distributions of performance measures that characterize these situations, the definitions and meaning of their pdf and CDF forms of representation remain the same. However, in these cases performance risk is expressed by the probability that the performance measure values will exceed, rather than be lower than, a given performance marker value. Figure C-4 illustrates this, by comparing side by side the representations of performance risk for two hypothetical performance measures, PMX and PMY, for which goodness of performance grows in opposite directions. For the Case 1 shown in the figure (the case where goodness increases with the value of the performance measure), risk with respect to a *performance marker*, PMK, is expressed by the CDF value at the abscissa value PMK, indicated in the figure as A1:

$$\text{Risk [Performance Marker not met]} = p(\text{PMX} < \text{PMK}) = A1$$

For Case 2 (goodness decreasing as performance measure value increases), performance risk with respect to the target PMK is, instead, expressed by the complement to unity, B2 in the figure, of the CDF value, A2, at the abscissa value PT:

$$\text{Risk [Performance Marker not met]} = p(\text{PMY} > \text{PMK}) = 1 - p(\text{PMY} \leq \text{PMK}) = 1 - A2 = B2$$




Figure C-4. CDF Representations for Performance Measures with Opposite Directions of Goodness

In situations as in Case 2 it is common practice to use for risk representation the complement to unity of the uncertainty CDF, which is referred to as the Complementary Cumulative Distribution Function (CCDF). Figure C-5 completes the illustration of this type of situation. It shows the same case depicted in Figure C-4 Case 2, but alongside the pdf it uses, instead of a CDF, a CCDF representation of the performance measure uncertainty distribution.




Figure C-5. Pdf / CCDF Representation of Risk / Probability of Exceedance of a Performance Value (Inverted Direction of Goodness Situations)

C.6 Alternative Means of Defining Risk Tolerance Levels and Classifying Risk

The discussion in preceding sections has used as its basic reference case a situation where an *Acquirer* and a *Provider* have established and/or negotiated two distinct *performance markers* in relation to a specific PM, in the form, respectively, of a *performance requirement* (or *constraint*) and of a *performance goal*. This is a common, but not the only type of situation that may be of interest in regard to *risk tolerance* and *risk classification*, as a couple of different types of contexts are also relatively common. The first such alternative context concerns an activity or project situation in which a PM and corresponding objective are targeted by means of a single *performance*

marker, which in such a case might be set either as a *performance requirement* or as a *performance goal*, depending on the nature of, and risk posture applied to, the activity of concern. The second situation is one encountered in certain regulatory contexts, mostly related to public and or environmental safety, where risk classification boundaries and regions may be defined by means of entire “limit CDF” or “limit CCDF” curves. These two cases are further illustrated below.

C.6.1 Single Performance Marker with Multiple Related Risk Tolerance Levels

From the preceding discussion, and more specifically using for reference the illustration provided in Figure C-1, it is evident that, once a *performance marker* is established for a given PM, the range of probability of PM exceedance or non-exceedance of the marker value, which by definition represents the **performance risk** of concern with respect to the objective “measured” by the PM, is a continuum in the range between 0 and 1. Consequently, if the decision maker or analyst desires to classify the risk in that range via the traditional three-color scheme, they may define two distinct boundaries: one to define the distinction between *acceptable* and *marginal risk* – the Green to Yellow risk boundary – and one to define the distinction between *marginal* and *unacceptable risk* – the Yellow to Red risk boundary. Figure C-6 illustrates this RTL definition scheme.




Figure C-6. Illustration of Dual RTL Definition for a Single Performance Marker

It should be noted that in this case with a single performance marker, the RTL for the goal is set at a lower value than the RTL for the requirement, due to its aspiratory nature. Additionally, instead of expressing risk posture by definition of different tolerances for two different marker levels of performance, defines it by splitting the level of tolerance into an “*watch boundary*” and a “*response boundary*” for the possible values of the risk/tolerance parameter itself.

C.6.2 Risk Tolerance Level Boundaries Defined as Two-Dimensional Curves

There exist situations where the potential for a shortfall with respect to some objective may span a wide range of performance measure values, and the identification of specific performance targets is not very definitively established within that range. These are cases where a desirable level of

performance may exist, however, rather than drawing “lines in the sand” in regard, the affected organization seeks primarily to control the magnitude of a shortfall. For example, a small cost-overrun may appear to be almost inevitable due to its high probability in a given project, and therefore is to be realistically tolerated, whereas a large cost-overrun, although low enough in probability to represent the same “expected-dollar-loss” (i.e., probability weighted loss value) may be viewed as utterly unacceptable because of its consequences, if these are realized.

For these situations the tolerance level for the total risk affecting the performance objective – i.e., risk from all “individual risk scenario contributions” – is not necessarily defined in terms of probability of a performance measure exceeding or non-exceeding pre-defined and fixed performance markers. Instead, risk tolerance levels may in such cases be expressed as probability limits on exceeding possible performance measure shortfalls of varying magnitude. The definition of such limits and boundaries can be accomplished by use of cumulative distribution functions (CDFs) or complementary cumulative distribution functions (CCDFs) of the performance measure (PM).

The choice between the use of a CDF or CCDF is primarily a matter of convenience, depending on the PM “direction of goodness.” It is also always possible to define a performance measure in such a way as to show its direction of goodness as being from right to left on the x-axis as in Figure C-7, so that its possible shortfalls are shown as increasing in the positive direction. A CCDF is then well suited to express the performance measure risk as the probability that a performance measure shortfall exceeds a given magnitude in its range of possible outcomes.




Figure C-7. Risk Tolerance for Performance Markers on a Continuum

The situation depicted in Figure C-7 is one conceptually equivalent to establishing a continuum of performance markers rather than a discrete set. Accordingly, the corresponding **risk tolerance boundaries** are expressed not just in terms of a single-value threshold on the probability of failing to meet a marker value, but in terms of full CCDF probability-threshold-curves for exceeding a given shortfall magnitude.

Figure C-7 shows cases of PM distributions that, consistently with the above, are classified as representing *acceptable*, *marginal*, or *unacceptable risk*, based on their position with respect to the predefined CCDF risk boundaries shown in the figure. It should be noted, however, that the risk classification and labeling exemplified in the figure essentially corresponds to the criterion that a PM distribution is classified in one way or another based on whether or not it impinges on any risk region boundaries. This is not the only risk classification criterion possible for a case like the one exemplified, as a classification could be based instead on the confidence level at which an impingement occurs or not. That is, depending on risk posture and tolerances adopted by an organization and its decision makers, a PM distribution could be classified as acceptable risk if its median value (i.e., its 50th percentile, or 50% percent confidence value) falls within the acceptable risk region, even though its higher percentiles (e.g., its 95% confidence level value) were to fall in the marginal, or even unacceptable risk regions.

Response and Watch Boundaries

In the context of the decision-making activities conducted within the risk management framework, the boundaries between the risk regions may have a special terminology. The demarcation between the yellow and red risk regions, which represents the boundary between “marginal” and “unacceptable” risk is sometimes also called the “Response Boundary,” since the yellow-to-red transition usually implies the need for an action to mitigate the risk. Similarly, the demarcation between the green and yellow risk regions, which represents the boundary between “acceptable” and “marginal” risk, is sometimes referred to as the “Watch Boundary,” implying that an action may be considered but is not imminently needed. In this handbook, the two types of boundaries are generally referred to as the “Acceptable Risk Boundary,” signifying that risk below the marked level is deemed to be acceptable, and the “Unacceptable Risk Boundary,” signifying that risk above the marked level is deemed to be unacceptable. In essence, the terms “acceptable” and “unacceptable” refer to risk from a perspective of Risk Posture, while the terms “watch” and “response” refer to it from the point of view of actionable decision making.

In Figure C-7 the values of the performance measures corresponding to the intersection of any given risk tolerance level – i.e., in the language of Section 3.3.1, an equivalent of either the RTL-G or RTL-R probability values – with the acceptable and unacceptable risk boundaries may be considerate as the continuous equivalents of the “performance markers” first introduced and discussed in Section 3.3.1. More specifically, the “acceptable risk boundary” corresponds to a continuum of possible “performance target/goal” values, whereas the “unacceptable risk boundary” corresponds to a continuum of possible “performance constraint/requirement” values.

C.7 Assessment and Display of Risk Relative to an Objective

The discussion on the definition of *performance markers* and *risk tolerance levels* has illustrated the conceptual quantitative underpinning of how both *individual risk scenarios* and *aggregate risk* can be assessed and evaluated in relation to the affected *objective(s)* and *performance measure(s)*. It should be understood, however, that such a quantitative underpinning does not prevent the

application of a graded approach, whereby risk(s) can be estimated and displayed on the basis of characterizations expressed in more qualitative form than the full derivation of pdfs, CDFs, or CCDFs. This section, and the subsections therein, seek to provide further discussion and insight into the correspondences and relationships that exist between the possible means of risk assessment and evaluation, proceeding from the conceptual baselines that have been insofar introduced.

The process of assessment and display of individual risk scenarios and/or aggregate risk can be carried out in conceptually straightforward successive steps. These may differ to some degree depending on whether the assessment can be executed directly in quantitative mode, or whether a more heuristic semi-quantitative approach is necessary, thus the discussion that follows is organized accordingly. The cases where the assessment can be carried out in full quantitative mode are presented first, as the more qualitative assessment processes are more readily explained and understood after the quantitative ones have been presented.

CDFs, CCDFs, and Probability as Frequency or Degree of Belief / Confidence Level

In probability theory, the cumulative distribution function (CDF) of a random variable X , evaluated at x , is the probability that $X < x$, and the complementary cumulative distribution function (CCDF) evaluated at x is the probability that $X > x$; hence the respective alternative names of “probability of non-exceedance” and “probability of exceedance,” and their relationship $CDF + CCDF = 1$. CDFs, CCDFs and their relationship with an underlying probability density function (pdf) have been discussed in Sections C.3-C.5. Figure C-2 illustrates the pdf-CDF relationship, while Figure C-5 illustrates the pdf-CCDF relationship.

CDFs and CCDFs can be used to represent physical/objective probability, which relates to the frequency of random physical events (e.g., radioactive decay or the rolling of dice) and thus represents “aleatory uncertainty”; or evidential/subjective probability, which represents “degree of belief” or “confidence level,” i.e., the degree to which a hypothesis or statement is deemed credible based on some form of evidence, and is thus a representation of “epistemic uncertainty” and uncertainty in knowledge [1] (see also Section 4.7.3.2).

From an evidential probability perspective, one can talk about the “probability of a probability.” In that phrase, the second probability is a physical/objective parameter expressing an expectation of frequency in time or in trials, whereas the first refers to a degree of belief, or confidence, that the second will take on certain values. For example, a CDF or CCDF for $P(LOC)$, expresses degrees of belief in values of the frequency of a mission failure causing loss of crew: thus $CDF=0.8$ (i.e., $CCDF=0.2$) for a $P(LOC)$ value = 0.01 indicates an 80% confidence/belief that the frequency of a LOC event is less than 0.01 per mission (i.e., a 20% confidence/belief that it is greater).

C.8 Communication of Risk has Individual and Aggregate Dimensions

Risk communication involves both individual risk scenarios and aggregate risks to objectives. As illustrated in Figure C-10, the status of the risks to an activity’s objectives can be communicated via a spider chart (see the left-hand side of the figure), and the top individual risk scenarios threatening each objective can be communicated via a set of spider charts – one for each objective (see the right-hand side of the figure).




Figure C-10. Complementarity of Individual and Aggregate Risk

The two sides of Figure C-10 have different, complementary, anchorings. The left-hand side is anchored to the tolerability of the total risk, whereas the right-hand side is anchored to the adequacy of the risk management effort to address each of the discrete vulnerabilities represented in the individual risk scenarios. An organization whose risks to its objectives are acceptable (left-hand side) can validly claim that its objectives are worth pursuing despite the risks. An organization whose top individual risks are acceptable can validly claim to have adequately controlled the known vulnerabilities.

The complementary nature of the left-hand and right-hand sides of Figure C-10 is closely related to the notion of “adequate safety” in the NASA System Safety Handbook [2]. In [2], adequate safety is defined as 1) achieving a system that meets or exceeds the minimum tolerable level of safety, and 2) achieving a system that is as safe as reasonably practicable (ASARP). The status of these two components of adequate safety is communicable via spider charts addressing aggregate risk to objectives and individual risk scenarios respectively.

Figure C-10 reinforces the point that spider charts are planning and communication tools, not analysis tools. They map the organization’s already-assessed risks against its RTLs, graphically making the case for their acceptability. When risk is in the red, a risk management response is needed (potentially including elevation) to rectify its intolerability. When risk is in the green, stakeholders can be assured of adequately low risk. All this is contingent, of course, on proper prior analysis of risk (including risk margins to account for UU risk) and prior establishment the organization’s risk posture.

C.9 Special Considerations for Adversary-Initiated Risks

As discussed in NPR 8000.4 and in Section 3.2.5 of this handbook, adversary-initiated risks (i.e., intentional scenarios) may be dealt with on a conditional basis because of the difficulty of

predicting threat likelihoods. When analyzing conditional risks, it is assumed that the attack has occurred. The analyst's tasks in relation to these scenarios are to evaluate the likelihood of the scenario unfolding as the assailant intends, based on the adequacy or inadequacy of the system defenses and controls, and to determine the character and magnitude of the resulting consequence(s). Later, information about threat likelihoods – e.g., from intelligence sources – may be incorporated into the conditional analysis so as to have a more informed basis for determining whether further mitigation of the risk is needed.

It is noted that, because of the fact that intelligence on potential attacker threats might tend to be more qualitative than quantitative, the process for incorporating this information into the assessment may also need to be flexible enough to accommodate this more qualitative than quantitative nature of portions of the available information. In this perspective, for each scenario the task would be to identify the value of the probability of attack which permits the target value of the performance measure to be met at a specified *risk tolerance level*. If the probability of attack for any scenario is clearly greater than that value, then additional mitigation may be needed.

As an example, suppose that in the presence of adversary-initiated Threat X, the set constraint for the realization of undesired consequences during the mission is 0.01 and the risk tolerance for exceeding that constraint is 0.5 (meaning that the decision-maker wants to be at least 50% confident that the likelihood of adversary-induced consequences is no greater than 0.01). Suppose also that the initially calculated value for the likelihood of adverse consequences in the actual presence of the adversary-initiated Threat X is 0.04 at a confidence level of 50%. Intelligence information is then assessed to determine whether the (absolute) probability of the Threat X actually materializing can be deemed to be less than 0.25 (the ratio of 0.01 to 0.04). If that is the case, then the (absolute) probability of adverse consequences at a confidence level of 50% can be inferred to be less than the targeted value of 0.01. If, however, the probability of the Threat X actually materializing is believed to be greater than 0.25 using the same reasoning and the same confidence level, then the inferred likelihood of adverse consequences would be greater than 0.01 at the stated risk tolerance level, and the cybersecurity defenses, i.e., the cybersecurity control provisions adopted against a Threat X type of attack would need to be strengthened.

C.10 Risk Tolerance for Interdependent Performance: The Joint Confidence Level

Risk tolerance values are sometimes defined in terms of multiple correlated objectives and associated performance measures. At NASA, the Joint Confidence Level (JCL) requirement for cost and schedule is an example of this. Except where exceptions are granted, NPR 7120.5 [3] specifies that it shall be demonstrated at each review, through analysis, that the probability of a program or project being completed within *both* cost *and* schedule is at least 70% (i.e., the joint confidence level is 70%). This is equivalent to stating that the risk tolerance for *either* the cost *or* the schedule having a greater-than-zero shortfall with respect to its mandated expectation is 30%. The risk attitudes for cost and schedule, therefore, are in this case not independent and, similarly, there may be other cases where certain objectives are highly correlated and therefore the risk attitude pertaining to these may be expressed in terms of joint risk tolerance levels.

C.11 References for Appendix C

1. Hajek, A., Stanford Encyclopedia of Philosophy, Interpretations of Probability. First Published October 2002, Revised August 2019.
2. NASA Special Publication, NASA/SP-2010-580 Version 1.0, NASA System Safety Handbook, Volume 1: System Safety Framework and Concepts for Implementation. November 2011.
3. NASA Procedural Requirements, NPR 7120.5F, NASA Space Flight Program and Project Management Requirements w/Change 3. August 2021.

Appendix D Aggregate Risk Effects of Cumulative-Consequence Individual Risk Scenarios

Given a performance marker, an aggregate RTL for the performance marker, and an estimate N of the number of individual risk scenarios that threaten performance, Section 3.3.4.1 has provided the heuristic rule for the determination of individual risk scenario RTLs. In general, such a rule works equally well for HCIRSs of a non-cumulative nature and for CCIRSs. Given that the nature of the latter may be of special concern, in the sense that the potential magnitude of the consequences of a multiple materialization of CCIRSs may appear to be almost unbounded, it is useful to present some examples that show how the heuristic IRTL-setting rules of Section 3.3.4.1 work for representative cases of CCIRS cumulation, in terms of the resulting Aggregate Risk profiles.

To illustrate such applications of the heuristic IRTL rules defined above, let us thus consider a realistic situation concerning a Project Cost, which is defined as follows:

- Project Cost (PC) is the PM being looked at – a situation representative of the worst type of aggregate risk, i.e., one where individual risk scenario consequences, if realized, are cumulative.
- The baseline cost (BC) is estimated to be $BC = \$750M$.
- A Cost Requirement Marker, PMK-R, is set at $\$900M$ with a corresponding Aggregate Risk ARTL-R set at 5% (low risk tolerance).
- A Cost Goal Marker, PMK-G, is set at $\$800M$ with a corresponding Aggregate Risk ARTL-G set at 15% (medium-low risk tolerance).
- The initial estimate of the number of AR contributors is that in the course of execution of the project there may be $N = 5$ CCIRSs that would be significant contributors (in the following these will simply be referred to as IRSs).
 - Accordingly with the above, per the heuristic rule presented in Section 3.3.4.1, the IRS risk tolerance levels (IRTLS) corresponding to the markers are set as:
$$IRTL-G = ARTL-G / N = 15\% / 5 = 3\%$$
$$IRTL-R = ARTL-R / N = 5\% / 5 = 1\%$$

The two examples ensuing from the above which are representative of realistic situations are presented in the following as “Case A” and “Case B.” The intent of these examples is to show that the IRTL-setting criteria of Section 3.3.4.1 remain practically applicable and useful even in limit-cases of cumulation of IRS consequences.

Case A

In Case A, besides the conditions defined by the bullets listed above, it is further estimated that each of the IRSs may have the potential of causing an added cost (AC) of about $\$200M$ over the baseline cost value (BC). This reflects a potentially severe situation because the occurrence of any one of the IRSs would result in a project cost $C = BC + AC = \$950M$, above the Cost Requirement Marker, i.e., it would result in $C > PMK-R$.

According to the adopted risk tolerance criteria, even if the potential for a violation of the requirement marker exist, an IRS that carries that violation as its consequence can be considered to be at an “Acceptable / GREEN” level of classification if its probability of occurrence – which we refer to in the following as p_{IRS} – is not greater than 1%.

For the limit “worst case” of this assumed situation, i.e., the case where:

$$p_{IRS} = 1\% ,$$

the Aggregate Risk resulting from the possible cumulative effects of the underlying five IRSs is represented by the following distribution of mutually exclusive outcomes:

- A. No IRSs occur – PC = BC (no cost overrun), with probability = 0.951
- B. Any One IRS occurs – PC = \$950M (\$200M cost over BC), with probability = 0.048
- C. Any Two IRSs occur – PC = \$1.15B (\$400M cost over BC), with probability = 0.001
- D. Any Three IRSs occur -- PC = \$1.35B (\$600M cost over BC), with probability < 1.0E-5
- E. Any Four IRSs occur -- PC = \$1.55B (\$800M cost over BC), with probability < 1.0E-7
- F. Any Five IRSs occur -- PC = \$1.75B (\$1B cost over BC), with probability < 1.0E-10

A more complete representation of the spectrum of probability and cost-consequences of this AR is provided below in Figure D-1, in both plot and tabular form.




Figure D-1. Aggregate Risk Produced by Cumulative Effect of “GREEN” IRSs

The significance of the above is that, due to the cumulative nature of the initial five IRSs being considered, the resulting project cost (C) Aggregate Risk (AR) is represented by the Probability Mass and CCDF distributions of the above outcomes A through F, as shown in the graph and table included in Figure D-1. The key take-away from the example is that, if each of the AR-contributing IRS satisfies the IRTL-R criterion for an “Acceptable / GREEN” classification, the resulting probability distribution of cumulative outcomes is such that the key parameter of interest for AR classification, i.e., $p(C > PMK-R)$, the probability of a cost requirement violation, remains below the ARTL-R value of 5%.²¹

²¹ Given the discrete nature of the possible outcomes assumed for this example this probability coincides with the CCDF value @ $C = \$750M$, which is 4.9%. To avoid confusion in this regard, one must recall that, for a given value

The conclusion that can be drawn from the above example is that the adoption of and compliance with the IRTL setup heuristic rule presented in Section 3.3.4.1, results in a “GREEN” AR classification when the contributing IRSs are individually kept at a GREEN level, based on the corresponding IRTL definitions. This holds in this example despite the conservative assumption made, i.e., that not only the IRSs are of the CCIRS variety, but also that the occurrence of a single one of them would be enough to cause the AR requirement marker, PMK-R, to be violated.

Case B

In Case B we consider a situation where an IRS is not in the “Acceptable / GREEN” category, but in the “Marginal / YELLOW” one. If we call p_R the probability that $C > \text{PMK-R}$ by occurrence of a single IRS and p_G the probability that $C > \text{PMK-G}$ for a similar occurrence, and we refer to both the IRTL-R = 1% and IRTL-G = 3% values, a “Marginal / Yellow” IRS classification, per the IRTL criteria previously defined, corresponds to a combination of p_R and p_G values such that:

$$\begin{aligned} p_R &\leq 1\% \\ p_G &> 3\% , \end{aligned}$$

To make the above realistically possible it is assumed for this case that each single IRS, if realized, may have two distinct additional cost outcomes, AC1 = \$75M with conditional probability $p_{1/\text{IRS}}$, and AC2 = \$150M with conditional probability $p_{2/\text{IRS}}$. Starting from these assumption, and also assuming the limit-value $p_R = 1\%$, one can apply simple calculations to derive a corresponding limit-value of the probability of occurrence of a single IRS, which as in Case A we refer to as p_{IRS} , and corresponding values of $p_{1/\text{IRS}}$ and $p_{2/\text{IRS}}$. The applicable mathematical formulas for such derivations are:

$$\begin{aligned} p_R &= 1\% , && \text{as the assumed “limit condition”} \\ p_G &= 4\% , && \text{as an assumed } p_G \text{ value that satisfies the condition } p_G > 3\%; \text{ together with} \\ &&& \text{the } p_R = 1\% \text{ condition this results in the initially intended IRS} \\ &&& \text{classification of “Marginal / YELLOW”} \\ p_G &= p_{\text{IRS}} \times p_{1/\text{IRS}} && \text{-- this formula follows from the fact that the occurrence of a single IRS} \\ &&& \text{with consequence AC1 = \$75M comports the violation of the goal marker} \\ &&& \text{PMK-G = \$800M} \\ p_R &= p_{\text{IRS}} \times p_{2/\text{IRS}} && \text{-- this formula follows from the fact that the occurrence of a single IRS} \\ &&& \text{with consequence AC2 = \$150M is at the boundary of violating the} \\ &&& \text{requirement marker PMK-R} \\ p_{1/\text{IRS}} + p_{2/\text{IRS}} &= 1 && \text{-- signifying that, if one of IRSs occurs, it can only be of one of the two} \\ &&& \text{types that, respectively, would cause either the AC1 or AC2 type of} \\ &&& \text{consequence} \end{aligned}$$

of its argument, the corresponding CCDF value represents the probability that the parameter being considered be greater than that argument value. Given the discrete nature of the possible outcomes, the CCDF value at $C = \$750\text{M}$ represents the probability that C may be at any of the other discrete values that are greater than \$750M, i.e., that it may be \$950M, or any of the other four possible outcomes for which C is greater than \$750M.

Use of the above assumptions results in the below values of the IRS unconditional probability of occurrence and of the two associated conditional sub-outcomes, i.e., IRS with consequence AC1 and IRS with consequence AC2:

$$p_{\text{IRS}} = 5\%$$

$$p_{1/\text{IRS}} = 80\%$$

$$p_{2/\text{IRS}} = 20\%$$

The resulting project cost (C) Probability Mass and CCDF distributions are shown in the graph and table included in Figure D-2 below. More specifically, the figure shows that the values of the cost probability distribution relative to the pre-established ARTL-R and ARTL-G values are, respectively:

$$p(C > \text{PMK-R}) = .89\% < 5\% \text{ (ARTL-R value)}$$

$$p(C > \text{PMK-G}) = 23\% > 15\% \text{ (ARTL-G value)}$$




Figure D-2. Aggregate Risk Produced by Cumulative Effect of “YELLOW” IRSs

The above combination of probability criteria corresponds to a classification of the AR that mirrors that of the underlying and contributing IRSs, i.e., “Marginal / YELLOW.”

It must be noted that, while the two illustrations of the heuristic IRTL-setting criteria provided above correspond to realistic situations of possible IRS cumulative effects on AR, and show the practical usefulness of the criteria for those situations, they nevertheless cannot be interpreted as being a proof of general validity of the criteria under any other alternative circumstances of IRS aggregate effects. As previously mentioned, the assessment and evaluation of IRSs according to criteria that consider them individually must be complemented by a parallel and well-organized assessment of aggregate risk for each of the organizational objectives and performance dimensions of concern.

Appendix E Stepwise Recap and Example of Risk Leadership and Objectives-Driven Risk Management Application

In order to provide a clear understanding of the conceptual and operational flow of Risk Leadership and Objectives-Driven Risk Management (ODRM) action and implementation processes, this example highlights the key steps that are involved in linking together the fundamental concepts introduced in Chapters 2 and 3. It should be noted that this example does not aim to cover all the intricacies and details of an execution, but rather focuses on the implementation of the principal concepts of interest.

One of the crucial aspects of effective ODRM is *risk leadership*, as applied by the individual leaders and management teams responsible for overseeing and guiding the ODRM activities within an organization. The example provided here emphasizes the importance of establishing *risk leadership* as a foundation for successful ODRM implementation. When project and activity leaders with the necessary authority and risk perspective provide the necessary guidance, the organization can ensure that risk management efforts are properly coordinated and aligned with its overall strategic objectives.

Another significant aspect addressed in the example is the identification and definition of *risk posture*. This refers to the organization's overall stance towards risk, including its tolerance of or “appetite” for it. By clearly defining its *risk posture*, the organization can establish a consistent and unified approach to managing risk within an activity or project, including allocation of risk posture to any supporting organizational units. The example highlights the need to carefully assess and articulate the organization's *risk posture* as it sets the tone for subsequent risk management activities.

In relation to the operational implementation of *risk posture*, the example underscores the importance of identifying and defining *risk tolerance levels (RTLs)*. Risk tolerance refers to an organizational unit's limit of acceptable likelihood of falling short of achieving a defined top-level objective. By establishing risk tolerance levels, organizations can make informed decisions regarding the acceptability of specific risks and the allocation of resources for their mitigation. The example emphasizes the need for a systematic approach for the definition of RTLs, which, besides consistency with the above cited identification of an internally desired *risk posture*, may also involve considering legal and regulatory requirements, and aligning with industry best practices.

To illustrate the operational application of the above concepts, the example focuses on the risk acceptance classification of identified and assessed individual risk scenarios, as well as of overall aggregate risk. This step involves evaluating each identified risk based on predefined criteria and categorizing them according to their level of acceptability. By classifying risks, organizations can prioritize their mitigation efforts and allocate resources effectively. The example highlights the significance of considering both individual risk scenarios and the overall aggregate risk to gain a comprehensive understanding of the organization's risk landscape.

Lastly, the example emphasizes the use of an effective communication format to display the analyzed risks. By representing risks in a summary format, organizations can visualize for managers' decision-making purposes the relationship between their likelihood and potential

impact. This facilitates the identification of high-risk priority areas that require immediate attention and enables effective communication of risk information to stakeholders. The example highlights the importance of a clear and intuitive visualization format to enhance understanding and decision-making.

Overall, this example demonstrates how the implementation of key concepts such as *risk leadership*, *risk posture* definition, *RTL* identification, *risk acceptance* classification, and *risk display* contribute to a robust and systematic approach to RM.

The above-mentioned concepts and elements of RM implementation are sequentially interconnected, as illustrated by the diagram in Figure E-1. The diagram illustrates the flow of some of the related ODRM actions and identifies the principal agent(s) responsible for executing them. It should be noted that the illustration is conceptual and does not aim to represent the various possible variations in organizational hierarchies and management interactions that can exist.




Figure E-1. Conceptual Flow of Risk Leadership and Management Implementation

The diagram is divided into two main portions. The top portion symbolizes the primary flow of directives and information within the organization that shape and determine how the key ODRM actions, identified in the lower portion of the diagram, are executed. This illustration is intended to emphasize the importance of effective communication and direction from higher levels of management to ensure the successful implementation of ODRM practices throughout the organization. The lower portion of the diagram identifies the key ODRM actions that implement the *risk leadership* directives and illustrates their logical and sequential flow. These actions represent core activities that implement ODRM within the organization via a systematic and self-consistent approach, as it is required to manage risks effectively and in full alignment with the strategic objectives of the entire organization.

Example scenarios and explanations of the individual elements corresponding to the concepts and actions illustrated in summary form by Figure E-1 are provided in the following subsections. These examples aim to offer practical illustrations of how each element contributes to the overall implementation of ODRM practices and facilitate a clearer understanding of the concepts and actions outlined in the diagram.

E.1 Assertion of Risk Leadership

The application of *risk leadership* within an organization is predicated upon the directives outlined in NPD 1000.0, which addresses the foundations of effective risk management practices. In practical terms, the application of *risk leadership* requires a clear and consistently-shared

identification and communication of high-priority *activity/project objectives* (A/P-Os) and the associated *risk posture* by leaders, managers, and execution staff involved in an Activity/Project (A/P).

As an example, consider a "Mission X," defined as a fast-track demonstration project aimed at utilizing a new type of radioisotope power source in planetary missions. In this example, we assume that NASA leadership envisions this project to be a "demo mission" with a short design and procurement cycle, a moderate level of cost, and full compliance with the stringent nuclear safety requirements for missions involving radioactive material in space. It is therefore assumed in our example that these specific goals and constraints are explicitly identified and communicated as top-level A/P-Os that should be reflected in the adopted A/P's *risk posture*, as summarized below:

- The primary Technical A/P-O is to demonstrate the viability of the new technology in a mission context
- A key Programmatic A/P-Os is to achieve a fast schedule and manage the project within a moderate budget.
- The main Safety A/P-O is to ensure strict compliance with safety requirements applicable to a nuclear-powered mission.

The responsibility of asserting *risk leadership* lies with managers at all levels of the organizational hierarchy, however it is the highest-level managers who are responsible for providing clear directives for its definition and application throughout the organization. By clearly articulating the A/P-Os and associated *risk posture*, leaders can guide and align the efforts of the teams involved in the project, enabling a coordinated and effective approach to ODRM. In summary the following can be asserted as the *risk leadership* cardinal principles:

- *Risk leadership* is established based on the directives of NPD 1000.0 as a crucial element for successful ODRM
- It involves the identification and communication of high-priority A/P-Os and the associated risk posture.

The example of "Mission X" demonstrates how these principles are applied in the context of a fast-track demonstration project, with specific A/P-Os related to technology demonstration, project schedule and budget, and safety compliance.

E.2 Definition of Activity Risk Posture

Defining and explicitly declaring a project or activity *risk posture* is a critical aspect of affirming and communicating risk leadership by activity leaders. The declaration of risk posture should specifically reference the key Activity/Project Objectives (A/P-Os) that have been established and to which it therefore applies. In the Mission X example, let's consider the *risk posture* choices that we can assume to have been made by the Executive Leaders for that specific project:

- In this project, the Technology Demonstration A/P-O is assigned a relatively high risk tolerance. The rationale behind this decision is based on the understanding that the mission

is primarily focused on research and development, where even partial success holds significant value. Therefore, the project leaders acknowledge that taking on a higher level of risk is acceptable in pursuit of the technology demonstration objective.

- On the other hand, the Schedule and Budget A/P-Os are assigned a medium to low risk tolerance. This decision is justified by the fact that the project operates under a paradigm of fast development and relatively modest project execution budgets. To ensure that the project remains on track and within budgetary constraints, a moderate to low level of risk tolerance is deemed appropriate for these objectives.
- Finally, any non-compliances with Public Safety A/P-Os are assigned a very low risk tolerance. This decision is based on the understanding that missions involving nuclear power sources are subject to special safety requirements. Any accident or failure to comply with safety protocols could result in radiological release, potentially impacting the public. Recognizing the severe consequences and political repercussions of such an event, the project leaders prioritize maintaining a very low risk tolerance for public safety-related objectives.

By explicitly defining and declaring the *risk posture* in relation to the key A/P-Os, activity leaders provide clear guidance to the project team and address stakeholders' risk concerns. This enables consistent ODRM-informed decision-making and risk management throughout the project life cycle, aligning with the organization's *risk posture* and strategic objectives. Openly communicating the rationale behind the *risk posture* choices helps foster a shared understanding of it, and a commitment to the application of consistent and balanced ODRM practices within the project or activity.

E.3 Identification of Risk Tolerance Levels

Risk Management analysts and specialists working within a specific activity or project play a crucial role in translating the *risk posture* directives received from project leaders into operational definitions of *risk tolerance levels* (RTLs) for each declared objective. This process involves two interconnected processes that may require iterative feedback and collaboration between the analysts and project leaders to reach a full convergence and consensus:

- To begin, the analysts identify and define *performance measures* (PMs) that can effectively capture the degree of success achieved for each (A/P-O). These PMs serve as quantifiable indicators that reflect the progress and achievement of the objectives.
- Additionally, *performance markers* (PMKs) are established to represent target levels of minimum performance within the scales of the PMs. These PMKs provide specific benchmarks against which the actual performance can be compared.
- Finally *risk tolerance level* values are defined and established in regard to the achievement of A/P-Os as expressed by the performance target represented by the PMKs.

According to the guidelines provided by NPR 8000.4, ***RTLs are set as threshold values for the maximum acceptable probability of missing each pre-defined PMK***. In other words, RTLs define the level of risk that the organization is willing to tolerate regarding the achievement or not of each

objective. The determination of RTLs is itself risk-informed – it is based on preliminary assessments of risk made at the outset of the activity or project. This allows for the establishment of feasible RTLs that the activity or project can credibly be expected to meet.

It is noted process of defining RTLs for each PMK may involve iterations and discussions between the analysts and project leaders. The analysts propose initial definitions based on their expertise and understanding of the project context, while the project leaders provide feedback and input to ensure alignment with overall project goals and *risk posture*. This iterative process allows for a collaborative approach, where both the managerial and technical branches of a project contribute their perspectives and insights to arrive at RTLs that are appropriate and acceptable to all stakeholders.

By establishing operational definitions of RTLs, analysts and specialists enable the project team to effectively manage and monitor risks in alignment with the organization's objectives. This systematic approach ensures that risks are assessed and managed based on specific performance indicators and associated performance targets, providing a structured framework for decision-making and risk mitigation throughout the activity or project life cycle.

When defining RTLs, the following considerations come into play within the context of a specific Activity/Project (A/P):

Possible existence of multiple PMKs for a given PM:

In many cases, two types of Performance Markers (PMKs) are established for a given A/P and Performance Measure (PM). The first is the "Required PMK" (PMK-R), which represents the minimum performance level set by the *Acquirer* that the Provider must deliver. The second is the "Goal PMK" (PMK-G), which is more flexibly negotiated between the Provider and *Acquirer* and sets a target performance level beyond the PMK-R. Additionally, Risk Analysts may define additional PMKs as "watch-margins" to provide an added safeguard against potential deviations from the primary PMK-R and/or PMK-G targets.

RTL for Aggregate vs. Individual Risk Scenarios:

It is generally advisable to define distinct RTLs for Aggregate Risk and Individual Risk Scenarios. An Aggregate Risk RTL (ARTL) should be identified first for each given A/P-O and its associated PM(s)²². The ARTL should reflect the *risk posture* directives concerning the achievement of the identified targets (e.g., the PMK-R and PMK-G markers) in the corresponding PM dimension(s), in consideration of the potential risks impacting that particular performance dimension. Once an ARTL is established, a corresponding Individual Risk Scenario RTL (IRTL) value can also be identified and set for the individual risk scenarios that may exist within the overall risk spectrum. The setting of such IRTL value may reasonably be based on a heuristic preliminary estimate of the number of significant Individual Risk Scenarios which might be contributing to the Aggregate

²² It is assumed here that a single PM can be identified as the metric for achievement of a specific A/P-O. If, however, an A/P-O is "measured" by multiple PMs, the corresponding multiple risk dimensions and the overall risk level resulting from contributions in each of those will need to be considered. In some other cases a composite PM may alternatively be formulated, e.g., via the use of utility theory, and risk may be evaluated in terms of such a composite metric.

Risk, and of the magnitude of uncertainties and unknown risks affecting the overall Aggregate Risk for the specific PM under consideration.

Multiple RTL Values for a Single PM:

For each PM, separate RTL values should be set for PMK-R and PMK-G. PMK-G should be set at a higher (in terms of desirability) value of performance than PMK-R, and the RTL value for PMK-R should be set at a lower risk tolerance (i.e., a lower RTL value) compared to PMK-G. This distinction recognizes the need for stricter risk management when it comes to meeting minimum performance requirements versus achieving higher performance targets.

Multiple RTLs for a Single PMK:

As mentioned in Section C.6.1, there might be situations where just one PMK is set in a given performance dimension and in relation to an A/P-O, and different RTL values are set to define limits of acceptable probability (i.e., performance risk) that the magnitude of a possible PM shortfall with respect to the PMK be greater than certain acceptable or tolerable values. Examples of this type of RTL and risk classification settings are discussed in Chapter 5.

RTLs for Margin PMKs:

If additional "Margin PMKs" (PMK-Ms), such as "PMK-R Margin" (PMR-M) and "PMK-G Margin" (PMG-M), are defined, any associated RTL values should follow a specific order, i.e., they should increase in a stepwise manner from PMK-R to PMR-M, from PMR-M to PMK-G, and from PMK-G to PMG-M, reflecting the logic by which levels of tolerance should reasonably decrease as one is moving from consideration of the violation of "alarm levels" – the "margins" – versus the violation of an actual "goal," or finally of a requirement or constraint.

By adhering to these principles, organizations can establish a consistent set of RTLs that covers all PM dimensions in alignment with the A/P-Os and the associated *risk posture*, and provides a framework for assessing and classifying risks at both the individual and aggregate levels.

To illustrate the above concepts, consider again the example of the Tech Demo A/P, where one of the programmatic objectives is to keep the overall cost at a predefined moderate level. The Performance Measure (PM) that represents the achievement of this objective is the "overall cost" of the A/P. For the purposes of the example, it is assumed that the *Acquirer* has set the "maximum allowable budget" for the Tech Demo A/P at \$300M, which serves as the Required PMK (i.e., is the PMK-R) for the Cost PM.

It is also assumed that, in addition to the PMK-R, the *Acquirer* and Provider have mutually agreed upon an overall cost/price goal of \$275M, which represents the Goal PMK (i.e., the PMK-G) for the Cost PM. Considering the A/P Leadership's indication of a "Medium-to-Low Risk Tolerance" for the A/P Cost PM, it can be thus assumed that the Risk Analysts might propose the following values for the marker ARTLs:

- Aggregate Risk Tolerance Level of 10% for the PMK-R (ARTL-R)
- Aggregate Risk Tolerance Level of 20% for the PMK-G (ARTL-G).

With the ARTLs established, the next step is to define the Individual Risk Scenario Tolerance

Levels (IRTLs) for the Cost PM.

A heuristic derivation of the IRTL value is obtained accounting for the combined effect of two basic elements:

- a. the possible presence of *U/U risk (UUR)*, which leading indicators suggest could be as high as X% of the aggregate of significant identifiable and known risks, based on consideration of the novelty and complexity factors relative to the radioisotope power generator technology the project seeks to demonstrate;
- b. the estimated maximum number, N, of significant Individual Cost Risks (ICRs) that may concurrently contribute to the Known Aggregate Cost Risk (KACR).

Considering these two elements, and the ARTL value set at the tolerance limit for the Total Cost Risk (TCR), inclusive of U/U risk, the following formulas may be used to heuristically obtain an IRTL value compatible with the established ARTL:

$$\begin{aligned} \mathbf{TCR \leq ARTL} , \quad & \text{where} \quad TCR = KACR + UUR, \\ & UUR \approx X\% \cdot KACR, \\ & \text{and} \quad KACR \approx N \cdot \text{Max(ICRs)} \end{aligned}$$

From the above one obtains:

$$\text{Max(ICRs)} \leq \mathbf{ARTL / [(100 + X)\% N]}$$

Thus, remembering that it must be by definition:

$$\mathbf{Max(ICRs) \leq IRTL} ,$$

if one chooses for IRTL a value that satisfies:

$$\mathbf{IRTL \leq ARTL / [(100 + X)\% N]} ,$$

and the above risk tolerance condition in bold cursive for individual risk scenarios is satisfied, then there is reasonable assurance that the tolerance condition for TCR can be satisfied.

Thus, if in the demo Mission X example $N = 3$ and $X = 50\%$, an IRTL-R value consistent with the previously established ARTL-R value of 10% would be calculated as $\mathbf{IRTL \leq 10\% / 4.5}$, resulting in an IRTL value of approximately 2.22%. Similarly, the IRTL-G would be $\mathbf{IRTL-G \leq ARTL-G / 4.5}$, suggesting an IRTL-G value of approximately 4.44%.

By establishing IRTL values in this fashion, individual risk scenarios associated with the Cost PM can be monitored and managed within tolerances that are consistent both with the tolerance established for *Total Aggregate Risk* and with the estimated impact of *U/U risk*. An IRTL setting established in this fashion thus provides a framework to assess and control individual risk scenarios at a desired level, consistently within the overall risk management strategy for the Tech Demo A/P.

E.4 RTL Approval and Determination of Risk Acceptability

The process of RTL determination can be based on “technical” considerations of the type discussed in the preceding section, but the final setting of the RTL values may involve iterations and what-

if thought processes involving both the risk analysts and the A/P decision makers. In some cases, only after an initial risk profile concerning an A/P-O and its PM(s) has been identified by the technical analysts and submitted to the evaluation of the A/P decision makers the initial selection of RTL values can be approved and confirmed.

Individual Risk Scenario (IRS) Classification

Whether at any specific time the definition of RTLs can be considered tentative or definitive, the process of determining acceptability for Individual Risk Scenarios (IRSs) -- purely on a technical basis that considers compliance or not with the risk tolerance criteria established by the setting of RTL values -- involves a systematic assessment carried out by the Risk Analysts. For each combination of Performance Measure (PM) and Individual Risk Scenario (IRS), the analysts must in fact:

- a. Evaluate the probability that the PM will exceed or fall below any PMK values of concern, such as the Required PMK (PMK-R), PMK-R Margin (PMR-M), Goal PMK (PMK-G), and PMK-G Margin (PMG-M), as a result of the effect of the IRS to be classified;
- b. Compare the value of such probability with the corresponding, pre-established RTL value.

This type of assessment is focused on the specific IRS's impact on the PM. To determine the PM-relevant probability values, there are two approaches that can be employed. The first approach involves deriving the probability distribution functions for the PM, preferably in cumulative distribution function (CDF) or complementary cumulative distribution function (CCDF) form. These functions provide direct information about the probability of non-exceedance or exceedance of any PM values of interest.

Alternatively, the assessment can directly address the question of whether the PM outcomes at the PMK values have a probability of exceedance or non-exceedance greater or smaller than the previously established Individual Risk Scenario Tolerance Levels (IRTLs) for those PMKs. This approach represents a more focused analysis, which may be based in part on qualitative considerations and expert judgment when detailed distribution functions are not readily available.

Once an assessment has been carried out, its results can be summarized with the aid of a decision table that considers the probabilities of exceedance or non-exceedance of the RTLs used in the risk classification, as shown in Table E-I for the more complex case where four different types of PMKs ("requirement," "goal," and respective "margins") may have been set.

While the final decision on risk acceptability rests with the accountable A/P decision-maker(s), a standardized pre-determination can be made based on the classification scheme outlined by the table, which can then be used as a decision aid by the Activity/Project (A/P) decision-makers responsible the determination of risk classification. The classification scheme typically includes three categories: "Acceptable," "Marginal," and "Non-acceptable," reflecting the level of risk associated with the IRS, as represented by the satisfaction or non-satisfaction of the established RTL criteria. This scheme provides a consistent framework for assessing and classifying risks, facilitating the decision-making process and ensuring a common understanding of risk acceptability across the organization.

Table E-I. Decision Table for Individual Risk Scenario Classification

IRTL-R	IRTL-RM	IRTL-G	IRTL-GM	Risk Classification
<i>Satisfied</i>	<i>Satisfied</i>	<i>Satisfied</i>	<i>Satisfied</i>	Acceptable
<i>Satisfied</i>	<i>Satisfied</i>	<i>Satisfied</i>	<i>Not-satisfied</i>	Acceptable (watch for Marginal)
<i>Satisfied</i>	<i>Satisfied</i>	<i>Not-satisfied</i>	<i>"Don't care"</i>	Marginal
<i>Satisfied</i>	<i>Not-satisfied</i>	<i>"Don't care"</i>	<i>"Don't care"</i>	Marginal (watch for Not-acceptable)
<i>Not-satisfied</i>	<i>"Don't care"</i>	<i>"Don't care"</i>	<i>"Don't care"</i>	Not-acceptable

Aggregate Risk (AR) Classification

The process for determining risk acceptability for Aggregate Risk (AR) follows an approach similar to the classification of Individual Risk Scenarios (IRSs). However, there are some key differences in the assessment criteria and classification process. Unlike the IRS classification process, which focuses on the risk effect of each individual IRS on the PM outcome distribution, the assessment of Aggregate Risk considers the cumulative impact on the PM of all significant known IRSs, and of *U/U risk*. The key criterion for risk classification in the case of Aggregate Risk is the satisfaction or non-satisfaction of the Aggregate Risk Tolerance Level (ARTL), as opposed to IRTL, thresholds, and, while in the case of IRS classification the probability to be compared with IRTL values is estimated by considering the effect of single IRSs, taken one at a time, on the PM of concern, the classification of TAR considers the PM outcome distribution estimated with inclusion of all known risk contributions and *U/U risk* projections.

Similarly to the case of IRS classification, the risk classification for Aggregate Risk is made by comparing the probabilities of PM outcomes at the PMK values with the ARTL thresholds. If these probabilities fall within the acceptable range defined by the ARTLs, the risk is classified as "Acceptable." If the probabilities are such that RTLs for requirements (PMK-Rs) are satisfied, but goals (PMK-Gs) are not, the risk is classified as "Marginal" – i.e., risk that may require further analysis or mitigation. If neither PMK-R nor PMK-G RTLs are satisfied the risk is classified as "Non-acceptable" – i.e., some type of risk control measure must be identified and implemented to permit continued A/P progression towards a successful completion and mission-execution. A decision table may again be used to aid the risk classification process for different PM-AR combinations, as illustrated in Table E-II.

Use of such a table, together with the corresponding one for IRS classification, helps the achievement of consistency in risk assessment and classification across the activity or project, and therefore also in the associated decision processes.

Further discussion of the above concepts via demo Mission X examples is provided in the following. In the first example, which illustrates the classification of an Individual Risk Scenario (IRS) threatening the Mission X cost objective, the risk arises from the non-negligible probability of incurring additional expenses for the production of the radioisotope required to fuel the demo mission nuclear power source. To determine the classification of this risk scenario, "IRS1," the Cost Performance Measure (PM) Cumulative Distribution Function (CDF) is assessed considering

just the IRS1 effect on the overall cost of the activity or project, as shown in Figure E-2.

Table E-II. Decision Table for Aggregate Risk Classification

ARTL-R	ARTL-RM	ARTL-G	ARTL-GM	Risk Classification
<i>Satisfied</i>	<i>Satisfied</i>	<i>Satisfied</i>	<i>Satisfied</i>	Acceptable
<i>Satisfied</i>	<i>Satisfied</i>	<i>Satisfied</i>	<i>Not-satisfied</i>	Acceptable (watch for Marginal)
<i>Satisfied</i>	<i>Satisfied</i>	<i>Not-satisfied</i>	<i>"Don't care"</i>	Marginal
<i>Satisfied</i>	<i>Not-satisfied</i>	<i>"Don't care"</i>	<i>"Don't care"</i>	Marginal (watch for Not-acceptable)
<i>Not-satisfied</i>	<i>"Don't care"</i>	<i>"Don't care"</i>	<i>"Don't care"</i>	Not-acceptable




Figure E-2. Example of PM CDF Assessment for “Acceptable” IRS1 Classification

The figure visually represents how IRS1 affects the PM CDF outcomes. However, it must be emphasized that the classification process only requires comparing the probability values (u, v, x, and y) at the PMK points with the corresponding Individual Risk Scenario Tolerance Level (IRTL) values (IRTL-R, IRTL-RM, IRTL-G, and IRTL-GM), thus a simplified assessment could be limited to the estimation of just those four probability values, and, if the number of PMKs were to be limited to just one requirement and one goal, or even altogether to just one single marker, the risk assessment for classification purposes would be even further simplified. In the figure, the element of probability that drives the risk classification as being “Acceptable,” although with need to guard against slipping into “Marginal,” is highlighted as being “y” (the CDF values at the PGM - Performance Goal Margin), which is the only probability value non-complying with the RTL threshold criteria. As previously discussed, the classification of IRS1 is determined by examining the PM CDF and comparing probabilities at the PMK values with the pre-set RTL values. The

resulting classification of “Acceptable” as per the decision rules previously shown in Table E-I is illustrated in Table E-III.

Table E-III. Classification of Example Risk Scenario IRS1

IRTL-R	IRTL-RM	IRTL-G	IRTL-GM	Risk Classification
<i>Satisfied</i>	<i>Satisfied</i>	<i>Satisfied</i>	<i>Not-satisfied</i>	Acceptable (watch for Marginal)

The next example illustrates the classification of another Individual Risk Scenario (IRS), referred to as IRS2. The Cost PM CDF is now therefore re-drawn considering the effect of IRS2 on the overall cost of the activity or project, as visually represented by Figure E-3. Similarly to the previous example, the classification of IRS2 is determined by comparing the probability values (u, v, x, and y) at the PMK points with the corresponding Individual Risk Scenario Tolerance Level (IRTL) values (IRTL-R, IRTL-RM, IRTL-G, and IRTL-GM). The classification of IRS2 is determined in the present case by the probability value, “x,” at the PMK-G marker, which exceeds the PMK-G RTL, RTL-G. Thus, IRS2 is classified as "Marginal," as shown in Table E-IV per the criteria of Table E-I. This indicates that IRS2 falls within a range of risk that calls for careful monitoring and management.




Figure E-3. Example of PM CDF Assessment for “Marginal” IRS2 Classification

Table E-IV. Classification of Example Risk Scenario IRS2 Classification

IRTL-R	IRTL-RM	IRTL-G	IRTL-GM	Risk Classification
<i>Satisfied</i>	<i>Satisfied</i>	<i>Not-satisfied</i>	<i>"Don't care"</i>	Marginal

In the next and final example, the focus is on the classification of yet another IRS, referred to as “IRS3.” Again, the Cost PM CDF distribution is now reassessed to specifically identify the effect of IRS3 on the overall cost of the activity or project, which is assumed to be as shown by Figure E-4. The resulting IRS3 classification, illustrated by Table E-V, is driven by the probability “u,” at the PMK-R marker, which is greater than the PMK-R RTL, RTL-R. This makes IRS3 “Not-Acceptable” per the criteria of Table E-I. This indicates that IRS3 risk exceeds the acceptable or tolerable thresholds of severity and requires immediate attention, for appropriate identification and application of mitigation strategies.




Figure E-4. Example of PM CDF Assessment for “Not-Acceptable” IRS3 Classification

Table E-V. Classification of Example Risk Scenario IRS3 Classification

IRTL-R	IRTL-RM	IRTL-G	IRTL-GM	Risk Classification
Not-satisfied	"Don't care"	"Don't care"	"Don't care"	Not-acceptable

The examples provided above illustrate the practical application of the classification process for Individual Risk Scenarios (IRSs), utilizing Cumulative Distribution Function (CDFs) representation of the effect of IRSs on the relevant Performance Measures (PMs). By comparing the IRS-induced probabilities at the PMK values with the corresponding Individual Risk Scenario Tolerance Level (IRTL) values, decision-makers can effectively assess the acceptability of such risks within an A/P, in alignment with the assigned and accepted A/P risk posture. The classification outcomes help inform risk management strategies and ensure that appropriate actions are taken to mitigate or address risks based on their severity and impact on the desired objectives.

E.5 Display and Communication of Risk Classification Information

It can be said that in practical and synthetic terms the ultimate focus of the classification of risks is to determine to which of the three basic categories "**Acceptable**," "**Marginal**," or "**Non-acceptable**" the IRSs of concern and the TAR should be assigned, so that appropriate ODRM actions may be applied to such risks. By assigning risks to these categories, decision-makers can prioritize their attention and allocate resources accordingly. In its conciseness and simplicity, the three-level classification allows for a clear differentiation between risks that are within pre-established tolerance levels (*acceptable risk*), risks that require careful monitoring and management (*marginal risks*), and risks that pose significant threats and demand immediate action (*non-acceptable risks*).

It follows from the above that the three-level categorization is also the most essential portion of risk information that needs to be displayed and communicated effectively for an understanding and overall perspective on the risk status of an activity or project. For Mission X example case and the related three hypothetical IRSs discussed in the preceding section, Figure E-5 shows the simple "risk-bar" format of display that communicates this "bottom-line" information.

Mission X Classification Info	
IRS3	"High" Probability of PM Shortfall w respect to PMK-R $\equiv$ Unacceptable Risk
IRS2	"High" Probability of PM Shortfall w respect to PMK-G $\equiv$ Marginal Risk
IRS1	"Low" Probability of PM Shortfall $\equiv$ Acceptable Risk

Figure E-5. Risk-Bar Display of Mission X IRS Classifications

In general, decision-makers and other project stakeholders may desire a more detailed understanding of specific risks or of the overall aggregate risk. To communicate risk at a more complete level, a risk distribution, such as the ones depicted in the preceding charts, is a more appropriate and comprehensive form of risk display and communication. A risk distribution showcases the probability of various PM outcomes (or events) occurring, providing a more nuanced and complete representation of the risk landscape, together with the additional connection links to even more in-depth risk information – e.g., the data and analyses by which the PM distribution derivation has been obtained.

By access to risk distribution and underlying information displays, decision-makers can develop a perspective on the likelihood and potential impact of different risk scenarios, which in some cases may be necessary to enable a more informed decision-making process, allowing for the identification of critical areas of concern in the risk assessment process itself, the evaluation of risk trade-offs, and the development of targeted risk mitigation strategies. Figure E-6 concludes the present recap and example thread with the illustration of a possible combined "CDF + Task-Bar" display format of the Mission X IRS set discussed earlier. For the sake of clarity, the CDF

side of the display is simplified in this example figure by showing only the PMK-R and PMK-G markers and corresponding RTLs.




Figure E-6. Example of Combined Risk-Bar / PM-CDF Display of IRS Classification

Appendix F Content Guide for the Technical Basis for Deliberation

F.1 Technical Basis for Deliberation Content

The Technical Basis for Deliberation (TBfD) document is the foundation document for the risk-informing activities conducted during Part 1 and Part 2 of the RIDM Process. The TBfD conveys information on the performance measures and associated imposed constraints (including proposed PMKs) for the analyzed decision alternatives.

Because the TBfD provides the specific risk information to understand the uncertainty associated with each alternative, this document serves as the technical basis for risk-informed selection of alternatives within the activity. The risk analysis team, working under the overall activity management guidance, develops TBfD documentation and updates the information provided as necessary based upon questions and/or concerns of stakeholders during deliberation. The risk analysis team works with the deliberators and decision-maker to support deliberation and alternative selection.

Depending on the decision context, the TBfD can be graded. For key direction-setting decisions, a fully developed formal TBfD report may be appropriate. For more minor decisions the information needed to adequately inform the decision may be less. In general, the formality of the TBfD will correlate with the rigor of the analysis it documents. In all cases, the information documented in the TBfD should support robust decision-making insofar as is practicable.

The TBfD includes the following general sections:

- Technical Summary: This section describes the problem to be solved by this effort and each of the general contexts of each of the alternatives.
- Top-level Requirements and Expectations: This section contains the top-level requirements and expectations identified in Step 1 of the RIDM process. In cases involving diverse stakeholders, a cross reference between expectations and stakeholder may be presented.
- Derivation of Performance Measures: This section shows the derivation of performance measures for the decision conducted in Step 1 of the RIDM process. Typical products are the objectives hierarchy and a table mapping the performance objectives to the performance measures. When proxy performance measures are used, their definitions are provided along with the rationale for their appropriateness. When constructed scales are used, the scales are presented.
- Decision Alternatives: This section shows the compilation of feasible decision alternatives conducted in Step 2 of the RIDM process. Typical products are trade trees, including discussion of tree scope and rationales for the pruning of alternatives prior to risk analysis. Alternatives that are retained for risk analysis are described. This section also identifies any imposed constraints on the allowable performance measure values, and a map to the originating top-level requirements and/or expectations.
- Risk Analysis Framework and Methods: This section presents the overall risk analysis framework and methods that are set in Step 3 of the RIDM process. For each analyzed

alternative, it shows how discipline-specific models are integrated into an analysis process that preserves correlations among performance measures. Discipline-specific analysis models are identified and rationale for their selection is given. Performance measures are identified for each alternative.

- Risk Analysis Results: This section presents the risk analysis results that are quantified in Step 4 of the RIDM process.
 - Scenario descriptions: For each alternative, the main scenarios identified by the risk analysis are presented.
 - Performance measure pdfs: For each alternative, the marginal performance measure pdfs are presented, along with a discussion of any significant correlation between pdfs.
 - Imposed constraint risk: For each alternative, the risk with respect to imposed constraints is presented, along with a discussion of the significant drivers contributing to that risk.
 - Adherence to the Risk Posture: For decisions made under an established risk posture, the risk with respect to each PMK is presented and compared to the associated RTL. Where the risk exceeds the RTL, a discussion of the significant drivers contributing to that exceedance is presented.
 - Supporting analyses: For each alternative, uncertainty analyses and sensitivity studies are summarized.
- Risk Analysis Credibility Assessment: This section presents the credibility assessment performed in accordance with [1].

F.2 Reference for Appendix F

1. NASA Standard. NASA-STD-7009A w/ Change 1, Standard for Models and Simulations. December 2016.

Appendix G Content Guide for the Risk-Informed Selection Report

G.1 Risk-Informed Selection Report Content

The Risk-Informed Selection Report (RISR) documents the rationale for selection of the selected alternative and demonstrates that the selection is risk-informed. The decision-maker, working with the deliberators and risk analysis team, develops the RISR.

Depending on the decision context, the RISR can be graded. For key direction-setting decisions, a fully developed formal RISR may be appropriate. For more minor decisions a decision memo may be sufficient. In all cases, the RISR should clearly communicate and defend the decision rationale.

The RISR includes the following general sections:

- Executive Summary: This summary describes the problem to be solved by this effort and each of the general contexts of each of the alternatives. It identifies the organizations and individuals involved in the decision-making process and summarizes the process itself, including any intermediate downselects. It presents the selected alternative and summarizes the basis for its selection.
- Technical Basis for Deliberation: This section contains material from the TBfD (see Appendix F).
- Performance Targets: This section presents the performance measure ordering and risk tolerances used to develop the risk-normalized performance targets (RPTs) during Step 5 of the RIDM process, with accompanying rationale. It tabulates the resultant performance targets for each alternative.
- Deliberation: This section documents the issues that were deliberated during Step 6 of the RIDM process.
 - Organization of the deliberations: The deliberation and decision-making structure is summarized, including any downselect decisions and proxy decision-makers.
 - Identification of the contending decision alternatives: The contending alternatives are identified, and rationales given for their downselection relative to the pruned alternatives. Dissenting opinions are also included.
 - Pros and cons of each contending alternative: For each contending alternative, its pros and cons are presented, along with relevant deliberation issues including dissenting opinions. This includes identifying violations of significant engineering standards, and the extent to which their intents are met by other means.
 - Deliberation summary material: Briefing material, etc., from the deliberators and/or risk analysts to the decision-maker (or decision-makers, in the case of multiple downselects) is presented.

- Alternative Selection: This section documents the selection of an alternative conducted in Step 6 of the RIDM process.
 - Selected alternative: The selected alternative is identified, along with a summary of the rationale for its selection.
 - Performance targets: The finalized performance targets for the selected alternative are presented, along with the final performance measure risk tolerances and performance measure ordering used to derive them.
 - Adherence to the Risk Posture: For decisions made under an established risk posture, the risk of the selected alternative with respect to each PMK and RTL is presented. Where the risk exceeds the RTL, a discussion of the significant drivers contributing to that exceedance is presented.
 - Risk list: The RIDM risk list for the selected alternative is presented, indicating the risk-significant conditions extant at the time of the analysis, and the assessed impact on the ability to meet the performance targets.
 - Decision robustness: An assessment of the robustness of the decision is presented.

Appendix H Estimate the Performance Measure Margins Needed to Accommodate Implied Risk Scenarios (The U/U Risk)

Graded Approach Guidance			
Activity Class A+	Activity Class B	Activity Class C	Activity Class D
Historical experience is used to develop correlations between pertinent leading indicators and performance measures. Attention is given to all identifiable leading indicators that relate to the potential for U/U risk. Estimates of U/U risk are obtained by comparing the historical (total) risk to the aggregate risk from the known individual risk scenarios, at their respective risk tolerance levels.	Same as for Activity Class A+, except reduced coverage is acceptable based on experience and expert judgement to deemphasize areas known historically to be not significant for achieving safety, mission success, programmatic requirements, or other requirements important to the activity being pursued.	Same as for Activity Class A+, except coverage focuses on readily evaluated performance measures and leading indicators that historically are significant to safety, mission success, programmatic requirements, or other requirements important to the activity being pursued.	Same as for Activity Class A+, except coverage focuses on safety and major equipment damage caused by interactions with interfacing systems of higher value.

The rationale and general framework for determining margins to account for U/U risk was presented in Sections 3.2.3, 3.2.6, 3.3.3 and 3.3.4.

As discussed earlier (see Section 3.2.4.1), *leading indicators* that characterize degrees of design or process complexity, design or process novelty, schedule pressure, nonadherence to quality principles, deficiencies in management culture, lack of defense in depth, lack of human factors consideration, national or regional economic conditions, and national or regional political trends are of interest here because they correlate with U/U risk. The challenge is to deduce which of these types of leading indicators are of most significance for the activity in question. This is done mainly by researching historical data and eliciting expert judgment.

Other indicators of the health of an activity are characterized as *lagging indicators*. These include, for example, the record of milestone achievement, number of unresolved issues, rate of depletion of financial and schedule reserves, and results from audits. The preferred approach for handling lagging indicators is to treat them as conditions and include them in the context of defined scenarios. In this way, they contribute to the known risk rather than the U/U risk.

As mentioned in the above table, estimates of the magnitude of U/U risk for specific values of the relevant leading indicators are obtained by comparing performance markers (e.g., PMK-R, PMK-G) that are derived from historical experience (and thus include U/U risk) with those that are derived from explicit assessment of the known individual risk scenarios (and thus do not account for U/U risk). This can be done quantitatively or qualitatively. If there is a quantitative correlation,

such as the example shown in Figure 3-6, then one way to estimate the magnitude of U/U risk is illustrated in Figure H-1.




Figure H-1. Example Quantitative Derivation of a Performance Measure Margin to Account for U/U Risks Associated with a Particular Leading Indicator.

The U/U risk in this figure is taken to be the difference between two performance measure values: (1) $Performance_H$, which is a historically-based value using similarity analysis and accounting for the current value(s) of the leading indicator(s), and (2) $Performance_K$, an analysis-based value derived from the aggregation of the known individual risk scenarios for the current activity. As shown in Figure H-1, both $Performance_H$ and $Performance_K$ are set at the established RTL of their respective pdfs (see Section 3.3.1).

For example, suppose the performance measure is total program cost, the leading indicator is a composite index which includes complexity, newness, and time pressures, the current value of the leading indicator is 0.9 on a scale of zero to one, $Performance_K$ for program cost is estimated as \$2.0 B with a corresponding RTL-R level of 0.10, but $Performance_H$ is estimated as \$3.0 B based on past experience when the leading indicator composite index is 0.9. Then the estimated U/U risk is \$3.0 B minus \$2.0 B, or \$1.0 B.

A qualitative correlation might be one for which the leading indicator is a ranking of a combination of attributes expressed qualitatively. For example, in a retrospective look at a variety of catastrophic accidents both within the space programs and within other non-space programs [1], it was determined that the probability of such an accident occurring from unknown and/or underappreciated sources was highly dependent upon a collection of diverse leading indicators, most particularly including: (1) newness of the integrated system, (2) time pressure, (3) priority of

management objectives (e.g., cost compared to safety), (4) management communication structure (e.g., hierarchical vs. open-door), (5) newness of the technology, (6) extension of an existing technology to conditions beyond its design basis, and (7) the tightness of coupling between subsystems. Five categories were defined, each containing a qualitative description of these leading indicators. One category, for example, was defined as: “New systems that are developed or operated under significant time pressure, and with a design philosophy that involves either new technology or new integration of an existing technology or scaling of an existing technology beyond the domain of knowledge or tight coupling, but with reliability and safety having a higher priority than cost and schedule, and with an inclusive management structure.” The retrospective then proceeded to provide historically based estimates of the ratio of the U/U risk to the known risk for each of these categories. A schematic showing the nature of the results to be expected from such a qualitative assessment is provided in Figure H-2.




Figure H-2. Example Qualitative Derivation of U/U Risks Associated with a Composite Leading Indicator.

The ordinate in Figure H-2 is the ratio of U/U risk to known risk based on retrospective similarity analysis and accounting for the leading indicator values of the current activity. The abscissa is a discrete set of categories containing qualitative descriptions of leading indicators such as those cited in the preceding paragraph.

For example, suppose that for the leading indicator category described above, the ratio of the total probability of loss of crew, $P(\text{LOC})$, to the value attributable to known risks is expected to be 5.0. Also, suppose that the value of $P(\text{LOC})$ predicted based on known risks is 0.01. Then the expected total $P(\text{LOC})$ is $5.0 \times 0.01 = 0.05$, and the estimated U/U risk is $0.05 - 0.01 = 0.04$, or four times the value predicted from known risks.

H.1 References for Appendix H

1. Benjamin, A., Dezfuli, H., and Everett, C., Developing Probabilistic Safety Performance Margins for Unknown and Underappreciated Risks, *Journal of Reliability Engineering and System Safety*, Vol. 145 (329-340), January 2016.

Appendix I Develop Risk, Opportunity, and Leading Indicator Taxonomies

Graded Approach Guidance			
Activity Class A+	Activity Class B	Activity Class C	Activity Class D
Involves a comprehensive identification of categories for individual risk scenarios, opportunities, and leading indicators. Includes coverage across programmatic, engineering, institutional, and enterprise activity domains; across safety, technical, security, cost, and schedule execution domains; and within these categories across potential types of departure events, assets, and consequences.	Same as for Activity Class A+.	Same as for Activity Class A+.	Same as for Activity Class A+.

Taxonomies can be used for two primary purposes: 1) as a brainstorming tool for the comprehensive identification of individual risk scenarios, and 2) as a categorization tool for the identification of the cross-cutting nature of individual risk scenarios that are already identified.

A *taxonomy* is a tree structure of classifications that begins with a single, all-encompassing classification at the root of the tree, and partitions this classification into a number of sub-classifications at the nodes below the root. This process is repeated iteratively at each of the nodes, proceeding from the general to the specific, until a desired level of category specificity is reached.

Within the overall taxonomic framework, there may be taxonomy subclasses. For example, a departure event taxonomy, or departure taxonomy, focuses on categorizing the types of events that can initiate an undesirable scenario. A work breakdown taxonomy (or work breakdown structure taxonomy) concentrates on the categories of tasks whose successful completion might be threatened by a risk scenario; an asset taxonomy on the categories of system components, organizational entities, and portfolio items; a consequence taxonomy on the categories of loss that may be encountered by the affected component, entity, or item; and an objectives taxonomy on the categories of objectives that might be impacted. Each taxonomy focuses on a different individual risk scenario attribute.

As an example, Figure I-1 shows a representative objectives taxonomy. For illustration purposes, Figure I-1 includes a brief mention of the kind of individual risk scenario whose Affected Objective might be categorized as a member of the associated taxon. Figures I-2 through I-5, respectively, illustrate a representative departure taxonomy, work breakdown taxonomy, asset taxonomy, and asset-level consequence taxonomy.




Figure I-1. Example Affected Objective Taxonomy (Notional, Not Intended as Prescriptive)




Figure I-2. Example Departure Taxonomy (Notional, Not Intended as Prescriptive)




Figure I-3. Example Work Breakdown Element Taxonomy (Notional, Not Intended as Prescriptive)




Figure I-4. Example Asset Taxonomy (Notional, Not Intended as Prescriptive)




Figure I-5. Example Consequence (Notional, Not Intended as Prescriptive)

In order for taxonomies to be useful in identifying cross-cutting risks, taxonomies of a given type should be uniform across the scope of similar activities and/or organizational units, and should be an integral part of an organizational unit's risk database. The registering of an individual risk scenario into the risk database would then include the step of categorizing each individual risk scenario in terms of the relevant taxa of each taxonomy. This enables the individual risk scenarios in the database to be associated with one another in terms of their common attributes. Ideally, the risk databases should be linked or integrated in such a way that enables the identification of related individual risk scenarios across NASA. Such individual risk scenarios can then be assessed to determine whether or not they are in fact separately identified manifestations of a common cross-cutting risk. If so, the potential exists for them to be managed collectively, as described in Section 6.3.

Taxonomies for *leading indicators of potential risk* (a type of taxonomy not shown in the preceding figures) generally include attributes such as excessive complexity, novelty of a design, excessive schedule pressure, nonadherence to quality principles, deficiencies in management culture, lack of defense in depth, and lack of human factors consideration. These and other candidates for leading indicators were discussed in Appendix B.

Taxonomies are subject to modification over time as risks are identified that suggest revisions to the categories, further partitioning of categories, or the addition of new categories. Because they integrate elements of risk that cross organizational lines, they should ideally be maintained at high levels of the NASA organizational hierarchy, so that all organizational units work from common sets of taxonomies. Otherwise, if each organizational unit manages its own taxonomies they can diverge over time, reducing their ability to identify cross-cutting risks. Because the taxonomies are used to communicate risk characteristics throughout the relevant organizational entities, modifications to the taxonomies must be coordinated among the entities and kept uniform throughout.

