

The Fundamentals of Secure Aviation Design

A Guidebook for the Industrial Base



Version History

Date	Version Identifier	Version Notes
September 20, 2023	-	Initial Public Release of Part I – Foundations
December 24, 2024	0.1	Completed Guidebook Baseline ***Lockheed Martin Internal Use Only***
July 2, 2025	1.0	Complete Guidebook as submitted to the Lockheed Martin Public Release Workflow
July 22, 2025	1.11	Approved for Public Release PIRA #AER2025071030
August 19, 2025	1.12	Section 1.5 updated to reflect the deployed location of the FSAD portal and to add the public email address for feedback.

Photo Credit:

The cover art and derived use of that image in the graphics throughout this document are provided by Shutterstock under a Premier All Media license.

Table of Contents

A Message from the Project Sponsors	x
Acknowledgements	xi
Part 1 – Foundations.....	1
Chapter 1 Introduction.....	2
1.1 Cyber Resiliency	3
1.2 Cyber Resiliency and the Supply Chain.....	4
1.3 Using This Book.....	6
1.4 References and Additional Resources.....	7
1.5 The Lockheed Martin FSAD Portal	8
Chapter 2 The Threat is Real.....	9
2.1 Cyber Contested Operational Environments.....	9
2.2 The Aviation Mission Stack and Cyber Contested Environments	11
2.3 Threat Intelligence and Aviation Systems	12
2.4 Aviation Threat Actors.....	13
2.4.1 Advanced Persistent Threats.....	13
2.4.2 The Insider Threat.....	14
2.5 The Supply Chain Threat	16
2.6 Unique Challenges for Aviation Systems Cybersecurity	17
2.6.1 Scalability	17
2.6.2 Attribution.....	18
2.6.3 Espionage.....	18
2.6.4 Subversion and Sabotage	18
2.6.5 Battlespace Preparation.....	19
2.6.6 Accidents, Errors, and Attacks.....	19
2.7 Conceptualizing Vulnerabilities in Aviation Weapons Systems.....	19
2.7.1 Safety and Adversarial Cyber Effects.....	21
2.7.2 Understanding Cyber Threat Agents against Aircraft Systems	23
2.7.3 The Ever-Increasing Capabilities of Cyber Threat Agents	24
2.7.4 Emerging Aviation Vulnerabilities (Zero Days).....	25
2.8 Summary	25
2.9 Additional Reading	26
Chapter 3 Our Collective Cyber Responsibilities	28
3.1 A Positive Cyber Culture	28
3.2 Transitioning to a Positive Cyber Narrative.....	29
3.2.1 Cybersecurity is everybody’s job, including mine.	29
3.2.2 Cybersecurity is essential to everything we do.....	29
3.2.3 When we see something, we say something.....	30
3.2.4 How can we prevent that from happening here?	30
3.2.5 We don’t punish cyber incident or vulnerability disclosure.	31
3.3 Success Factors of a Healthy Aviation Cyber Culture.....	32
3.3.1 Building Cybersecurity In.....	32
3.3.2 Compliance and Security Control Requirements.....	32

3.3.3	Accountability and Responsibility	33
3.3.4	Management Commitment.....	33
3.3.5	Cyber Risk Assessment and Managing Cyber Risks	33
3.3.6	Vigilance of Potential Cyber Events / Continuous Monitoring	34
3.3.7	Incident Response	34
3.3.8	Security Education, Awareness, and Training.....	34
3.3.9	Continuous Improvement and After-Action Reviews	35
3.3.10	Cyber Considerations for Purchasing Equipment, Software, and Materials	35
3.4	Summary	35
Chapter 4	Securing the Development and Production Infrastructure	37
4.1	The Aviation Supply Chain and Critical Infrastructure	37
4.2	Risk Management and Critical Infrastructure	38
4.2.1	Identify	38
4.2.2	Protect	39
4.2.3	Detect	39
4.2.4	Respond.....	39
4.2.5	Recover	39
4.3	ICS and SCADA Systems	40
4.4	Enterprise IT and Cyber Hygiene.....	41
4.4.1	Practice Defense in Depth.....	41
4.4.2	Maintain the Inventory/Configuration Management	41
4.4.3	Run Antivirus.....	42
4.4.4	Apply Security Patches and Updates	43
4.4.5	Physical Access Control	43
4.4.6	Account Management	43
4.4.7	Authentication.....	44
4.4.8	Least Privilege	44
4.4.9	Separation of Duties.....	45
4.4.10	Protect Data.....	45
4.4.11	Automated Backup System.....	45
4.4.12	Monitor and Analyze Audit Logs	46
4.4.13	Use Firewalls and Network Segmentation.....	46
4.4.14	Include ICS and SCADA Systems.....	47
4.4.15	Minimize Attack Surface	47
4.4.16	BYOD/Mobile Device Strategy	48
4.4.17	Remote Access.....	48
4.4.18	Security Awareness Training.....	48
4.4.19	Disposal Procedures.....	49
4.5	DFARS	49
4.6	Air-Gapped Networks	50
4.7	Incident Response	50
4.8	Summary	51
Part 2	Secure Architecture and Design for Air Systems	52

Chapter 5	The Foundations of Secure Architecture and Design of Air Systems	53
5.1	Cyber Survivability, Cyber Resiliency, and Cybersecurity	53
5.2	The Confidentiality-Integrity-Availability (C-I-A) Triad	55
5.3	Intrinsic Cyber Resilience	55
5.4	Build Systems Better / Build Better Systems	56
5.5	Systems Engineering	57
5.6	Security Systems Engineering / Cyber Resiliency Systems Engineering	58
5.7	Lifecycles, Methodologies, and Model Based Systems Engineering	59
5.8	Airworthiness	60
5.9	System Assurance	61
5.10	Additional Reading and Resources	62
Chapter 6	Program Inception and Requirements.....	64
6.1	Cyber Resiliency at Program Inception	64
6.2	Technical Program Reviews.....	65
6.3	Systems Requirements Engineering.....	66
6.3.1	Requirements Documentation.....	66
6.3.2	Requirements Allocation, Decomposition, and Traceability	67
6.4	Elicitation of Cyber Resiliency Objectives and Requirements	68
6.5	Mission Impact Assessment	70
6.6	Use and Abuse Cases	71
6.7	Requirements for Security Sustainment.....	72
6.8	Requirements for Cyber Active Defense and Response	72
6.9	Cyber Resiliency Adjacent Requirements.....	73
6.10	Aviation Cyber Threat Intelligence.....	75
6.11	The Iterative Nature of Cyber Resiliency Requirements	76
6.12	Additional Reading and Resources	77
Chapter 7	Cyber Resilient Architecture and Design	79
7.1	Introduction to System Architecture and Design	79
7.2	Architectural Allocation of Cyber Resiliency Requirements.....	80
7.3	The Foundational Objectives of Cyber Resiliency	80
7.3.1	Prevent	81
7.3.2	Mitigate.....	82
7.3.3	Recover	82
7.3.4	Adapt.....	83
7.4	Additional Objectives of Cyber Resilient Architecture and Design	83
7.4.1	Minimizing Cyber Attack Surface	83
7.4.2	Design Anticipating Nation State Cyber Adversarial Activity	84
7.4.3	Implement Defense in Depth and Coordinated Protection	84
7.5	Techniques for Cyber Resilient Aviation System Architecture and Design.....	85
7.5.1	Techniques for Preventing Cyber Adversity.....	86
7.5.2	Techniques for Mitigating Cyber Adversity	92
7.5.3	Techniques for Recovering from Cyber Adversity.....	94
7.5.4	Techniques for Adapting Against Future Cyber Adversity	97

7.6	Pitfalls and Fallacies of Cyber Resilient Design for Aviation Systems	99
7.6.1	Assuming Aircraft are Protected By an Air Gap	99
7.6.2	Overvaluing Physical Access Controls	100
7.6.3	“Insider” Attacks Against Aviation Systems	100
7.6.4	The Paradox of “Hardened” Systems	101
7.6.5	Inherent Flaws/Zero Days	102
7.6.6	Security by Obscurity Doesn’t Work	102
7.7	Advanced Topics in Aviation Systems Cyber Resiliency	103
7.7.1	Exercising Operational Security for Aviation Systems	103
7.7.2	War Reserve and Combat Mode	103
7.7.3	Modular Open Systems Architecture (MOSA)	104
7.7.4	Vulnerability Management for Aviation Systems	105
7.7.5	Aviation Cyber Kill Chain	106
7.7.6	Zero Trust Concepts for Aviation Systems	108
7.7.7	Virtualization	109
7.7.8	Architecting and Designing for Worst-Case Survivability	109
7.8	Summary and Additional Resources	110
Chapter 8	Threat Modeling and Assessing Cyber Resiliency	112
8.1	Threat Modeling	112
8.2	Threat Modeling in the Systems Engineering Lifecycle	113
8.3	Trust Boundaries and Attack Surface In Threat Modeling	113
8.4	A Threat Relationship Model for Assets	114
8.5	Limitations of Threat Modeling	116
8.6	Threat Identification	116
8.6.1	STRIDE	116
8.6.2	Attack Trees	117
8.6.3	Attack Libraries	119
8.6.4	Criticality Analysis	119
8.7	Mitigating Threats	120
8.8	Integrating Threat Modeling with Risk Management	120
8.9	Alternative Threat Modeling Approaches	121
8.10	Summary and Additional Resources	122
Chapter 9	Cyber Risk Assessment	123
9.1	Characteristics of Effective Cyber Risk Assessment	123
9.2	Cyber Risk Assessment Process	125
9.2.1	Preparation for Cyber Risk Assessment	125
9.2.2	Conducting the Cyber Risk Assessment	126
9.2.3	Analyzing and Reporting Cyber Risk Assessment Results	126
9.2.4	Maintaining Cyber Risk Assessment Results	128
9.3	Cyber Risk Assessment Integrated With Systems Engineering	128
9.4	Selection of a Cyber Risk Assessment Methodology	128
9.4.1	The Case for Mission Based Cyber Risk Assessment	129
9.4.2	Cyber Table Tops (CTT)	129

9.4.3	Legacy CRA Methods.....	130
9.5	Cyber Risk Assessment is not a Precise Measurement of Risk	130
9.6	Cyber Risk Assessment and Broader Business Impacts	131
9.7	Cultural Attributes of Effective Cyber Risk Assessment.....	131
9.8	Summary and Resources	132
Chapter 10	Cyber Test and Evaluation	133
10.1	Integrating Cyber Test and Evaluation in the Systems Development Lifecycle	133
10.2	Cyber Test and Evaluation and Cyber Resiliency.....	135
10.3	Security Control Testing/Compliance.....	136
10.4	Blue Team Testing	136
10.5	Red Team and Penetration Testing	136
10.6	Cyber Verification and Validation	137
10.7	Summary and Additional Resources	137
Chapter 11	Securing the Complete Air System.....	138
11.1	Enterprise IT and Aviation Cybersecurity	138
11.2	Air System Components.....	139
11.2.1	Aircraft Boundary	139
11.2.2	Mission/Flight Planning Boundary	139
11.2.3	Maintenance System Boundaries.....	140
11.2.4	Logistics Systems Boundary.....	140
11.2.5	Training Boundaries.....	140
11.3	Summary and Additional Resources	141
Chapter 12	Assessment and Authorization of Aviation Systems.....	143
12.1	The Evolutionary and Interpretative Nature of A&A Standards.....	143
12.2	The Enterprise IT Foundations of RMF	144
12.3	RMF and the Systems Engineering Lifecycle.....	145
12.4	The Pitfalls of A&A Standards	147
12.5	Summary and Additional Resources	148
Part 3 – Cyber Resiliency Specialty Domains		149
Chapter 13	Software Assurance	150
13.1	Fundamentals of Secure Software Development	150
13.2	People in Software Assurance.....	151
13.2.1	Secure coding mindset/culture.....	152
13.2.2	Avoidance of Known Software Issues	152
13.2.3	Knowledge, Skills, and Attributes	153
13.2.4	Training the Software Development Role	154
13.3	Process in Software Assurance	155
13.3.1	Software Development Plan (SDP)	155
13.3.2	Software and Security Policies	156
13.3.3	Secure Coding Standards and Guides	157
13.3.4	Rigorous Code Reviews.....	158
13.3.5	Software Security Controls and Cybersecurity Models.....	159
13.4	Technology in Software Assurance.....	161

13.4.1	Static Code Analysis / Static Application Security Testing.....	161
13.4.2	Dynamic Code Analysis / Dynamic Application Security Testing	162
13.4.3	Automated Vulnerability Scanning.....	163
13.4.4	Security-Focused Software Testing	163
13.4.5	Automated Security Testing	164
13.5	Full Stack Cybersecurity for Aviation Systems	164
13.6	Pipelines: Pulling It All Together.....	166
13.7	Technology and Tool Evaluation and Selection	167
13.7.1	Selection of the Run-Time Environment	167
13.7.2	Selecting Programming Languages	168
13.7.3	Selection of Software Evaluation Tools	169
13.7.4	Configuration and Change Management	169
13.8	Specific Techniques for Cyber Resilient Software	170
13.8.1	Address Space Layout Randomization	170
13.8.2	Substantiated Integrity	171
13.8.3	Checksums and Vertical Parity Checks	171
13.8.4	Hashing	172
13.8.5	Input Validation	172
13.8.6	Bounds Checking	173
13.8.7	Malicious Commands and Messages	173
13.8.8	Software Implementation and Zero Trust Architecture	174
13.8.9	Validate and Control Software Loading	174
13.9	The Software Supply Chain and the Software Bill of Materials (SBOM).....	175
13.10	Summary and Additional Resources	175
Chapter 14	Hardware Assurance	178
14.1	Introduction to Hardware Assurance.....	178
14.2	The Engineering Trade-Offs of Hardware Assurance.....	179
14.3	Introduction to Hardware Components	180
14.3.1	Integrated Circuits (IC)	181
14.3.2	Printed Circuit Board (PCB).....	186
14.4	Common Hardware Attacks	187
14.4.1	Exploiting Debug Interfaces	188
14.4.2	Monitoring/Manipulating Exposed Connections	189
14.4.3	Data Extraction Flowchart	190
14.5	Advanced Hardware Attacks.....	191
14.5.1	Side-Channel Attacks.....	192
14.5.2	Fault Injection Attacks	193
14.5.3	IC Deconstruction Attacks	195
14.5.4	Hardware Implant Attacks	195
14.5.5	Hardware Trojan Attacks	196
14.6	Additional Methods of Hardware Assurance	197
14.6.1	Hardware Assurance and Supply Chain Risks.....	198
14.6.2	Mutual Authentication	200

14.6.3	Hardware Root of Trust	200
14.7	Summary and Additional Resources	201
Chapter 15	Cyber Supply Chain Management.....	203
15.1	Introduction	203
15.2	The Full Significance of the Aviation Supply Chain	205
15.3	Specific Concerns for DoD Aviation	208
15.4	Supply Chain Risk.....	210
15.5	Supply Chain Risk in the Program Protection Plan	210
15.6	Potential Sources of Supply Chain Risk	212
15.6.1	Data and Cyber Risk	212
15.6.2	Hardware Manipulation/Product Tampering Risk.....	213
15.6.3	Geopolitical Risk	213
15.6.4	Procurement Risk.....	214
15.6.5	Physical Risks in Transportation and Distribution	214
15.6.6	Software Development Risk	214
15.6.7	Sustainment and Maintenance Risk	215
15.6.8	Commercial Off the Shelf (COTS) Product Risk	215
15.6.9	Product Traceability Risk	216
15.6.10	Natural Risk.....	216
15.6.11	Mapping Sources of Supply Chain Risk to the Aviation Mission Stack	216
15.7	Applied Cyber Supply Chain Risk Management (C-SCRM)	217
15.8	Cyber Supply Chain Risk Management (C-SCRM) Key Practices	219
15.8.1	Foundational Practices	219
15.8.2	Sustaining Practices	220
15.8.3	Enhancing Practices	221
15.9	Integrating C-SCRM with Risk Management Framework (RMF) Security Controls .	222
15.10	Summary and Additional Resources	222
Part 4	Program Management and Executive Leadership.....	225
Chapter 16	The Critical Role of Executive Leaders in Cyber Resilient Aviation.....	226
16.1	Senior Leadership as a Recipe for Success	227
16.1.1	Accepting Responsibility and Accountability	227
16.1.2	Tying Strategic Risk Exposure to Targeted Cybersecurity Risk Mitigation	227
16.2	Conducting Both Internal and Independent Assessments	228
16.3	Prioritizing the Mitigation and Tracking of Cybersecurity Risks	229
16.4	Placing Special Emphasis on the Microelectronics Supply Chain.....	230
16.5	Conclusion.....	231
Part 5	Appendices	232
Appendix A	Acronyms and Abbreviations	233
Appendix B	Bibliography	235

Table of Figures

Figure 1. Aviation Mission Stack	4
Figure 2. The Aviation Mission Stack and Cyber Contested Environments	11
Figure 3. Insider Threats in the Aviation Mission Stack	15
Figure 4. Examples of Supply Chain Risk.....	16
Figure 5. Vulnerability Lifecycle in Aviation Systems	20
Figure 6. Safety and Adversarial Cyber.....	22
Figure 7. Unknown Factors of Cyber-Attack Probability.....	23
Figure 8. The C-I-A Triad.....	55
Figure 9. Systems Engineering “V” Model	58
Figure 10. Aviation Platform Critical Mission Framework.....	70
Figure 11. Threat Modeling Relationships (51).....	115
Figure 12. Notional Attack Tree Against Aircraft Navigation	118
Figure 13. Risk Cube Example	127
Figure 14. CT&E and the DoD Acquisition Lifecycle (62).....	134
Figure 15. CT&E and the Systems Engineering Lifecycle.....	135
Figure 16. Aircraft Boundary.....	139
Figure 17. The Risk Management Framework Steps.....	145
Figure 18. RMF and the Systems Engineering “V” Model	146
Figure 19. Aviation Platform Critical Mission Framework.....	160
Figure 20. Data Extraction Methodology Flowchart	191
Figure 21. Aviation Mission Stack	204
Figure 22. Supply Chain Risk Management Flow Downs	206
Figure 23. Enterprise Visibility, Understanding, and Control of a Supply Chain. (74)	209
Figure 24. What Are We Protecting? (84)	211
Figure 25. NIST SCRM Framework Model	218
Figure 26. Four Aspects of Supply Chain Risk Management (74).....	219

Table of Tables

Table 1. Content Recommendations for Various Roles in Secure Aviation Design	7
Table 2. Adversarial threat tiers described in GAO-19-128. (11)	25
Table 3. Cyber Survivable Attributes (CSAs) (38).....	69
Table 4. Mapping the Cyber Kill Chain for Aviation Systems.....	107
Table 5. RMF, Systems Engineering, and FSAD Mappings	147
Table 6. Security Aspects of the Aviation System Stack.....	165
Table 7. Relationship Between the Aviation Mission Stack and Common Supply Chain Risks.....	217

A Message from the Project Sponsors

In an era where aviation systems rely on intricately intertwined cyber-enabled devices and digital infrastructure, ensuring the resilience of aircraft against cyber events has become an imperative. This guidebook was chartered in recognition of the clear and present threats against flight platforms in both the commercial and military sectors. It is a result of ongoing collaborative efforts of cyber aviation specialists from Lockheed Martin, our partners, and our peers. The consolidated information presented here represents a significant stride toward advancing the state of the practice of cyber resiliency within the aviation domain.

The Fundamentals of Secure Aviation Design (FSAD) is a groundbreaking initiative that signifies our dedication to aviation excellence and our commitment to cybersecurity. As a leader in aerospace and defense, Lockheed Martin takes immense pride in presenting this comprehensive guidebook. It is a testament to our unwavering pursuit of innovation and security.

At Lockheed Martin, we recognize that safeguarding aircraft assets extends beyond our own products and engineering development teams. The aviation industrial base relies on all stakeholders to implement the best practices required to produce cyber resilient systems. Aircraft platforms are only as secure as the weakest link. That is why this guidebook has been publicly released. It is simply too important not to share.

We extend our gratitude to the collaborative efforts that have gone into the creation of this resource and to our industry partners who will help us in our collective journey to “build security in.” Together, we will forge ahead into a future where all aviation systems are resilient against the challenges of an ever-evolving cyber threat landscape.



Jake Wertz

Vice President, F-35 Combat Systems
Lockheed Martin Aeronautics



Vince Frese

Director, Emerging Technologies
Lockheed Martin Aeronautics

Acknowledgements

The Fundamentals of Secure Aviation Design (FSAD) Guidebook is a collective endeavor, drawing on the expertise and dedication of numerous contributors from across the aviation and cybersecurity fields. As the final editor and lead author, I have had the privilege of distilling a wealth of knowledge that addresses the critical and evolving challenges of designing secure aviation systems in an increasingly interconnected world. I am deeply grateful to all the people whose insights and efforts have been instrumental in shaping this work. Together, we hope this book serves as a vital resource for professionals, researchers, and students dedicated to “building security in” to aviation systems.

Much of the credit for this product goes to Tambre Paster who was the original driving force behind this effort. Her extraordinary dedication and hard work provided a solid foundation from which this final guidebook emerged. Similarly, retired Lockheed Martin Cyber Fellow Gerry Ourada played a pivotal role by synthesizing the inputs from multiple contributors into an initial cohesive framework.

Lt Gen James Kevin McLaughlin (Retired), Eugene Moore, Marcus De La Garza, Tambre Paster, and Jeff Chang are credited with primary authorship of select chapters within this book. However, numerous other behind-the-scenes contributions significantly shaped this project, impacting it in ways that are not obviously visible in the final version. Christopher Goulet, Jeff Langham, Trey Jones, Sachin Kamath, and Scott Stokely elevated the quality of this book through their initial ideas and inputs. The editorial assistance and meticulous attention to detail of Lex Thiesse was similarly invaluable.

Dr. Bill “Data” Bryant and Dr. William “Dollar” Young of Modern Technology Solutions Inc. provided advanced insights and ideas that greatly influenced this current edition. They have illuminated the path toward future evolution of this material. Their commitment to advancing secure aviation engineering is unparalleled.

Finally, a special debt of gratitude goes to Jake Wertz, the visionary sponsor of this project. Without his exceptional leadership, strategic foresight, and unwavering support, the FSAD guidebook never would have come to fruition. I am also thrilled to welcome Vince Frese as the new sponsor of this initiative. His enthusiasm and dedication toward carrying this important work forward is greatly appreciated.

I am deeply humbled and honored to have been a part of this remarkable journey, collaborating with such talented and dedicated individuals.



Teresa Merklin
Cyber Fellow
Aeronautics Cyber Range
Lockheed Martin Aeronautics Company

Part 1 – Foundations



More than 300,000 parts from aerospace leaders across the industry go into the F-35's assembly. Here, the jet comes together at the Fort Worth, Texas, Lockheed Martin production facility.

Chapter 1

Introduction

Modern aviation is driven by complex hardware and software that powers cutting-edge flight control, navigation, and communication systems. Harnessing the potential of cyber technologies has enabled aircraft to transcend the limitations of traditional analog and mechanical methods. It unlocks unprecedented levels of performance, efficiency, and speed. In military aviation systems, cyber-enabled capabilities have become a game-changer, delivering a decisive combat advantage that has revolutionized modern warfare. Cyber is the most efficient, economic, and rapid means of developing and delivering capabilities that redefine the boundaries of what is possible in aviation.

Unfortunately, the advanced performance of modern aviation platforms also creates an inherent dependence on cyber-enabled systems. Failures, human errors, accidents, and hostile actions have always had the potential to cause loss of capabilities, platforms, and even human life. In the past, those acts were strictly physical or kinetic in nature. The emergence of cyber-enabled aviation technologies has added new electronic attack surface as a potential target for hostile threat actors.

Both military and commercial aviation systems represent a vital part of national critical infrastructure. On the commercial side, aviation is a key part of the transportation sector essential to both the physical and economic security of the nation. Military aviation platforms, on the other hand, are crucial to project power and defend national interests.

Hostile cyber-based attacks against both commercial and military aviation systems is a significant threat to national security. It is incumbent on the aircraft manufacturers and designers of these systems to build aviation platforms that are not vulnerable to adversarial cyber-attacks. In addition, those systems must be resistant to errors and accidents that may also lead to failures, loss of platform, or even loss of human life.

Aviation systems must be engineered and built to support cyber resiliency. This guidebook, The Fundamentals of Secure Aviation Design (FSAD), was created to help all stakeholders in the aviation industrial base understand what is required to create and sustain secure cyber resilient aircraft systems.

1.1 Cyber Resiliency

Developing systems for **cyber resiliency** requires a conceptual understanding of the meaning of that term. For the purposes of this book, cyber resiliency is defined as **the ability of a system to anticipate, withstand, recover from, and adapt to changing conditions to maintain the mission critical functions required for the system to achieve operational objectives.** (1)

Aviation cyber resiliency requires all the software and hardware in the aircraft to operate predictably, safely, and securely, even when subjected to failures, errors, accidents, and hostile actions. That means that all aircraft parts, ranging from the smallest subcomponents all the way up to the most complex assemblies, must be designed and procured in accordance with the cyber resiliency goals of the system.

The objectives of cyber resiliency are centered around the directive to “maintain the mission critical functions” to “achieve operational objectives.” That means cyber resiliency makes the most sense when considered in the context of the essential missions performed by the system. The imperatives to “anticipate, withstand, recover from, and adapt” must be performed within the context of mission execution.

Cyber resiliency requires all aviation system parts, ranging from the smallest subcomponents through the most complex hardware and software assemblies, be designed to maintain mission critical functions. However, the context of the mission performed will vary from component to component and likewise from system to system.

Aviation platforms include a multitude of very small electronic parts. This includes things like resistors, capacitors, inductors, diodes, transistors, relays, switches, and small integrated circuits. These parts were unlikely to have been developed for use on a specific aircraft or aviation application. Indeed, it would be economically foolish if that was the case in many instances.

As a simple example, the typical mission essential function of a resistor is to regulate the flow of electrical current when installed in a circuit. Additionally, a resistor may have additional performance constraints, such as maintaining a stable resistance value across a specified temperature range. The essential mission of a resistor is relatively simple. Without system context, it does not matter if the resistor is used in an aircraft or a garage door opener. It still must perform that mission essential function.

In contrast, the hardware and software assemblies in aviation systems will inevitably have mission essential functions that are considerably more complex. Those parts are much more likely to require cyber resiliency with direct relevance to the primary missions of the aircraft.

Engineering systems for cyber resiliency requires that the critical mission for the context of the component, subsystem, line replaceable unit, or entire aircraft are considered and supported as the system is conceptualized, designed, implemented, operated, and sustained.

1.2 Cyber Resiliency and the Supply Chain

Every part integrated on an aviation system has a foundational role in the cyber resiliency posture of the aircraft. Those components don't simply materialize out of thin air. The network of organizations that design, produce, and move goods to a final destination is colloquially referred to as the "supply chain."

Figure 1 presents a conceptual model of an aviation mission stack (2) that encapsulates the complex relationship between various aspects of the supply chain and ultimate mission of the aircraft.¹

Each layer in the stack simultaneously represents an opportunity to support cyber resiliency as well as a potential threat vector for vulnerability or attack.

All aviation systems are acquired to fulfill some purpose. For DoD aviation platforms, the objective is to project military power to defend our nation's interests. Commercial aircraft are procured for transportation of people or cargo. Even recreational airplanes exist to fulfill some mission, even if that is simply for entertainment.

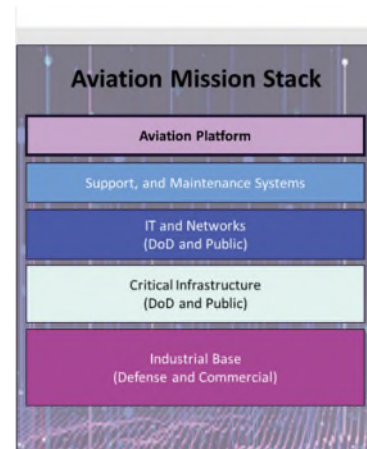


Figure 1. Aviation Mission Stack

This FSAD guidebook was created to support development and production of cyber resilient flying platforms. It defines a foundation of systems engineering and cybersecurity analysis that must be performed to accomplish that objective.

However, cyber resiliency for an aircraft cannot be achieved in isolation. The aviation mission stack is essential for understanding the full scope of what is required to achieve the ultimate mission of the platform. Each layer in the stack simultaneously represents an opportunity to support cyber resiliency as well as a potential vector for hostile cyber-attack. That goes well beyond the **aviation platform** situated at the top of the stack.

The operation of all aircraft relies on offboard systems that provide maintenance support. Aviation platforms that do not have electronic interfaces are becoming increasingly rare. Many systems connect to laptop computers or similar cyber-enabled devices to perform routine servicing and maintenance. Consequently, those **support and maintenance systems** must be developed and operated with cybersecurity objectives in mind to avoid compromising the cyber resiliency posture of the aircraft.

¹ This aviation mission stack was adapted from a concept originally developed and presented by Mr. John Garstka, Director, Cyber Warfare in Office of the Under Secretary of Defense for Acquisition and Sustainment. (2)

Sophisticated aviation platforms rely on data from **Information Technology (IT) and networked systems**. That information can originate from, or traverse through, the public internet or private intranets, such as those operated by the DoD. Secure aviation design requires understanding these data flows as cyber-attack surface that could potentially compromise the security posture of the aircraft. Information originating from external networks may not be trustworthy. That must be accounted for when developing and sustaining aviation platforms.

Aircraft are a part of the transportation sector, which is a vital component of national **critical infrastructure**. (3) We rely on flying platforms to transport people and goods throughout the country and overseas. Aviation systems depend on other parts of the critical infrastructure to operate. Secure aviation design may include analysis and recommendations for aircraft support in the event of extensive outages of power, water, and fuel delivery systems.²

The aviation **industrial base** includes manufacturers that support both defense and commercial platforms. The role of the Defense Industrial Base (DIB) in development and support of military platforms is well documented. While the commercial industrial base enjoys significantly less formal attention, it is equally important for the overall cyber resiliency of all aviation platforms.

Many aircraft manufacturers support both military and commercial aircraft. Additionally, sometimes military aviation platforms reuse parts and subsystems that were originally designed for commercial use. For that reason, the FSAD is targeted to the entire industrial base for aircraft systems.

The aviation mission stack is the perfect jumping off point to consider the full scope of the supply chain for flying platforms. The vast majority of the aviation industrial base have suppliers and are themselves a supplier of parts, software, and information. Those dual roles bring great responsibility.

It is imperative for all members of the industrial base to implement best practices that secure the parts, software, and services supplied to aviation platforms. At the same time, those manufacturers need to be aware of the impact their suppliers also have on the security of their own products.

Cybersecurity of every aviation system relies on locating and eliminating potential points of vulnerability. The vast supply chain of complex modern aircraft contains many potential weak points. It is the duty and responsibility of the entire industrial base to ensure products installed in aircraft systems are not the source of cyber resiliency failures. That demands vigilance in everything we do.

² Aviation platforms operating in the transportation sector are both at the top and the middle of the Aviation Mission Stack. Diving into the recursion is beyond the scope of this guidebook.

1.3 Using This Book

The Fundamentals of Secure Aviation Design (FSAD) is a guidebook created to address the informational needs of a diverse set of aviation system suppliers and manufacturers. This document is organized into parts, which are written to provide targeted information relevant to specific roles and responsibilities for both individuals and organizations.

Part 1 - Foundations. This describes the basic information required for all individuals and organizations with a role in the aviation supply chain. It is written for all suppliers, including those who deliver products that are not cyber-enabled. All aviation industry stakeholders should familiarize themselves with Part 1.

Part 2 – Secure Architecture and Design for Air Systems. This part of the FSAD guidebook provides information needed by all people and organizations involved in the design of cyber-enabled parts. It introduces unique concepts for developing, deploying, and operating aircraft systems in cyber contested environments.

Part 3 – Cyber Resiliency Specialty Domains. Securing the supply chain is a critical aspect of aircraft system cyber resiliency. This part of the FSAD guidebook is intended for individuals and organizations that develop or acquire cyber-enabled hardware parts, components, or software for aviation systems. It is also appropriate for personnel who implement or manage aircraft-related supply chains.

Part 4 – Program Management and Executive Leadership. Cyber resilient aviation systems require executive level commitment to the principles outlined in the FSAD. This part of the guidebook outlines how program managers and senior leaders can support and instill disciplined cyber security best practices within their organizations.

Table 1 provides a mapping between the contents of this guidebook against organizational roles and responsibilities. Every place marked “Required” indicates a section containing essential knowledge for that specific role. While it is a good idea for all stakeholders to be familiar with the entire contents of this guidebook, that is not strictly required for people serving in roles with a relatively narrow function in the aviation industrial supply chain.

It is recommended that all stakeholders have at least cursory knowledge of the entire contents of the FSAD. However, at the individual level, each person should first concentrate on what is directly relevant to their specific role and responsibilities.

Role	Part 1	Part 2	Part 3	Part 4
	Foundations	Secure Architecture and Design for Air Systems	Cyber Resiliency Specialty Domains	Program Management and Executive Leadership
Supplier / Developer Non-Cyber-Enabled Products	Required	Optional	Optional	Optional
Supplier / Developer Software	Required	Required	Required	Optional
Supplier / Developer Hardware	Required	Required	Required	Optional
Supply Chain Management	Required	Required	Required	Optional
Program Management and Executive Leadership	Required	Optional	Optional	Required
Prime Contractors and Platform Integrators	Required	Required	Required	Optional

Table 1. Content Recommendations for Various Roles in Secure Aviation Design

1.4 References and Additional Resources

Due to the complex and rapidly evolving nature of the cybersecurity landscape, there is an onslaught of directives, standards, policies, and procedures seeking to provide guidance for aviation and critical infrastructure systems. The evolutionary nature of cyber threats makes it crucial for government agencies and industry thought leaders to continuously update their recommendations to keep up with known and perceived cyber risks to aviation systems.

Consequently, this FSAD guidebook focuses on the timeless foundational concepts that remain constant over time. Appendix B contains a list of the resources that were referenced during the creation of this document. As a cautionary reminder, the items in that list will most certainly undergo revisions and updates as the industry understanding of cybersecurity and cyber resiliency continues to evolve. Additionally, Appendix A contains a list of useful acronyms and their definitions.

1.5 The Lockheed Martin FSAD Portal

This version of the Fundamentals of Secure Aviation Design (FSAD) is available for free public download from the FSAD Portal at <https://www.lockheedmartin.com/en-us/suppliers/business-area-procurement/aeronautics/fsad-guidebook.html>.

As the FSAD evolves over time, it may reference outdated or superseded versions of public standards. Due to the rapid pace of change in cyber policy directives, it is challenging to ensure that all referenced materials remain current. Therefore, when using any of the referenced source material, it is essential to verify that the information is up-to-date and validate its relevance to ensure accuracy and compliance with the latest standards and directives.

Comments and feedback on the FSAD may be submitted to **fsad.fc-aero@lmco.com**.

Chapter 2

The Threat is Real

One of the most significant challenges threatening the resiliency of aviation systems is a failure to recognize cyber risks. A cyber event resulting from a natural failure, human error, or adversarial cyber action could result in loss of mission performance, the aircraft platform, and even human life.

The United States military recognizes cyberspace as a warfighting domain. During times of conflict, our adversaries can be expected to deploy offensive cyber-attacks intended to generate a decisive advantage in the battlespace.

Additionally, cyber-attacks may be used as a precursor to escalation of traditional kinetic warfare. The cyber domain can be leveraged for information and intelligence gathering well in advance of declared conflict. In fact, cyber intrusion is already known to be in widespread use for espionage operations.

The pervasive nature of cyber-enabled systems presents a large and attractive attack surface to a hostile adversary. Cyber-attacks will certainly be waged against enterprise IT computing assets and commercial-off-the-shelf (COTS) components that are used to maintain and sustain aviation platforms. Additionally, the embedded cyber components in the aviation platforms are also attractive targets. Those systems are where the most dramatic and impactful cyber effects are possible.

The people who acquire, develop, operate, sustain, and maintain aviation systems cannot ignore the very real cyber threats that exist against aircraft platforms. Any cyber logic-bearing component that interacts with the external world either through direct communications interface or other environmental mechanisms is a potential vector of direct attack against the aircraft system on which it is installed.

From the largest most complex line replaceable units (LRUs) down to the smallest logic-bearing component, we all share a collective responsibility to recognize the potential of cyber threats against aviation systems and work in concert toward effective mitigation.

2.1 Cyber Contested Operational Environments

Aircraft platforms operate within a broad information environment. Humans and cyber-enabled systems work on aggregated data that is at risk of adversarial manipulation or interference. Military aircraft are flown in information contested environments. The DoD has developed doctrine, operational procedures, and technical frameworks designed to guide effective mission performance during the “fog of war.” (4)

Cyber contested spaces are a subset of the broader information contested environment. In fact, cyber-attacks and cyber-effects are an attractive attack surface for threat actors who are intent on disrupting the operational and mission performance of aviation systems.

The capabilities delivered by cyber-enabled aviation platforms provide a dominant performance advantage both in day-to-day operations and during times of military conflict. In fact, many modern aircraft could not be effectively operated by humans without cyber-enabled flight controls. An over-reliance on those advanced capabilities could turn dominance into a disadvantage, if cyber vulnerabilities are known and accessible to hostile adversaries.

Military strategists assert that in conventional warfare, victory goes to the side that deals most effectively with the “fog of war.” (5) That concept refers to an inherent absence of accurate situational awareness and uncertainty that forces decisions and actions to be taken with less than perfect knowledge.

Cyber is an enabling technology that can greatly enhance informational superiority that should, in theory, reduce the “fog of war.” However, if an adversary denies, degrades, disrupts, or compromises the integrity of cyber systems, that advantage can quickly become a liability.

Commercial aviation critical infrastructure also operates in a cyber contested environment. Violent extremists have previously demonstrated a willingness to commit acts of terrorism against civilian infrastructure using kinetic means. It would be naive to assume that these threat agents would not exploit cyber mechanisms for the same purpose.

Aviation companies that operate that segment of our critical infrastructure are also attractive targets for economic espionage and ransomware. Even attacks motivated only by financial gain could be crippling to a commercial aviation company or possibly the entire aviation system.

In early January 2023, the Federal Aviation Administration (FAA) Notice to Air Missions (NOTAM) system suffered an outage that caused a nationwide halt to all flight departures in the United States. While that outage was eventually traced to accidental corruption of a data file, the incident demonstrates the potential impact and widespread nature of a conceptual cyber-attack in the future. In addition, it also illustrates that failures of cyber resiliency can originate from unintentional action or error.

Our commercial aviation systems are recognized as legitimate military targets during times of declared conflict. As a part of the critical infrastructure, aircraft transportation supports movement of goods and personnel. Commercial cargo is a vital part of the supply lines that support military operations.

Aviation systems, whether they be military or commercial critical infrastructure, operate in cyber contested environments. The people who acquire, engineer, develop, operate, maintain, and sustain aviation systems must take steps to build cyber resiliency in the platforms that will operate as needed in a pervasive cyber contested environment.

2.2 The Aviation Mission Stack and Cyber Contested Environments

While operational use is where the most devastating cyber-attacks are likely to occur, it is not the only domain that is cyber contested. Figure 2 augments the aviation mission stack introduced in Chapter 1 with examples of the types of cyber adversarial action that could potentially occur at each tier.

The aviation platform can be a direct target of cyber-attacks intended to deceive, degrade, deny, or even destroy the aircraft. For example, jamming of the Global Positioning System (GPS) satellite signals could degrade or deny navigation capabilities. Spoofing a GPS signal could cause the pilot to lose accurate knowledge of the location of the aircraft. That could be devastating in mountainous terrain or when making instrument landings under low visibility conditions.

Some stakeholders erroneously believe that because aircraft do not usually rely on traditional IT networks for aviation, navigation, and communication, that these platforms are impervious or at low risk of direct cyber-attack. That is a dangerous belief because it can result in an understatement of the level of cyber risk to the platform. Such false assumptions could directly lead to a failure to engineer for cyber resiliency during development. That could directly lead to vulnerabilities accessible to hostile adversaries when the aircraft is in operational use.

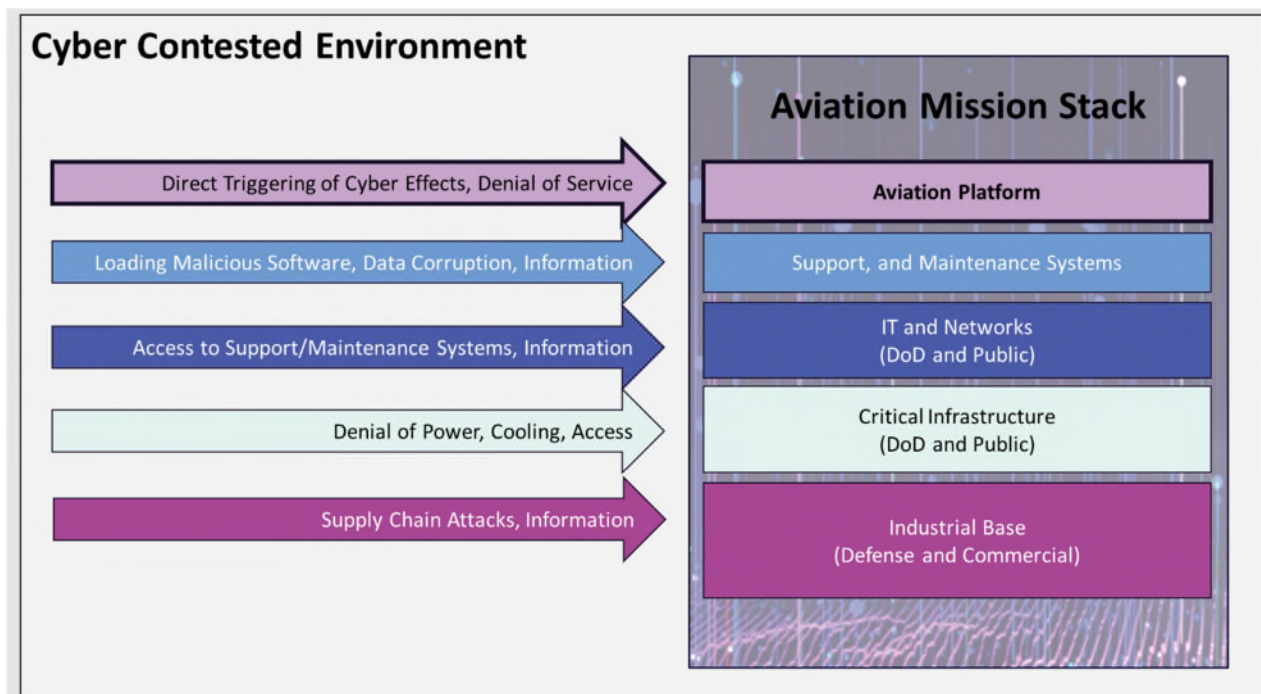


Figure 2. The Aviation Mission Stack and Cyber Contested Environments

It is important to establish a widespread appreciation that all tiers of the aviation mission stack present significant access vectors for adversarial cyber threat agents. In fact, most of those tiers are more accessible to attackers than the aircraft itself.

Support and maintenance systems connect directly to the aviation platform and are typically used to load software, data, and to determine its maintenance status. The access to insert or trigger a malicious payload on an aircraft during flight operations is challenging, but not impossible. However, it is significantly less complex for a cyber adversary to accomplish that objective when the aircraft is connected to ground support systems. Maintenance connections are also a valuable source of information that might be helpful as a part of an information gathering campaign and a precursor to future cyber-attacks against the aircraft.

Consequently, the people who acquire, engineer, design, operate, maintain, and sustain support and maintenance systems must be cognizant that those assets are both a target and an access opportunity for adversarial cyber actions. These systems operate in cyber contested environments.

While some support and maintenance systems may be used without active network connections, many exchange data and information with traditional enterprise IT platforms to execute its functions. Those systems may connect to DoD networks, corporate intranets, and in some cases even the public internet.

Traditional IT networks are a cyber contested environment with a well-documented and lengthy history of active cyber exploitation and incursion. The people who acquire, engineer, design, operate, maintain, and sustain enterprise IT systems must be cognizant that they are operating in a cyber contested environment. Additionally, the systems that have interconnections or rely on external data for maintenance must be designed with the understanding that adversarial cyber-attack could occur across those interfaces.

2.3 Threat Intelligence and Aviation Systems

Cyber threat intelligence consists of information or indicators of an attack against cyber-enabled systems. That knowledge can be used to prevent or mitigate an attack. The DoD and the private sector invest considerable resources collecting, generating, and disseminating cyber threat intelligence.

Standard IT equipment, which includes things like desktop computers, servers, and networking devices, operates under an abundance of threat intelligence. The relative wealth of data is available because those computing systems are ubiquitous in industrialized parts of the world. Many people have direct physical access to commercial IT systems, which makes it easier to conceptualize, develop, and test malicious payloads against those platforms.

Much of the threat intelligence for enterprise IT systems is detected and collected by the private sector. As such, sharing and dissemination of that information isn't restricted by DoD classification considerations or safety concerns.

In contrast, public threat intelligence for aviation systems is scant to nonexistent. Additionally, the information we do have is typically closely guarded. When shared, it is on a strict need-to-know basis. Cyber threats to military aviation systems are typically regarded as classified information and require special access authorizations.

Unfortunately, sometimes people and organizations interpret the dearth of specific threat intelligence against aviation systems as an indication that risk of catastrophic cyber-attack is negligible. However, nothing could be further from the truth.

The National Strategy for Aviation Security (NSAS) recognizes potential cyber disruption from nation state actors and other organized groups as a serious threat to the aviation ecosystem. Threat actors are known to be gathering information which could lead to disruptive cyber-attacks that are a profound threat to our national security. (6)

Our aviation systems, whether they be military or commercial critical infrastructure, are subject to adversarial cyber threats. The people who acquire, engineer, develop, operate, maintain, and sustain aviation systems must build cyber resilient platforms that are resistant to these pervasive cyber threats. To complicate matters, this must be achieved despite limited actionable threat intelligence. That does not eliminate the responsibility, but rather just makes it more of a challenge.

2.4 Aviation Threat Actors

Cyber threat intelligence definitions frequently include several categories of information required to assess risk. In addition to indicators of attack or compromise, some claim that it is important to characterize the attackers themselves. The identities, motivations, and capabilities are potentially useful when a successful cyber-attack is underway and during forensic analysis.

However, for aviation systems what is most important is anticipating cyber-attacks and building platforms that can withstand and recover from likely adversarial cyber activity. It is less important to know who the specific aviation threat actors are than it is to be aware of the general categories of threats against a system. Threat actors might be well-funded nation states. They might be terrorists or political anarchists. Some threat actors operate for criminal or economic gain. There are also recreational cyber threat actors who may attack systems simply for the challenge or personal recognition.

Insiders with direct physical access to aviation systems are potential threat actors potentially harbored deep inside the development supply chain. In fact, defending cyber-enabled systems from the insider threat is one of the most challenging and frequently overlooked consideration when developing cyber resilient systems.

Understanding and accepting that aviation platforms will be confronted with substantial cyber threats from various threat actors is crucial. Aircraft must be designed to withstand cyber adversity, regardless of the originating source.

2.4.1 Advanced Persistent Threats

Some threat actors operate using sophisticated techniques in a sustained adversarial cyber campaign. On traditional IT systems, these attackers frequently establish an undetected foothold into a system and then slowly and systematically pursue their objectives. These cyber adversaries have incredible patience and long-term operational horizons.

The phrase “Advanced Persistent Threats” and the associated acronym, APT, was first coined by the Air Force in 2006. The moniker was needed to describe characteristics of various groups for unclassified public dissemination. NIST formally defines the APT as follows:

“An adversary that possesses sophisticated levels of expertise and significant resources which allow it to create opportunities to achieve its objectives by using multiple attack vectors (e.g., cyber, physical, and deception). These objectives typically include establishing and extending footholds within the information technology infrastructure of the targeted organizations for purposes of exfiltrating information, undermining or impeding critical aspects of a mission, program, or organization; or positioning itself to carry out these objectives in the future. The advanced persistent threat: (i) pursues its objectives repeatedly over an extended period of time; (ii) adapts to defenders’ efforts to resist it; and (iii) is determined to maintain the level of interaction needed to execute its objectives.” (7)

Exfiltration of information is a common objective of the APT. In some instances, access to information is the exclusive objective of an APT campaign. The use of espionage for access to political or military information pre-dates cyber-enabled systems by centuries. Espionage can also be motivated by economic gain or theft of intellectual property.

In cyber-enabled systems, information exfiltration can be an early part of a multistage attack. For example, information gleaned from one system may be used to learn how to prepare and execute attacks against another. In aviation systems, design information or software exfiltrated from the developer can be used to identify potential vulnerabilities in the operational platform. Additionally, data collection and exfiltration from aircraft could be used to discover viable attack vectors against those systems.

Another known objective of the APT is to undermine the system by injecting vulnerabilities or making changes to impede performance. An APT with access to the development systems for components of an aircraft platform could implant means to trigger cyber-effects or system failures at a future time of their choosing.

At the fundamental level it is essential to identify and accept that aviation systems face significant cyber threats from nation state funded APT organizations. It is important for the people who acquire, engineer, develop, operate, maintain, and sustain aviation systems to build cyber resilient platforms that anticipate adversarial action from very sophisticated and highly skilled cyber threat actors.

2.4.2 The Insider Threat

Cyber threats posed by people with direct access to the development, sustainment, and operational aviation systems cannot be ignored. Threat actors with insider access can potentially cause a broad range of disruptions and damage to aircraft. NIST describes the insider threat as follows:

“The threat that an insider will use her/his authorized access, wittingly or unwittingly, to do harm to the security of organizational operations and assets, individuals, other

organizations, and the Nation. This threat can include damage through espionage, terrorism, unauthorized disclosure of national security information, or through the loss or degradation of organizational resources or capabilities.” (8)

It is important to understand that an insider may have agency in malicious cyber adversarial action. However, an insider might unwittingly serve as a conduit for adversarial access. A good example of an unwitting insider in an enterprise IT system is a user who clicks on a malicious file in a spear phishing email message. An analogous attack example in an aviation system is convincing an aircraft maintainer to load malicious software to the platform believing that it contains important security updates.

Figure 3 annotates the aviation mission stack with examples of insider threat agents that could be a conduit to disruptive cyber effects on the operational aircraft platform. That diagram reveals the potential breadth and scope of the insider threat.

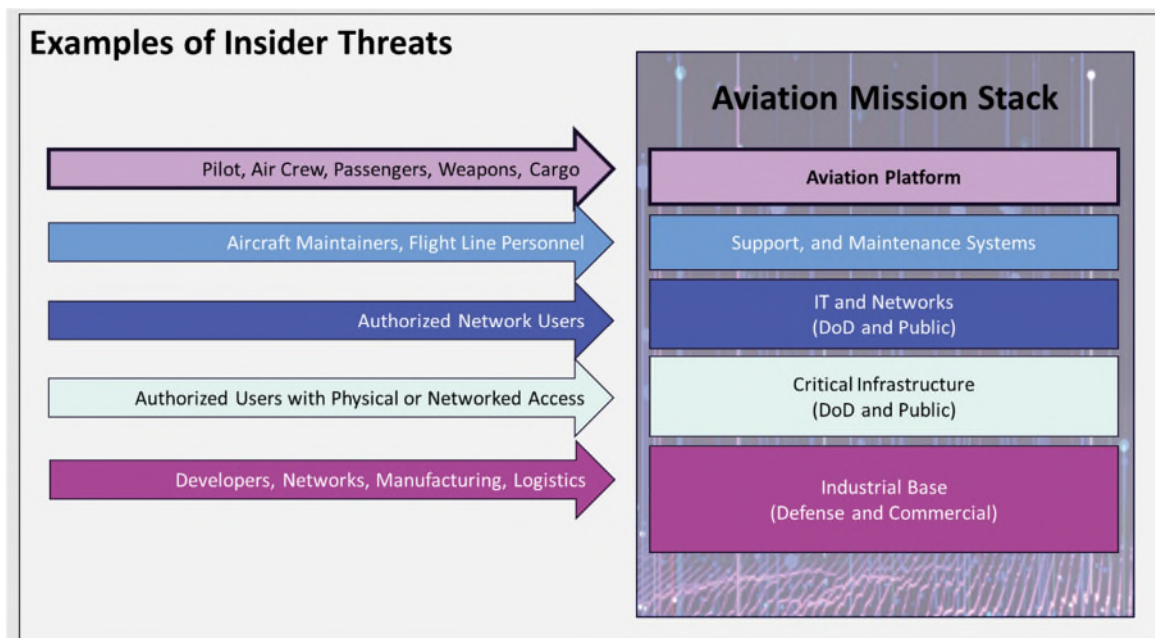


Figure 3. Insider Threats in the Aviation Mission Stack

Mitigating the insider threat is a challenge because system access is integral to developing, operating, maintaining, and sustaining the aviation system. Simply eliminating access is not a viable alternative in most scenarios.

In addition to direct malicious action, the insider threat could provide an adversary with system knowledge or even data exfiltration. That information could be an end-state objective of an adversary for the purposes of espionage. Additionally, stolen data could simply be the initial information gathering stage of a long-term sophisticated attack.

At the fundamental level it is important to understand and accept that aviation systems have a tremendous number of insiders who may be an unwitting or malicious cyber threat agent against the platform. It is imperative for the people who acquire, engineer, develop, operate, maintain,

and sustain aviation systems to build cyber resilient platforms that anticipate the threats and risks associated with insider access.

2.5 The Supply Chain Threat

As mentioned in the introduction of this guidebook, the aviation mission stack is a great framework for illustrating the cyber supply chain risk inherited by operational aviation platforms. Figure 4 overlays examples of supply chain risks against the aviation mission stack. While this illustration is not meant to be an exhaustive list, it is apparent that supply chain risk exists at all levels.

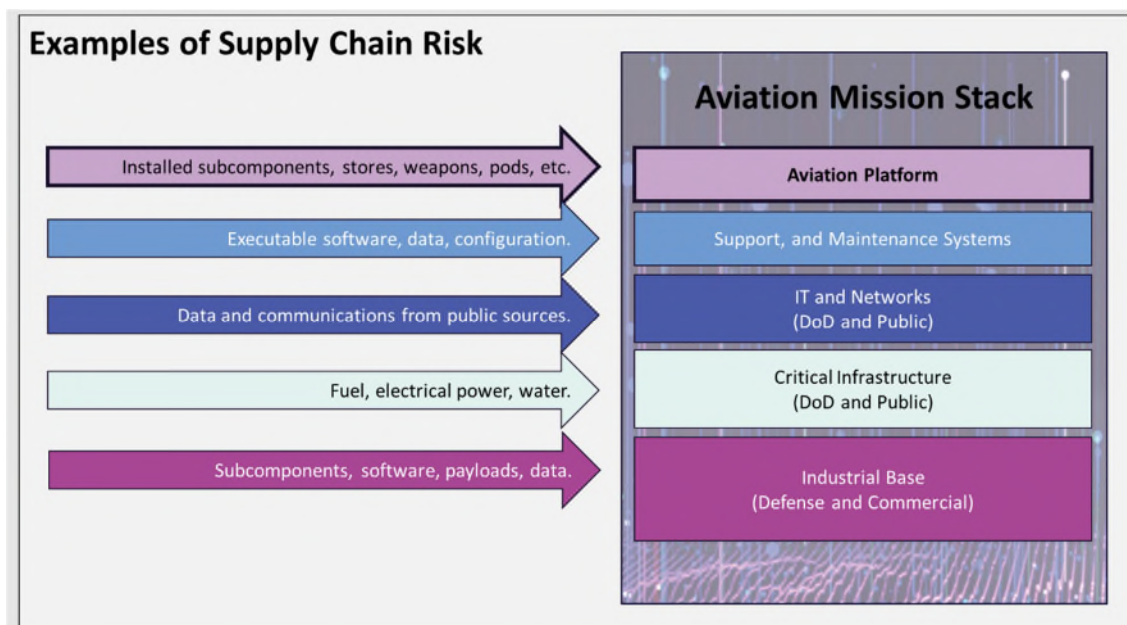


Figure 4. Examples of Supply Chain Risk

Chapter 15 of this FSAD guidebook provides more details on best practices for supply chain management of aircraft systems. That section is written for people and organizations who perform aviation supply chain management. However, all participants in the aviation mission stack must understand the basic supply chain management concepts.

The **Industrial Base** has both a role and responsibility in supply chain cyber risk management. That includes selecting and sourcing parts that do not introduce inherited risk to the system. The Industrial Base is also responsible for ensuring that development tools and environments are secure from cyber tampering. Finally, the industrial base is expected to design, develop, and deliver subcomponents, parts, and systems that are cyber resilient.

Development, production, and operation of aircraft systems relies on basic services from the **Critical Infrastructure**. In some instances, it may make sense to design aircraft systems that can operate under austere conditions where basic services such as power, water, and fuel are not readily available. Though it is beyond the scope of this document, there is an inherent need for the developers and operators of the critical infrastructure to mitigate risk and attack surface in

their own supply chain. Additionally, commercial aviation systems are regarded as critical infrastructure as a part of the transportation sector.

Networked IT in both the public and private domains presents a vast attack surface. At a high level, all participants in the aviation mission stack must ensure that their networks are secure and that all data from external sources is scrutinized to minimize risk associated with malicious cyber activity. Additionally, aviation systems must be designed with the understanding that any network connection must be viewed as a cyber contested environment. That should impact design decisions for aviation systems that rely on other networks for maintenance, planning, and operational use.

Similarly, the **Support and Maintenance Systems** that connect directly to the aviation platforms must be designed and operated in a way that prevents adversarial cyber-attacks from propagating to the aircraft. These systems typically use traditional IT hardware and operating systems. They also connect to other computing devices and networks for data required to operate, sustain, and maintain the aircraft. Consequently, support and maintenance systems are one of the primary vectors of cyber-attacks against platforms.

The security of the **Aviation Platform** depends on the cyber resiliency of each and every part in the system. Cyber risk incurred in any part of the aviation mission stack is inherited by the platform. Additionally, the aircraft will receive data, updated software/firmware, and replacement parts from the supply chain.

As stated in the introduction to this guidebook, the cybersecurity of an aviation system is profoundly impacted by its weakest link. The vast supply chain of complex modern aircraft contains many potential weak links. It is the duty and responsibility of the entire industrial base to ensure that delivered systems are not the source of cyber resiliency failures. That necessarily includes vigilance in managing and monitoring the supply chain.

2.6 Unique Challenges for Aviation Systems Cybersecurity

As mentioned in the introduction to this chapter, the United States military recognizes cyberspace as a warfighting domain. That means that during times of conflict, both military and civilian critical infrastructure may be subject to offensive cyber-attacks intended to generate a decisive advantage in the battlespace or to disrupt operations.

2.6.1 Scalability

When contrasted with conventional kinetic warfare, the cyber domain presents some unique challenges. If an adversary only has one missile, then he can potentially take down one airplane. Shooting down more than one target requires additional missiles. Ammunition is a consumable resource that must be replenished as it is expended. That constrains the scalability of traditional kinetic attacks.

In most cases, a cyber-attack may be repeated or replicated without limits. This means that a single cyber “bullet” or attack could simultaneously take down an entire squadron of aircraft or

even the entire fleet. To make matters even worse, some “wormable” cyber-attacks can spread independently of a command-and-control mechanism.

2.6.2 Attribution

In kinetic warfare, when an adversary fires a weapon it is possible to determine the origin of the attack with a high degree of certainty. When cyber weapons are deployed, attribution of the perpetrator is complex. Even when an attack is known to have been launched from a particular computer, it is frequently impossible to tell whether the owner of the system is responsible or if they themselves are victims of a cyber intruder illicitly using their computing resources.

2.6.3 Espionage

Cyber-enabled systems present a rich opportunity for military or economic espionage. Systems that are connected to public networks potentially offer an adversary remote access to information and data without exposure to personal risk. Even systems that are in isolated networks or operate with an “air gap” have been proven to be at risk from cyber-attack and espionage.

In the event that cyber data collection and exfiltration is detected on a system, attribution is a challenge. Even worse, these attackers intentionally cover their tracks by deleting log files and audit records. They frequently use strong encryption to hide the data and information that has been stolen.

Economic espionage is frequently motivated by monetary gain. Theft of intellectual property can “save” years of non-recurring engineering and development costs. Cyber espionage on commercial systems can also be driven by informational needs. For example, an airline ticket booking system could be used to track the travel patterns of a person of interest. Commercial freight logistics data could reflect a sudden surge of shipments of a new material to a factory. That could reveal that formulation of a product had been revised along with a significant clue as to the nature of the change.

Espionage in the military aviation system domain is used to learn as much as possible about an adversary before cyber or kinetic conflict is initiated. Exfiltrated information could be used to close a technological gap or to learn of vulnerabilities in fielded systems that could be exploited if tensions escalate.

Additionally, cyber espionage is frequently only part of a much broader sustained campaign. While data collection and exfiltration may be the ultimate goal, establishing a foothold and gathering information to inform lateral movement throughout a system is also a frequent objective. For an aircraft that could ultimately lead to subversion, sabotage, or destructive attacks.

2.6.4 Subversion and Sabotage

An adversarial threat agent with access to cyber systems in the industrial base could insert vulnerabilities or weaknesses for later exploitation in a contested environment. A malicious presence on support and maintenance systems could provide erroneous status to maintainers

negatively impacting fleet readiness or potentially allowing aircraft to operate under unsafe conditions.

Modification of messaging and data could have serious ramifications during times of kinetic warfare. For example, changing target coordinates could cause weapons to miss legitimate military targets and potentially put civilians at risk.

Once again, while subversion and sabotage are a part of traditional kinetic warfare, the scalability of cyber amplifies the effects of successful attacks. Additionally, cyber sabotage could potentially lie dormant and undetected for much longer periods of time than kinetic effects.

2.6.5 Battlespace Preparation

Military strategists believe that cyber-attacks will likely be used in the future as a precursor to escalation of traditional kinetic warfare. In other words, the opposing forces will attempt to use cyber methods to prepare the battlefield for impending escalation of conflict. The effects could include disruption, degradation, deception, denial of service, and even physical destruction.

Much of the technical superiority of military weapons systems in the United States is powered by cyber-enabled systems. An effective cyber-attack by an adversary could produce wildly asymmetrical effects and level the battlefield, or worse.

It is important for the people who acquire, engineer, develop, operate, maintain, and sustain aviation systems to build cyber resilient platforms that anticipate potential cyber adversarial action as a precursor to kinetic attacks. That awareness typically opens the aperture to broader classes of threats and risks to the aircraft.

2.6.6 Accidents, Errors, and Attacks

Cyber effects do not always originate from malicious cyber-adversarial action. Sometimes simple equipment failures and human error can lead to dramatic cyber incidents. Additionally, it is sometimes difficult to determine whether a cyber event was caused by malicious activity or simply an unfortunate coincidence.

In the midst of the “fog of war,” it is possible for the mere perception of a successful cyber-attack to foster confusion and a lack of trust in cyber-enabled systems. It is a paradox unique to the cyber domain that attacks do not necessarily have to be successful to create the desired effect. In other words, making the operator of a system believe that a cyber-attack is in progress may be enough to achieve adversarial objectives. System resets or a complete shutdown may be performed to “mitigate” the perceived attack. That is essentially a self-induced denial of service.

2.7 Conceptualizing Vulnerabilities in Aviation Weapons Systems

The general perception of cyber risk and vulnerability is heavily influenced by recent events and media headlines. As of this writing, the aviation industry is fortunate that a catastrophic cyber-attack has not resulted in an aviation disaster. However, the absence of dramatic headlines of that nature is not a reliable indicator that the risk isn’t real.

Aviation accidents involving the Boeing 737 Max have been attributed to malfunctions and design flaws. (9) However, the existence of those types of flaws should serve as a stark reminder of the possibility of conceptually similar design flaws that could be intentionally triggered by an adversarial cyber-attack.

Any assumption that the absence of substantiated malicious cyber activity in either commercial or military aviation platforms is a valid indication of a lack of cyber vulnerabilities is flawed. In fact, it is highly likely that a latent issue might currently exist in an aircraft that could be intentionally triggered by cyber action.

While both commercial and industry groups have been chartered to elicit and share information about potential aviation system vulnerabilities, there are many challenges to overcome to realize the full potential of those initiatives. There simply isn't much intelligence that can be shared, even when all parties are committed to transparency.

Unlike traditional enterprise IT systems, there are significant barriers to entry for ethical white-hat penetration testers seeking to assess and test aircraft for potential vulnerabilities. Additionally, a paucity of commercially available penetration testing tools means that aviation cyber testers frequently must develop their own tools and train themselves on the specifics of the aircraft. There are very few people and organizations capable of efficiently performing that level of work.

Additionally, both commercial and military aircraft are expensive. The owners of aviation systems are justifiably reluctant to allow cyber penetration tests on their assets.

It is important for all stakeholders in the aviation ecosystem to understand that any latent vulnerability in an aircraft system exists whether it has been detected or not. In fact, who detects potential vulnerabilities, and how that information is shared, has a profound impact on the overall risk and threat level to the aircraft.

Figure 5 illustrates one viewpoint of the vulnerability lifecycle of an aviation system.

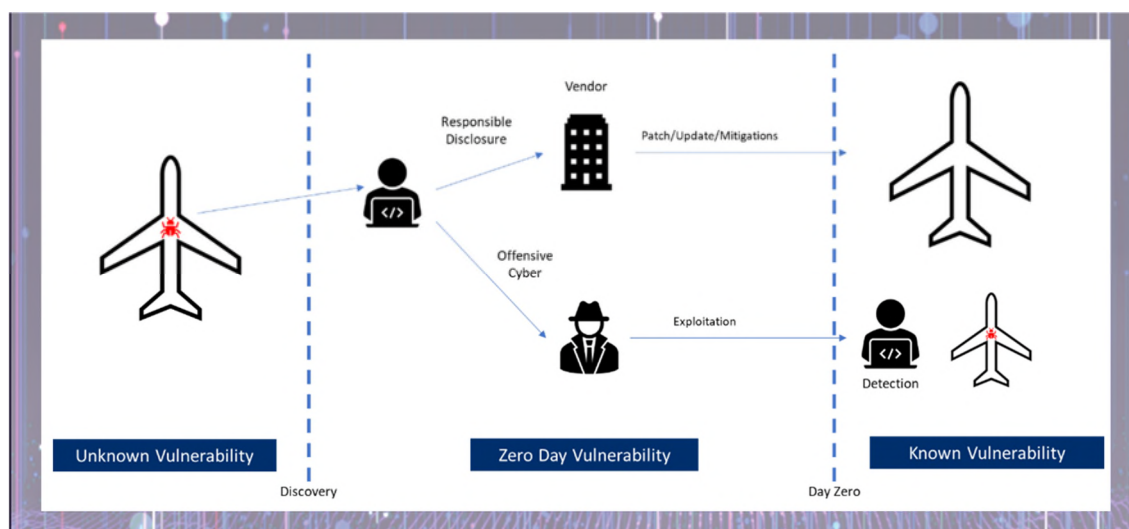


Figure 5. Vulnerability Lifecycle in Aviation Systems

A latent vulnerability may exist in an aircraft that is eventually detected by a person or organization. It is important to note that the weakness in this hypothetical scenario is present, regardless of whether it is ever discovered.

The person or organization that finds the vulnerability has a few options for communicating the existence of the problem. It is possible that the person could decide not to disclose it at all. There are some plausible scenarios where that might occur, but in general that isn't a very responsible course of action. Eventually, someone else will also find it, thus opening up the possibility of cyber exploitation.

The industry standard for responsible disclosure in enterprise IT systems is to notify the manufacturer or the operator of a system about the vulnerability so they can take steps to repair or mitigate the issue. That logically same course of action can be taken for aviation systems via disclosure to the manufacturer or the operating organization.

Unfortunately, another course of action is available to a person who discovers a latent vulnerability on any type of system. They could sell it to brokers and nation state buyers of previously undisclosed zero-day vulnerabilities. In fact, these organizations and nation states quite likely have their own efforts underway to detect vulnerabilities in systems, effectively cutting out the middleman. It would be naive to dismiss the possibility of a lucrative black market for aviation vulnerabilities.

A vulnerability crosses the threshold of day zero when it is publicly exposed or disclosed. One way that can occur is when manufacturers release patches that fix or mitigate it. In enterprise IT systems, the release of a security update sets off a race between adversarial actors and system defenders. Malicious threat agents seek to reverse engineer the software to uncover the vulnerability and then try to create malicious payloads that exploit it before system operators have time to fully patch their systems.

Sometimes a zero-day vulnerability is detected when it is first observed being exploited in "the wild," which is an industry euphemism for the operational environment. When that occurs, there is still a race, but one where the malicious actors enjoy a significant head start.

One of the most frightening aspects of defending aviation systems from catastrophic cyber-attack is the potential that an adversary may have already discovered or injected zero-day vulnerabilities into a platform. The highly sophisticated APT would most certainly elect to defer exploitation until a time of their choosing. That would most likely be associated with an escalation of conflict. In fact, that is a significant source of fear, uncertainty, and doubt in the aviation industry when it comes to cybersecurity.

2.7.1 Safety and Adversarial Cyber Effects

The aviation industry has rigorous regulatory standards and certification processes that contribute to the assurance posture of aircraft systems. Flight controls and safety-critical sensors are developed using well-established design practices, engineering standards, and extensive testing.

It is important to understand that existing safety certifications have historically focused only on errors and accidents. Consideration of intentional malicious adversarial cyber action has only recently become an industry concern. In fact, there are still many who believe that the existing safety standards are adequate to protect aviation platforms from malicious cyber-attack. Unfortunately, that idea is flawed.

Figure 6 presents a conceptual illustration of the limited overlap between safety and adversarial cyber activity. The blue circle on the left represents errors and failures that inevitably arise from aviation system operation. The pink circle on the right represents adversary-initiated cyber events.

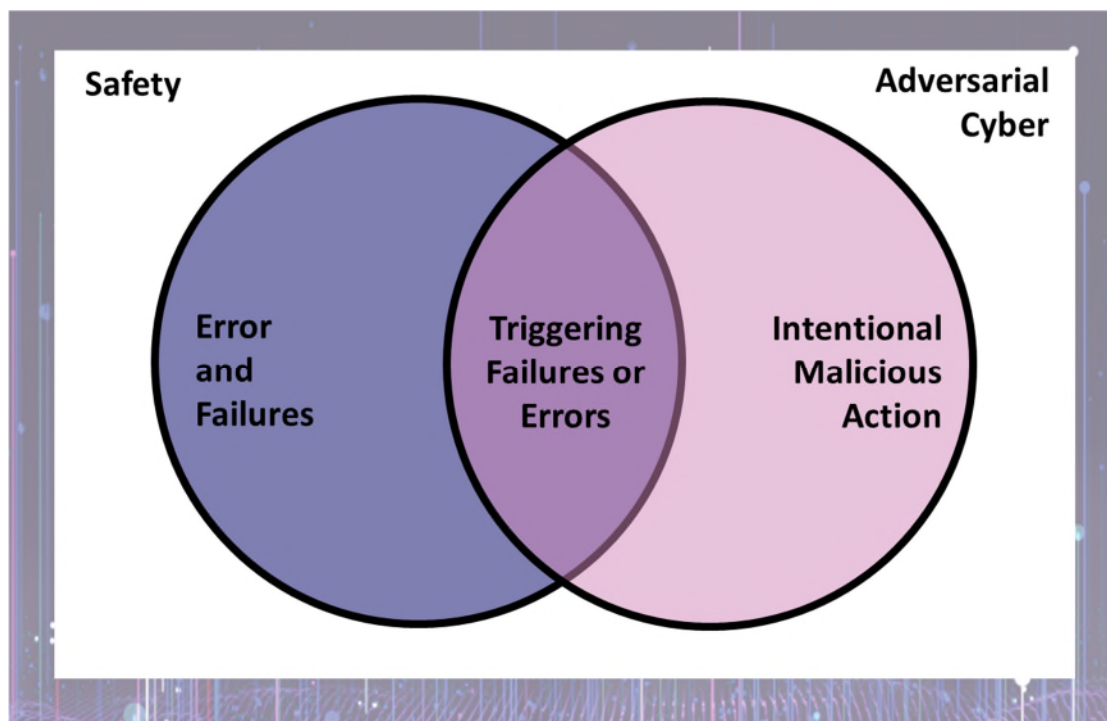


Figure 6. Safety and Adversarial Cyber

Safety focuses on eliminating error and failures. That is a noble and worthy cause that is foundational to cyber resiliency. However, sometimes safety analysis assumes that certain events cannot happen in combination or relies on interlocks to prevent unsafe actions. Unfortunately, the instances where those assumptions are in place are potentially fertile areas for triggering cyber effects that are desired by a malicious threat actor.

It should also be noted that some cyber-attacks that impact an aircraft can emanate offboard from the external environment. For example, spoofed GPS signals confuse the aircraft and the pilot about their actual location. That could be potentially devastating when using instrument landing systems under low visibility conditions. Cyber resiliency can be negatively impacted if the aircraft uses data from a compromised external source without validation.

The holy grail for intentional malicious action by a cyber adversary is a flaw in the system that can be externally triggered on demand. That vulnerability could be something that was latent to the system when it was designed and developed. More recently, there has been a rise of the concern that cyber adversaries could inject those types of flaws in a supply chain attack. A very real example of how that type of injection is known to have occurred is in “The Untold Story of the Boldest Supply-Chain Hack Ever,” an article recommended in Section 2.9.

The information in this section is provided not to sow fear, uncertainty, and doubt, but rather to instill vigilance in every stakeholder in the aviation industry. Every producer and supplier of parts and systems needs to understand the specific nature of threats against aircraft. The same holds true for the aviation operators.

2.7.2 Understanding Cyber Threat Agents against Aircraft Systems

One popular idea in cyber threat intelligence is that an organization needs to understand the capabilities and motivations of the threat agents who are likely to act against their system. While that can be useful information in some situations, a paucity of historic information of that nature for aviation systems can sometimes be a distraction.




Figure 7. Unknown Factors of Cyber-Attack Probability

NIST 800-30, “Guide for Conducting Risk Assessments” directs practitioners to determine the probability of exploitation when determining the risk of a given vulnerability. (10) The focus on probability is directly derived from best practices in the safety domain.

As an aircraft is operated over a long period of time, the predictions of rates of failures and errors becomes more accurate because it is based on historical data. These calculations work best in the absence of external variables outside the scope of the system.

To determine the probability of a cyber-attack using traditionally prescribed methods, a risk analyst requires insight into the intent, capabilities, and likely targeting of malicious cyber threat actors. It is simply not possible to make an accurate determination of probability when an autonomous threat actor has full agency in choosing the time of attack.

Well-intentioned people attempting to follow standard risk management processes sometimes become fixated or paralyzed seeking probability data that does not exist. To exacerbate the situation, the scarcity of threat intelligence for the aviation industry only increases the uncertainty.

Everyone in the aviation industrial base should be working under the assumptions outlined below. These are equally valid for both commercial critical infrastructure aircraft as well as military platforms.

- **Malicious cyber adversaries will one day act with the intent of inflicting the maximum amount of harm on aviation systems.**
- **Highly capable, well-funded nation state actors will one day launch cyber-attacks against aviation systems.**
- **All aviation systems will one day be the target of malicious adversarial action.**

Given the current threat landscape, it is incumbent on the manufacturers and operators of aviation systems to eliminate vulnerabilities and maximize cyber resiliency throughout the development and operational lifecycles. To do otherwise is simply irresponsible.

In fact, if the aviation industry waits to act for evidence of adversarial action that reveals the intent, capability, and targeting of malicious cyber threat agents, it will be too late. The catastrophe will have already occurred.

2.7.3 The Ever-Increasing Capabilities of Cyber Threat Agents

Cyber Risk Assessment policy and standards documents inevitably describe discrete tiers of threat agent capabilities. Each tier is populated with definitions and characteristics of the adversary threat level. Table 2 is one such example that was published by the Government Accountability Office. (GAO)

This information is intended to support the scoring of cyber risk by first identifying the likely threat actors against a system, and then estimating the capabilities of that organization. That exercise isn't a terribly good use of time and resources for aviation systems.

Aviation platforms have extraordinarily long service lifespans. Meanwhile, the rate of change within the adversarial threat landscape is rapid. The capabilities required to be categorized as the most advanced adversarial threat level is ever increasing. As the advanced tradecraft is revealed through ongoing operations, individuals and organizations learn how to perform at an increasing level of capability.

Once nation state level zero-days are used and detected, information about the vulnerabilities and exploits used are exposed and published. That makes it more accessible to threat actors who were previously scored at lower tiers. Additionally, the published exploits, tools, and scripts result in the constant erosion of the skill required to successfully attack a vulnerability. In other words, yesterday's sophisticated nation-state attack eventually becomes script-kiddie fodder over time.

Threat Level	Description
Advanced	May conduct complex, long-term cyber-attack operations that combine multiple intelligence sources to obtain access to high-value networks. May develop detailed technical and system knowledge of the target system to deploy more damaging cyber-attacks.
Moderate	Able to use customized malware to conduct wide-ranging intelligence collection operations, gain access to more isolated networks, and in some cases creates limited effects against defense critical infrastructure networks.
Limited	Able to identify—and target for espionage or attack—easily accessible unencrypted networks running common operating systems using publicly available tools. Possesses some limited strategic planning.
Nascent	Little-to-no organized cyber capabilities, with no knowledge of a network’s underlying systems or industry beyond publicly available open-source information.

Table 2. Adversarial threat tiers described in GAO-19-128. (11)

Aviation systems must be designed, operated, and sustained to anticipate the ever-increasing capabilities of threat agents along with the continuous erosion of the skill level required for exploitation. It only gets easier over time.

Consequently, it is important for the developers and operators of aviation systems to bring an inherent understanding that cyber threats against aircraft are dynamic and only get better and better all the time. The best way to defend a system is to deliver with no latent vulnerabilities at all.

2.7.4 Emerging Aviation Vulnerabilities (Zero Days)

Developers and operators of aviation systems can do absolutely everything right based on current cyber resiliency engineering and cybersecurity practices when an aircraft is under development, yet still have latent security issues or vulnerabilities that are discovered at a future time. That can happen through the emergence of a new class or type of cyber-attack that was unknown when the system was designed and developed.

A strategy of building a perfectly unhackable system is arguably naïve. Failure to design security update and sustainment mechanisms for aircraft systems cannot be justified with the idea that development was so perfect that no future issues will arise. At the same time, that reality is not an excuse for failing to use industry best practices to minimize vulnerabilities in the platform to the greatest extent possible.

2.8 Summary

This chapter has outlined the reasons why one of the most dangerous risks to the security of aviation systems is simply a failure to recognize the threats. Inaction could result in loss of mission performance, the aviation hardware, and even human life.

All aviation platforms, both commercial and military, currently operate or will certainly operate in cyber contested environments in the future. During times of conflict, our adversaries can be

expected to deploy offensive cyber-attacks intended to generate a decisive advantage in the battlespace and on our critical infrastructure systems.

The entire industrial base for aviation systems must recognize and respect the clear and present threat of adversarial cyber actions against our systems. The unique challenges of securing aircraft must be considered each and every day. The threat is real.

The people who acquire, develop, operate, sustain, and maintain aviation systems cannot ignore the very real cyber threats that exist against our platforms. Any cyber logic-bearing component that interacts with the external world, either through direct communications interface or other environmental mechanisms, is a potential vector of attack on aircraft systems.

From the largest most complex LRUs down to the smallest logic-bearing component, we all share a collective responsibility to recognize the potential cyber threats against aviation systems and work in concert for effective elimination and mitigation.

2.9 Additional Reading

Anyone who remains unconvinced that the cyber risks outlined in this section are real, or for those that are interested in more information, the following resources provide substantially more detail.

- **The Untold Story of the Boldest Supply-Chain Hack Ever**, Kim Zetter, Wired Magazine, May 2, 2023. (12)

A detailed account of how a state sponsored cyber threat organization compromised products at SolarWinds. That company's products were subsequently used to successfully attack many of the firm's clients, including several top-tier government agencies. This story should serve as a wake-up call to the potentially devastating impacts of supply chain injections.

- **Countdown to Zero Day: Stuxnet and the Launch of the World's First Digital Weapon**, Kim Zetter, Crown, November 11, 2014. (13)

This book traces the origins and development of Stuxnet, a sophisticated cyber campaign believed to be launched by the United States and Israel, against the Iranian nuclear program. This book reveals how malicious cyber-attacks can cause significant physical damage to cyber physical systems. It also serves as a cautionary tale for those who believe that air-gapped systems are impervious to this kind of attack. This book also illustrates how a nation state zero-day attack detected in the wild became widely accessible to less skilled threat actors once the conceptual vulnerability was understood.

- **The Colonial Pipeline Hack Is a New Extreme for Ransomware**, Andy Greenberg, Wired, May 8, 2021. (14)

This article outlines the significant impact that a ransomware attack had on the Colonial Pipeline. Though it is unlikely that critical infrastructure was intentionally targeted in this

attack, the widespread disruptions when the pipeline was shut down underscores the potential impacts of attacks on the critical infrastructure.

- **Hackers Remotely Kill a Jeep on the Highway – With Me in It**, Andy Greenberg, Wired, July 21, 2015. (15)

This article describes the terrifying impacts of remote cyber-attacks against a motor vehicle. It also includes a story of what happened when a known vulnerability was dismissed as not a concern until widespread exploitation started to occur with serious consequences.

- **Sandworm: A New Era of Cyberwar and the Hunt for the Kremlin's Most Dangerous Hackers**, Andy Greenberg, Doubleday, November 5, 2019. (16)

This is a detailed history of a series of cyber-attacks targeting utility companies and the electric grid in Eastern Europe. The paralyzing effects of cyberwarfare described ultimately impacted multiple industries and snarled shipping on a global scale. At the attack's epicenter in Ukraine, ATMs froze, the railway system shut down, the postal systems ceased to operate, and hospitals went dark. These attacks on critical infrastructure reveal the widespread impacts of cyber warfare.

- **This Is How They Tell Me the World Ends: The Cyberweapons Arms Race**, Nicole Perlroth, Bloombury Publishing, February 9, 2021. (17)

This book is about the secretive market for cyberweapons and zero-day exploits. It describes how various organizations are discovering, buying, selling, and hoarding zero-day vulnerabilities until they are needed for a cyber campaign.

Chapter 3

Our Collective Cyber Responsibilities

All stakeholders in the aviation industry must take an active role in the protection of aircraft systems against adversarial cyber-attacks. This is equally important for both military and commercial aviation platforms. The ability to project military power is essential for the defense of national interests and commercial aviation is a vital part of the nation's critical infrastructure.

Chapter 2 provided an overview of the malicious cyber threats faced by our aircraft systems. Maintaining resiliency against an ongoing and future onslaught of attacks requires the support and diligence of every stakeholder in the supply chain and operational environments. The cyber threats against aircraft platforms must be taken seriously.

Hostile cyber-based attacks against both commercial and military aviation systems is a significant threat to national security. It is incumbent on the designers and manufacturers of these systems to build aviation platforms that are not vulnerable to cyber adversity. In addition, aircraft must be resistant to errors and accidents that can lead to failures, loss of platform, or even loss of human life.

This chapter outlines the proactive attitudes and behaviors necessary for effective cyber defense of aviation ecosystems. Every stakeholder must embrace the necessary mindset required to build a proactive cyber defensive posture.

3.1 A Positive Cyber Culture

Peter Drucker, the highly regarded management consultant, is often credited with coining the phrase "Culture eats strategy for breakfast." However, according to the Drucker Institute, a more accurate quote is "Culture—no matter how defined—is singularly persistent." (18) Cyber initiatives and security controls cannot be effective without complete buy-in from absolutely everybody involved in the aviation industrial base.

At the core of the proactive cyber culture is an inherent understanding and acceptance of the threats faced by our aviation systems. However, it is equally important that all people with a role in the aviation industry also make a commitment to implement the best practices required to make that a reality. Cyber defense must be an absolute priority.

The International Civil Aviation Organization (ICAO) defines security culture as follows (19):

Security culture is a set of security-related norms, values, attitudes, and assumptions that are inherent in the daily operation of an organization and are reflected by the actions and behaviors of all entities and personnel within the organization.

A more succinct way to rephrase that definition is that the security culture is a combination of what an organization says it does, as embodied by what the organization actually does. It isn't enough for everybody to say that cybersecurity is valued by the organization. That must be backed up by actions every single day.

3.2 Transitioning to a Positive Cyber Narrative

In any organization, the stories that are told are a fundamental reflection of the culture. In fact, stories can be used to build, reinforce, and fine-tune the cultural norms. These stories help everybody in the aviation industry understand what is valued in the organization by telling them how to behave.

For many organizations, the culture is built on a strong history of prioritizing cybersecurity. For others, a cultural transformation is required. Issuing memorandums or sending personnel to annual security compliance training is not likely to permanently move the needle to where the organization needs to be.

The stories we tell ourselves and each other have a profound influence on the culture. The following subsections outline the themes of the stories that need to be told and retold within the entirety of the aviation industrial base.

3.2.1 Cybersecurity is everybody's job, including mine.

Some large organizations have personnel assigned to job roles that are responsible for cyber. Smaller organizations are less likely to have people assigned to that narrow role. Regardless of whether there are cyber specific jobs and resources within an organization, nobody can delegate their personal responsibilities for implementing cyber to other people.

Cybersecurity is everybody's job. That means that every one of us need to consider how the products and services we are providing contribute to, or potentially detract from, the cybersecurity posture of the aircraft system. Individuals cannot ignore cyber because it is someone else's job. The reality is that it is all our jobs.

Cyber specialists in an organization, or the broader industry, are available to provide insight and guidance. In many cases those experts can be brought in to assist in formulating things like cyber resilient designs or mitigations to advanced threats.

However, cybersecurity starts at the grassroots level. In essence, we are all cybersecurity personnel. The stories we tell must reinforce that expectation within the culture of the aviation industrial base.

3.2.2 Cybersecurity is essential to everything we do.

It is sadly not uncommon to uncover attitudes that resources and time spent on cybersecurity comes at the expense of features or capabilities for aircraft systems. Regarding cybersecurity as a waste of resources is one of the most dangerous ideas within the aviation industry.

In the consumer market, it is not uncommon to find indifference in buyers toward security features. In many cases, people are unwilling to pay more for goods and services with higher cyber assurance levels. That is a societal problem beyond the scope of this FSAD guidebook.

However, in aviation systems, the consumer and buyers must understand the importance of safety. Cybersecurity risks and vulnerabilities represent genuine safety concerns for both the aircraft and human life. Cyber resiliency should be at the core of everything we do.

Another way this dangerous idea can assert itself is via legacy oriented thinking. If a product or system has never had requirements for cybersecurity controls or analysis in the past, adding those things to updates and revisions can seem like gold-plating or extraneous work. However, our awareness of the risks associated with cyber vulnerabilities has changed. Our current understanding makes cybersecurity an essential requirement for doing business in the aviation industry.

The stories that we tell ourselves must be built on the theme that cybersecurity and resiliency are at the core of everything we do. The safety and security of the aviation ecosystem are reliant on embodying that within the culture.

3.2.3 When we see something, we say something.

In physical security, the principle of “see something, say something” encourages individuals to report suspicious activities or potential threats. The same vigilance is required in cyber. The aviation ecosystem will be stronger when we all exercise that principle.

Cybersecurity awareness and responsible online behavior are essential for safeguarding ourselves and the aviation systems that we develop. IT related security events in development systems can be a potential indicator of much more nefarious actions against aircraft platforms, as compromises and information gleaned from those environments can help an advanced cyber adversary formulate future attacks.

That means that it is important to report unusual or suspicious cyber events observed in the aviation industrial base. This includes phishing attempts, unusual network activity, or inconsistencies in software system behavior.

Gone are the days of assuming that cyber events are probably nothing. Cyber compromises of stakeholders or suppliers with seemingly minor roles in the aviation ecosystem could be significant.

When it comes to cybersecurity in the aviation industry, when we see something, we say something. The stories that we tell ourselves and our organizations must reinforce that belief.

3.2.4 How can we prevent that from happening here?

The news and media are increasingly full of stories about how companies, schools, and government organizations have experienced crippling cyber adversity. Ransomware attacks are one current example of widespread malicious activity against systems. How organizations respond to those reports reveals a lot about their culture.

Adopting a “this could never happen here” attitude is not constructive. A better culture is one where everybody receives news of recent cyber-attacks as an opportunity to think through what actions could be taken to prevent a similar compromise from occurring in our own systems.

Another dangerous variation in this genre is a belief that a company is so small or insignificant that they would never be targeted. Unfortunately, gaining access to a supplier system and using that foothold to propagate to other companies and organizations in the supply chain is a common tactic used by the APT described in Section 2.4.1. No one should harbor the belief that they are not, or will not be, targeted by malicious cyber actors. It is really only a matter of time.

Sometimes ransomware and other general forms of cyber-attack against a developer or supplier’s IT systems are not regarded as a direct concern to the aviation systems they develop. Unfortunately, should a successful attack of that nature occur, it is an issue. An organization that is susceptible to ransomware cyber-attacks is most certainly at risk of more advanced adversarial cyber-attacks. Garden variety enterprise IT malware is like a canary in a coal mine.

A positive cyber culture does not pass up on opportunities to learn from cyber misfortunes regardless of the source. The stories that successful organizations tell themselves is how they leverage learning opportunities and the experiences of others to prevent similar cyber adversity from occurring on their systems.

3.2.5 We don’t punish cyber incident or vulnerability disclosure.

One of the most significant cultural challenges for the cybersecurity industry in general is creating an environment where people and organizations that disclose cyber events and vulnerabilities are sometimes punished. When an institution is vilified for cyber breaches it creates the temptation to sweep notifications and disclosures under the rug.

When cyber events are not disclosed or hidden, other organizations with similar weaknesses or vulnerability do not have the opportunity to learn from those mistakes. Worse, indicators of attack or compromise aren’t shared which could prolong the impact or severity of a cyber intrusion on the systems of another organization.

At the local level, organizations should not punish self-reported security incidents that are the result of mistakes or errors. That kind of transparency is required to properly assess and mitigate any negative consequences of what occurred. The security posture is stronger when incidents are dealt with openly and directly. That must necessarily come with empathy and understanding that mistakes happen.

Additionally, self-reported cyber security incidents should be evaluated for systemic changes that could prevent future successful cyber-attacks or intrusions. Open and honest sharing of that information ultimately makes the entire aviation industrial base stronger.

3.3 Success Factors of a Healthy Aviation Cyber Culture

In addition to the stories that build and shape the cybersecurity culture for all stakeholders in the aviation industry, a core set of characteristics are evidence of a healthy organization. This section describes the success factors that must be instilled in all corners of the aviation industrial base.

3.3.1 Building Cybersecurity In

The critical nature of aviation platforms makes building security into products and services from the earliest stages of the development process of paramount importance. Due to the nature of the embedded cyber physical systems, cyber resiliency is most effective as an inherent attribute of the underlying system.

Cybersecurity and cyber resiliency are characteristics that cannot be bestowed on a system during the final stages of the development lifecycle. Cyber test and evaluation can only prove the existence of vulnerabilities. It can never provide sufficient evidence for an absence of latent issues.

Project plans, schedules, and resource allocation includes cybersecurity analysis at every step of the process. Those planning and management resources will track specific activities that contribute to the cybersecurity posture of the system.

Subsequent chapters of the FSAD provide more implementation guidelines and best practices for organizations developing complex hardware and software for aviation cyber physical systems. However, all stakeholders must be aware of the importance of building cybersecurity into every part and system developed and delivered.

3.3.2 Compliance and Security Control Requirements

All participants in the aviation industry likely have some form of cybersecurity compliance and security controls imposed on them. That includes the cybersecurity clauses and requirements in the Defense Federal Acquisition Regulation Supplement (DFARS) described in Section 4.5 as well as the Assessment and Authorization (A&A) process defined in Chapter 12. How organizations regard these compliance and security requirements is a cultural success factor.

The aviation industry's most mature organizations will embrace compliance and security control requirements as a beneficial structured approach for minimizing risk to both development environments and deliverable systems. Rather than an imposition, these organizations understand that compliance requirements provide succinct criteria that documents a good faith effort to be secure.

At the same time, successful organizations also recognize that security compliance controls represent the minimum baseline actions required. They understand that systems that do not meet the minimum compliance standards are certainly not cyber resilient. High performing organizations embrace the idea that going above and beyond is frequently necessary. Depending on the nature and the complexity of the aviation system, additional features or functions

supporting the cybersecurity posture of the development environment and deliverable products may be required.

The most sophisticated aviation industry organizations embrace rather than resent any compliance and security controls that are allocated to them. They understand the importance of meeting the minimum baseline requirements and look for opportunities to do more.

3.3.3 Accountability and Responsibility

There is an adage in management consulting that a task that is everyone's job is actually no one's job. The fact that cybersecurity is everybody's job is a recurring theme throughout this FSAD guidebook. However, those statements are made with the understanding that someone must be appointed to make sure that cybersecurity best practices are diligently applied and followed within each organization.

While everybody in the aviation industry must understand that cyber resiliency is one of their job duties, they need to know where and who to reach out to when advanced assistance in the cyber domain is needed. Additionally, every organization must have a designated person or team that is responsible to ensure that cybersecurity plans are maintained and followed.

3.3.4 Management Commitment

Advocacy for aviation cyber resiliency is critically important from the leaders at the highest levels of each organization in the aviation industrial base. While Chapter 16 provides additional details written specifically for executive management, all other members of the organization must be aware of what to expect from their leaders as a collective accountability mechanism.

Executive advocacy is necessary for organizations to allocate enough resources, including budgets and people, to effectively implement cyber security initiatives. If those leaders do not place any importance on the security posture of the development systems and products, then the organizations and people will follow suit.

Executive leadership advocacy sets the foundation for a cybersecurity-first culture. It is almost impossible for the right things to happen from the grassroots level without top-level support.

3.3.5 Cyber Risk Assessment and Managing Cyber Risks

Cyber Risk Assessment (CRA) is a systematic process of identifying, evaluating, mitigating, and in some cases accepting risk in a system. Chapter 9 goes into greater depth on that topic for people directly responsible for CRA performance. However, everybody in the organization must be aware of the importance of continuously and rigorously performing those assessments.

CRA is a best practice for both development systems as well as products that are developed for aviation platforms. For systems that require 3rd party cyber assessment and authorization, CRA is a fundamental part of achieving approval. For example, it is an integral part of the Risk Management Framework (RMF) which is the current authorization standard for military systems.

Regardless of the process used, mature organizations will have a system for continuously identifying and managing risk in place. Otherwise, the team is flying blind.

3.3.6 Vigilance of Potential Cyber Events / Continuous Monitoring

Organizations must diligently monitor their systems for evidence of potential cyber compromises. One powerful reason for that is because major disruptive cybersecurity events typically start with small incursions. Monitoring allows small cybersecurity issues to be detected and corrected before turning into large scale problems.

Additionally, some systems are most valuable to a hostile adversary as an access vector into other systems or as a conduit for launching other attacks. For example, the threat actors who perpetrated the SolarWinds attack described in Section 2.9 used that system to inject vulnerabilities into the supply chain. The presence of malicious threat actors in the SolarWinds development system was not detected until the downstream injected exploits were discovered in their customer's networks.

Organizations that are not looking for cyber incidents on their networks are unlikely to find any. That should never be confused with an absence of cyber compromises. Continuous monitoring for potential cyber events is an imperative.

3.3.7 Incident Response

Effective response to cyber security incidents is necessary for maintaining a secure supply chain for aviation systems. Adversarial attacks and mishaps are becoming more frequent, damaging, and disruptive over time. It is the responsibility of all organizations to recognize and respond appropriately when these events occur.

At the present time, cyber incidents cannot be completely prevented. Consequently, an effective incident response capability is an imperative for all organizations. Additionally, it is important to report and share information about new and novel cyber events promptly as they occur. That information could prevent recurrence of a successful cyber-attack against other parts of the aviation system supply chain.

The planning and resource allocation for effective incident response is outlined in NIST Special Publication 800-6, "Computer Security Incident Handling Guide." (20) That document is a great starting point for creation of an organizational incident response plan.

3.3.8 Security Education, Awareness, and Training

All stakeholders in the aviation supply chain are responsible for cybersecurity. That means recurring security awareness education and training is necessary, so everybody understands what that entails. Additionally, users that have privileged or administrative access on systems require additional training above and beyond the basic information that is provided to all stakeholders.

For personnel who work on aviation systems, training is required for both the enterprise IT development networks as well as the domain of products and services delivered into the aviation

supply chain. That may include cybersecurity practices specific to the embedded cyber physical systems used.

3.3.9 Continuous Improvement and After-Action Reviews

The ever-increasing capabilities of malicious cyber threat actors was described in Section 2.7.3, and the almost certain emergence of new classes of vulnerabilities was shared in Section 2.7.4. That places the onus on organizations to exhibit a culture of continuous improvement for cybersecurity. The lessons learned should be used to update and guide local policies and procedures.

When a cyber event occurs, mature cyber cultures will invest in understanding what happened and how it could be better contained or prevented in the future. These after-action reviews are focused on lessons learned rather than assigning blame.

The most effective cyber security cultures learn from mistakes and reflexively leverage every opportunity to get better. These organizations inherently have a culture of continuous improvement that supports and enhances their security practices and posture.

3.3.10 Cyber Considerations for Purchasing Equipment, Software, and Materials

The principle that cybersecurity is essential to everything we do was outlined in Section 3.2.2. That includes procurement and purchasing decisions. While this topic is covered in much greater detail in Chapter 15 of this FSAD guidebook, this is also a cultural characteristic in addition to a supply chain management best practice.

The enterprise IT systems used in development and production of products and services consider cybersecurity in the acquisition process. Software, hardware, and services purchased for inclusion in aviation systems are vetted for secure operation. Often this means purchasing a more expensive product with a substantiated security pedigree rather than the lowest cost solution.

It should be noted that there is currently a dearth of provably secure products viable for use on aircraft platforms. Until that changes, organizations must ensure that the equipment, products, software, and materials used in aviation systems are obtained from reliable sources that also pay attention to the cybersecurity posture of their products. Aircraft operators and their customers are dependent on the aviation industrial base to select and use the most cyber resilient products available.

3.4 Summary

The proverb that a chain is only as strong as its weakest link applies to cybersecurity. (21) Consequently, the idea that cybersecurity is everybody's job is critically important to the aviation industry not only to protect ourselves, but also our customers and peers. A single lapse at one supplier can set off a cascading series of events that could have catastrophic impacts to aviation platforms.

All stakeholders in the aviation industry must take an active role in the protection of our aircraft systems against adversarial cyber-attacks. Hostile cyber-based actions against both commercial and military aviation systems is a significant threat to national security. It is incumbent on the designers and manufacturers of these systems to build aviation platforms that are not vulnerable to cyber adversity. In addition, those systems must be resistant to errors and accidents that can lead to failures, loss of platform, or even loss of human life.

While proactive attitudes and behaviors are not singularly sufficient for successful cyber defense of aviation aircraft systems, failure to adopt the proper mindset can cripple the most well-intentioned technical efforts.

The people who acquire, engineer, develop, operate, maintain, and sustain aviation systems must embrace the necessary attitude required to protect the production development systems as well as aviation system products and services.

Chapter 4

Securing the Development and Production Infrastructure

Secure aviation products rely on a strong cybersecurity posture in our underlying development and manufacturing systems. Consequently, those devices are subject to adversarial cyber-attack from advanced nation state threat actors. The network connected business computing assets operated by every supplier in the aviation supply chain represents exploitable attack surface in the eyes of the adversary.

Additionally, industrial control systems (ICS) and cyber-physical systems (CPS) in our manufacturing and supplier networks cannot be overlooked as potential access vectors for highly skilled cyber adversaries.

There is no universal solution for identifying, managing, and mitigating cyber threats and risks to our development and manufacturing systems. Organizations will have unique equipment, configuration, implementations, and practices which can impact the resilience and operation of their private networks.

However, everyone in the aviation industry shares one thing in common: Highly skilled advanced nation state adversaries are targeting our business systems that play a vital role in the foundation of our country's critical infrastructure. It is our collective duty and responsibility to protect our business and production systems, which in turn protects the aircraft we develop.

4.1 The Aviation Supply Chain and Critical Infrastructure

In 2023, Microsoft detected what is believed to be malicious code in telecommunication systems in Guam. (22) The Pacific island is home to two strategic US military bases, so disruptions in communications on the island could be tactically significant should tensions in that region escalate. There is considerable speculation that critical infrastructure attacks may be a significant part of future warfare. In fact, the book "Sandworm: A New Era of Cyberwar and the Hunt for the Kremlin's Most Dangerous Hackers," recommended in Section 2.9 outlines how Russia used critical infrastructure attacks against Ukraine long before the Russian invasion of 2022. (16)

The aviation supply chain both depends on, and is a part of, our national critical infrastructure. Organizations should have contingency plans in place for continuity of operations if power, water, transportation, and other essential services were disrupted for an extended length of time.

Additionally, a more surgical attack against military systems could be instigated by cyber threat actors against the aviation supply chain. Disruptions or destruction of production and sustainment capabilities could be distracting during the "fog of war." It could even be tactically significant.

In recognition of that ever-present risk, this chapter describes the baseline best practices that all organizations participating in the aviation supply chain must take to protect their own computing systems.

4.2 Risk Management and Critical Infrastructure

Every organization must implement a cyber risk management process and cybersecurity program that is appropriate and tailored for its business or production systems. From the smallest manufacturer, that may produce simple metal brackets, to large scale system integrators, each organization and team must take prudent steps to secure their systems.

The National Institute of Standards and Technology (NIST) has developed a cybersecurity risk management framework that outlines how to identify, scale, and implement cybersecurity programs for critical infrastructure operators. The Framework for Improving Critical Infrastructure Cybersecurity is a voluntary guide to help operators of critical infrastructure assets identify, assess, and manage cyber risks. (23)

The framework recommends five categories of cybersecurity functions. Operators of critical infrastructure must **Identify** the assets that require protection in their environment. Steps must be taken to **Protect** those systems. Organizations should also put mechanisms in place to **Detect** potential cybersecurity events. When necessary, organizations must **Respond** to cyber incidents. Finally, they must have methods to **Recover** normal operations.

The remainder of this chapter focuses on the baseline actions that all individuals and organizations must ensure are in place to protect the infrastructure. These practices universally apply to everybody who is a provider or stakeholder in the aviation supply chain.

Members of the aviation industrial base that provide more complex cyber-enabled systems have responsibilities beyond the basics outlined in this guidebook. It is recommended that all organizations familiarize themselves with the Framework for Improving Critical Infrastructure Cybersecurity and make a documented determination of the extent to which the recommended controls apply to them. (23)

We all have a key role in protecting the critical infrastructure that supports our aviation platforms. Securing the infrastructure is a fundamental requirement for doing business in this industry.

4.2.1 Identify

Successful defense requires a clear understanding of what is being defended. Members of the aviation supply chain must develop and maintain a detailed inventory of their cyber assets used to produce products and services for aviation systems. That should include physical devices such as computers, printers, and other peripherals. Furthermore, software and applications must also be inventoried for configuration management.

It is important that risk management plans identify both tangible and intangible assets. Tangible assets have a physical existence which includes things like computer servers, land, and buildings.

While intangible assets don't have a physical existence, intellectual property, reputation, and business plans can be critical to the future success of the company. Data and software are also intangible assets.

All participants in the aviation industrial base must understand how their business, development, and manufacturing systems are networked. That includes the data flows necessary to support conducting daily operations.

4.2.2 Protect

Once the tangible and intangible assets to be defended are identified, steps and strategies must be implemented to protect those systems. At a minimum, the cyber hygiene practices identified in Section 4.4 must be implemented for all computing assets used to produce aviation system products and services.

4.2.3 Detect

The cultural imperative that organizations must monitor their systems for evidence of potential cyber events was outlined in Section 3.3.6. Failing to look for cyber incursions does not prevent them from happening and in most instances exacerbates the severity when they occur. Detection mechanisms are absolutely required to protect production and development systems.

Many of the cyber hygiene practices identified in Section 4.4 are intended to support detection and forensic analysis of potential cyber events.

4.2.4 Respond

When cyber-attacks occur in aviation development and production systems, the organization must be prepared to properly respond. First and foremost, the attack must be stopped and contained. Once that occurs, damage assessment must be performed including any downstream impacts of the attack. That could include exfiltration of critical data or potential injection of vulnerabilities or backdoors into development products or systems.

Once the cyber event is contained, it is crucial to analyze it for identification of controls and implementations that might prevent the same type of attack from succeeding again in the future. Additionally, sharing the findings with the vendor and broader community, when appropriate, is essential to protecting other participants in the aviation supply chain, since they are likely to be similarly targeted.

4.2.5 Recover

Finally, every organization must have a backup and restoration plan in place that supports returning to normal operation. Additionally, that plan must be executable within timelines that do not disrupt peer suppliers or customers. The backup and restoration plan should be exercised periodically to ensure that recovery can be successfully performed when necessary.

4.3 ICS and SCADA Systems

Some manufacturers in the secure aviation design supply chain may leverage Industrial Control Systems (ICS) and/or Supervisory Control and Data Acquisition (SCADA) systems within their production environment. These systems are used in industrial settings to provide and regulate things like electricity, water, and fuel. ICS and SCADA systems are used in manufacturing of components for aerospace and aviation systems, as well as parts for several related and adjacent industries.

ICS and SCADA systems were once thought to be unsusceptible to cyber vulnerabilities and security concerns that are traditionally associated with enterprise IT systems. The idea was that proprietary control protocols, specialized software, and custom hardware were less likely to have exploitable weaknesses than more ubiquitous business systems. Additionally, ICS and SCADA systems were regarded as physically isolated and rarely connected to IT networks. Those assumptions and beliefs were never valid.

Network connected ICS and SCADA systems are now widely used in the manufacturing environment. That opens the possibility of exploitation of vulnerabilities and incidents in that equipment. Additionally, ICS and SCADA systems may be connected to business operational networks, which potentially further exposes those systems to remote cyber-attack and incursion.

Unfortunately, commercial off the shelf (COTS) cybersecurity solutions are not as mature for ICS and SCADA components than they are for enterprise IT systems. That leaves this important part of the critical infrastructure at an elevated level of risk. Special precautions must be taken when securing ICS and SCADA systems.

New security solutions are needed that are tailored specifically to cyber physical systems. Cyber-enabled logic executing in ICS has a direct effect on the physical world. Consequently, it could have a profound impact on human health and safety. Adversarial exploitation of ICS and SCADA systems could damage the environment.

Additionally, successful attacks on cyber physical systems could cause them to literally self-destruct, or worse. The book “Countdown to Zero Day: Stuxnet and the Launch of the World’s First Digital Weapon” recommended in Section 2.9, provides a detailed account of how a cyber-attack crippled the Iranian nuclear program.

Production systems that include ICS and SCADA components should be specifically addressed in the overall cybersecurity programs at every organization. Furthermore, due to the essential nature of these systems to the manufacturing process, the sources of threats against these systems cannot be understated. Adversarial events can emerge from nation state cyber actors, terrorists, and disgruntled insiders. Additionally, accidents and natural disasters could compromise the safety and security of these systems.

NIST Special Publication 800-82, “Guide to Industrial Control Systems (ICS) Security” provides more guidance and information for establishing and operating secure ICS and SCADA systems. (24) All stakeholders in the aviation supply chain should be cognizant of the potential for threats and risks against these systems and secure them appropriately.

4.4 Enterprise IT and Cyber Hygiene

Cybersecurity experts universally extol the virtues of good “cyber hygiene” practices. Ironically, there is no single authoritative source defining exactly what that means. At the personal level, hygiene is the term for conditions or practices conducive to supporting health and preventing disease. It is generally associated with cleanliness.

In the cyber domain, hygiene refers to the practices computer owners, administrators, and users adopt to maintain the safety and security of their systems and data. It keeps those systems healthy and supports prevention of cyber events. In essence, cyber hygiene is maintaining systems in a state that supports systemic health.

There is widespread belief that if everybody practiced good cyber hygiene that the majority of malicious cyber-attacks would be prevented. It would certainly make it more difficult. However, as outlined in Section 2.7.3, cyber-attacks and the skill of threat agents get more sophisticated and stronger over time. As a consequence, good cyber hygiene practices are a continuous process of improvement.

Cyber hygiene alone is not sufficient to completely prevent successful cyber-attacks against a system. However, when good cyber hygiene is in place, the probability and consequences of cyber adversity is reduced. From a business standpoint, it reduces costs and downtime associated with recovering from attacks.

The following sections describe good cyber hygiene practices at the conceptual level.

4.4.1 Practice Defense in Depth

Defense in depth is a cybersecurity philosophy that embodies the idea that multiple layers of security are better than relying on any single control for protection. This strategy uses a variety of mechanisms to protect computers, mobile devices, peripherals, networks and data. Each layer provides a barrier that adversarial threat agents must overcome to achieve their objectives.

The defense in depth approach is intended to minimize the success of cyber intrusion through a combination of preventative, detection, and response mechanisms. Those controls are distributed in and around the critical computing assets of an organization. When cyber adversaries successfully defeat one mechanism, other measures in place can detect the attack and contain the impact.

An organization that embraces the defense in depth philosophy is in a better position to prevent, detect, and contain cyber-attacks than one that does not. It is a critical mindset when protecting the aviation system industrial base.

4.4.2 Maintain the Inventory/Configuration Management

Configuration Management (CM) is performed by organizations to establish a baseline understanding of enterprise assets. CM is also used to control modifications and updates to documented baselines. Organizations with a robust CM process will understand what is on their system as well as what precipitated updates and who implemented the change.

CM is an essential information security practice required to secure the critical infrastructure. While not every change and update to an IT system will be implemented due to security considerations, every update could potentially have security implications.

For example, a rarely used software application might be installed on a computing asset. If a security issue arises with the software application that requires urgent patching or update, those actions are unlikely to be performed if no one is aware that it is installed. CM provides an authoritative inventory for cross check of emerging vulnerabilities and issues.

Harm can potentially arise from the hardware and software installed on organizational computing assets. The risk can be exacerbated by disruptions, human errors, and intentional adversarial cyber-attack. Every organization participating in the secure aviation design and manufacturing supply chain must manage risk associated with their cyber-enabled critical infrastructure.

An organization that does not have a detailed inventory of its processing assets will be unable to identify and respond to risks associated with those systems. Additionally, CM can also provide a roadmap for restoration to the exact previous state when recovering from an incident. Poor CM practices could jeopardize not only an organization's own operations, but also the customers and suppliers that interact with them.

NIST Special Publication 800-128 "Guide for Security-Focused Configuration Management of Information Systems" provides a resource for organizations to understand their full responsibilities for adequate CM practices within their environment. (25)

4.4.3 Run Antivirus

Antivirus software is a special purpose computer program designed to detect, prevent, and remove malicious software from computing systems. It protects enterprise IT against various types of cyber threats that have been previously detected either through the signature characteristics of the malware or heuristic patterns.

Antivirus software is available from commercial vendors with subscriptions for periodic updates as new types of malware are detected. This software is generally good at detecting and preventing attacks that have been previously encountered.

Antivirus software is less effective against zero-day vulnerabilities, since the signature and pattern matching against previously unknown attacks is not reliable. Nation state level actors may have zero-day exploits that can circumvent antivirus protection mechanisms. However, that is not a valid justification for failing to use antivirus software.

In addition to antivirus, security products for antispyware, personal firewall protection, and host-based information protection systems (IPS) may also be available for the computing platform. This layered defense in depth approach is key to isolating breached systems more quickly.

Additionally, if a nation state level actor can gain access to a system using a well-known exploit without "burning" a valuable zero-day vulnerability, they will most certainly take the less elegant

attack path into a system. Consequently, when it is available, there is no valid reason to not run commercial security software as a cyber hygiene best practice.

4.4.4 Apply Security Patches and Updates

Computer operating systems and applications are complex software products that frequently harbor security flaws. When a vulnerability is detected in that software, the manufacture releases security patches that update it to close or mitigate the particular issue.

As discussed in Section 2.7, software vulnerabilities are sometimes identified via responsible disclosure of a security researcher. However, sometimes vulnerabilities are detected that are already being exploited in the wild.

Additionally, when a security patch is released against a product, malicious actors rush to reverse engineer the update to uncover the underlying vulnerability it corrects. That starts a race between adversarial threat agents, who are seeking to exploit new and emerging vulnerabilities, and cyber defenders who are charged to implement security updates and patches before that can occur.

It is extremely important to expedite the application of security updates once they are released by the manufacturer of the software operating system or application. For enterprise IT systems, there are many commercial products that will automate this process.

Good cyber hygiene dictates automating and prioritizing deployment and installation of security patches to the greatest extent possible.

4.4.5 Physical Access Control

Cyber threat agents with physical access to computing assets have more avenues of attack available to them than strictly networked based attacks. Controlling who has physical access to computing equipment, and under what conditions, is an essential cyber hygiene practice.

Unauthorized personnel should not have unfettered access to computing assets. That means that the physical office environment must be controlled including rooms with wired network connectivity. Additionally, the location of network security devices and servers should only be accessible by a limited subset of authorized administrative users.

Good cyber hygiene means that physical access to computing devices is tightly controlled.

4.4.6 Account Management

Accounts that enable access to enterprise IT assets and other computing devices should only be created for authorized users. Additionally, as personnel depart the organization or change roles, accounts that are no longer needed should be promptly disabled. Deleting accounts immediately is not recommended because that might cause the loss of valuable business or forensic information.

Any default accounts delivered or configured within a new system should either be deleted or renamed. Additionally, default passwords must be changed upon installation. Another good

cyber hygiene practice is periodically reviewing each system to make sure the active accounts are all authorized. Cyber-attackers sometimes create an account for themselves on systems for future backdoor access. Extraneous unauthorized accounts are a potential sign of compromise.

It is particularly important to remove account access for personnel who leave the company under duress. There are many examples of disgruntled former employees destroying data or systems after their employment is terminated.

Managing accounts is an essential cyber hygiene practice.

4.4.7 Authentication

In addition to account management, authenticating the person who accesses the account is essential. Cyber hygiene practices dictate having policies and technical enforcement mechanisms for strong passwords. Additionally, users should be prohibited by policy from sharing account passwords or writing them down in an unsecure location.

Multi-factor authentication adds an extra layer of security to account access. The three primary factors are “something you know, something you are, or something you have.” For example, a password is “something you know.” “Something you are” includes biometrics, such as fingerprints and facial recognition. “Something you have” can be a physical token with a continuously updated passkey. Using multi-factors for authentication makes theft and reuse of account credentials more challenging for adversarial cyber actors.

If remote access is allowed on a system, then multi-factor authentication is strongly recommended as a part of good cyber hygiene. In addition, screen locking mechanisms after long periods of inactivity and session timeouts should be configured. That can prevent an unauthorized user from accessing the system through an unattended session.

4.4.8 Least Privilege

System users should not have administrative or privileged access unless it is needed to carry out their job duties. Additionally, users that legitimately require privileged access, should have two accounts, one with the elevated privileges and one with standard user permissions. The account with administrative access should be used only when performing tasks that require it.

Users should not have access to data or information on the system that is not required to perform their job duties. There are many business reasons for that practice. For example, all employees should not have access to payroll data or financial transaction systems of a company.

Account privileges and data access should be reviewed on a periodic basis. As organizational roles and responsibilities evolve, sometimes users retain access to systems or information that is no longer driven by business need. Periodic review can identify and correct that issue.

From a cyber hygiene perspective, limiting access to data can prevent a cyber-attacker from gaining access to widespread information from a single compromised account.

4.4.9 Separation of Duties

Separation of duties is an effective practice for mitigating insider threats. For example, separating privileged access from monitoring and auditing duties could prevent a person from using privileged access to cover up their own malicious activities. Separation of duties is also an important fraud control in banking. In that context, a financial institution might require separate people when initiating and authorizing financial transactions. The same principle can be used to manage data transfers to and from sensitive systems.

Separation of duties means that no single person has access to all the “keys to the kingdom.” It also reduces the risk of bad behavior without collusion. That practice also helps minimize potential impacts of the compromise of a single account.

Separation of duties is an essential cyber hygiene principle.

4.4.10 Protect Data

Data encryption is a best practice for maintaining the confidentiality of sensitive data. It ensures that only authorized parties with access to the decryption key can access and use the information. Even if a malicious cyber actor gains access to a system, and exfiltrates or accesses encrypted files, access to the contents is denied.

Data should be encrypted when “at rest” or stored in a system. It should also be encrypted when transmitted over external networks.

Good data encryption practices can prevent unauthorized data exfiltration and espionage operations. It can also protect the company from extortion attempts to not reveal stolen data. For aviation systems, encryption can protect intellectual property and trade secrets.

In addition to encryption, it is generally a good practice to refrain from retention of sensitive data that is not needed. For example, recent data breaches have resulted in leaks of sensitive personal data that wasn’t truly necessary for the company to have on their systems in the first place.

Data protection is an essential cyber hygiene practice.

4.4.11 Automated Backup System

Backups of data, software, and system configuration are essential for recovery following a loss of cyber-enabled assets. When the stakes are very high, sometimes it even makes sense to also maintain backups of hardware.

It should be noted that backups are essential as a mitigation of potential loss from cyber events. The need for system restoration is not limited to malicious cyber-attack. Loss can also occur due to natural disaster, accident, or human error. It is always a good practice to have backups.

For that reason, it makes sense to keep backup media and hardware in a physically separate location away from the primary system. In the event of natural disaster, backups can be destroyed when co-located with the assets that were lost. Additionally, for systems that cannot

tolerate short outages, some organizations implement “hot” backup mirrored systems that can take over processing in the event of localized outages.

There are two additional important practices regarding backups. First, the restoration process should be exercised from time to time. The enterprise IT industry is full of horror stories about companies that thought they were diligently creating backups only to discover that all the media was corrupted when it was actually needed.

Additionally, it is critically important to protect the backup media. Encryption is a best practice for all backups. Many recent data breaches involved data that was compromised through backups that were not sufficiently protected.

Automated backup systems can help speed recovery after loss from a cyber event, when done well. It is an essential cyber hygiene practice.

4.4.12 Monitor and Analyze Audit Logs

Modern IT enterprise operating systems and network management software comes with options for collection of audit information for monitoring systems. It is critically important that these systems be configured to collect data and store it in case it is needed later for forensic analysis of actual or suspected malicious cyber events.

It is also a best practice to have some method in place for periodic review of audit data that is collected. There are many instances of highly capable threat agents gaining access to a system and going undetected for extended periods of time. The longer a cyber incursion goes undetected, the greater the consequences.

Intrusion Detection Systems (IDS) monitor assets real time and notify administrators of anomalous events or traffic. Monitoring a system with an IDS will flag potential cyber-attacks as they occur which is considerably faster than subsequent examination of audit logs.

As previously mentioned, organizations that are not actively looking for security events and failures within their system are not likely to detect any. To protect the critical infrastructure and the aviation systems we develop, collecting audit data and periodic review is an essential cyber hygiene practice.

4.4.13 Use Firewalls and Network Segmentation

It is a very bad practice to connect devices directly to the internet. Using a security device known as a network firewall to protect information flows in and out of internal networks and systems is an absolute minimal requirement.

IT firewalls inspect packets of data traversing external and internal networks. Firewalls determine if each packet is allowed to pass from one part of the system to another. Those devices can protect the internal network from unauthorized access and many types of external network-based cyber-attacks. On the internal side, firewalls can detect and prevent violations of the organizations security policy.

Depending on the size and complexity of cyber-enabled systems in an organization, dividing the network into isolated segments is a best practice that significantly mitigates risk. Network segmentation can be implemented using virtual local area networks configured at the switch, or firewalls. In the event of malicious cyber incursions, network segmentation can make it more challenging for a threat actor to pivot to all parts of the system. In many cases, segmentation can isolate the consequences of compromise.

Network segmentation and firewalls are a valuable security practice that protects cyber-enabled assets and network integrity. Effective deployment of these devices can also reduce the severity of security incidents.

4.4.14 Include ICS and SCADA Systems

While much of this section uses language specific to enterprise IT computing assets, it is also important to protect ICS and SCADA devices using these same cyber hygiene principals. In fact, these systems can be more of a challenge to protect due to less mature COTS security mechanisms targeted specifically to those types of devices.

Many legacy ICS and SCADA systems were developed prior to industry recognition that they are also targets of hostile cyber threats. In many cases those devices are more vulnerable than enterprise IT systems, and the resulting consequences to human life and the critical infrastructure can be devastating.

Isolating ICS and SCADA systems from external networks is recommended, but not always possible. Consequently, network segmentation that isolates these systems to the greatest extent possible is a best practice when a physical air gap is not practical.

An organization that operates ICS and SCADA systems must include these device types in infrastructure protection plans. Cyber hygiene best practices are absolutely essential for those devices.

4.4.15 Minimize Attack Surface

One of the best ways to protect cyber-enabled systems is to limit the places where a hostile threat actor has access to the greatest extent possible. This network design principle is typically referred to as “minimizing the attack surface.”

Organizations should not have internal or external network connections that do not serve a well-established and documented business purpose. Similarly, all computing and hardware assets should also be performing some beneficial function to the organization.

A proxy server between an organization’s internal network and the internet can significantly reduce the external cyber-attack surface of a system. By acting as an intermediary between users and the internet, a proxy server hides internal network details. It can also be configured to block access to external websites or servers, which may harbor malicious software or carry potential cybersecurity risks.

Removing unnecessary software, hardware, and network connections cuts down on the access vectors available to a malicious cyber adversary. Minimization of attack surface can greatly reduce the risk profile of a system. It is an essential cyber hygiene practice.

4.4.16 BYOD/Mobile Device Strategy

Many organizations have embraced a highly connected culture that allows mobile computing devices to connect to enterprise network systems. In fact, the ability to work remotely is a characteristic of an agile and responsive organization, which can be a benefit. However, mobile devices bring additional risk to an enterprise network. Organizations need to identify this risk and implement controls to minimize potential impacts.

Additionally, some organizations allow users to “Bring Your Own Device” (BYOD) to connect to the enterprise network. The security configuration and visibility into those devices may be weaker than for organizationally owned and managed endpoints. Personally owned devices bring even greater risk to the system.

Organizations that use mobile devices to access their computing resources must include policies and provisions for data protection such as encryption and multi-factor authentication. Additionally, mobile devices are more likely to be misplaced or lost. Mechanisms for remotely wiping data is highly recommended. Additionally, a method to disconnect mobile sessions after a period of inactivity is another best practice for those devices.

The benefits and convenience of mobile devices can be a productivity enhancer for organizations. However, the potential risks associated with these devices must also be understood, monitored, and managed.

Effective management of mobile devices is an important practice of cyber hygiene.

4.4.17 Remote Access

Remote work has become a common practice for many organizations. Consequently, many companies allow access of enterprise systems from non-business locations. It is essential that remote access be tightly monitored and controlled to prevent unauthorized access.

Virtual Private Networks (VPNs) create a secure encrypted tunnel between remote users and the organizations computing assets. VPNs can be configured to identify users using multi-factor authentication methods and also to collect audit data.

Remote access has become an essential part of doing business for many organizations. Due to the ubiquitous nature, securing it is essential. Hostile adversaries will attempt to exploit remote access mechanisms as a vector into the organization’s systems.

4.4.18 Security Awareness Training

Organizational computing systems have a heavy reliance on its human users and administrators to maintain the security posture of the system. Consequently, cultivating a security conscious culture fosters the human behavior required to support secure operations.

Recurring security awareness training is an established mechanism for ensuring that users understand the security threats and risks facing organizational systems. Additionally, training can help personnel understand policies, procedures, and expectations.

At a minimum, users should be educated on the importance and best practices of selecting strong passwords. They should also be aware of the policies prohibiting account credential sharing. Training on the risks of phishing and opening suspicious emails is essential. Additionally, users should know the procedures and points of contact to report suspicious email traffic.

Security awareness training should also address the insider threat. Organizational personnel are often in the best position to identify malicious insiders.

Cybersecurity relies on our weakest link which is frequently the human operators. Educating the users of their role and responsibilities in protecting computing assets is a critical cyber hygiene best practice.

4.4.19 Disposal Procedures

Every organization should have clear policies and procedures for secure disposal of computing assets. Before systems are sold or scrapped, any sensitive data stored on the systems must be destroyed. That includes traditional computing assets as well as other office systems such as photocopying machines.

Simply wiping data could prevent a cyber adversary from gaining information that enables a successful direct attack against an organization's systems. In some cases, additional measures should be taken to ensure the data is not recoverable using current techniques of nation state level threat actors. Sometimes it is best to destroy storage media rather than scrapping it.

4.5 DFARS

All organizations in the aviation system supply chain should have an understanding of the cybersecurity clauses and requirements in the Defense Federal Acquisition Regulation Supplement (DFARS). It is important to emphasize that the cyber DFARS only apply to the DoD acquisition contracting process. However, there is benefit from the practices and standards imposed by the cyber DFARS for all aviation system suppliers and manufacturers.

Subpart 204.73 of the DFARS describes safeguarding covered defense information and cyber incident reporting. It requires that contractors provide "adequate security" on all covered contractor information systems. This section also imposes all nonfederal organizations that have Controlled Unclassified Information (CUI) on their systems comply with NIST SP 800-171. (26) It also requires that contractors and subcontractors rapidly report cyber incidents.

Subpart 204.75 of the DFARS describes the policies and procedures for including the Cybersecurity Maturity Model Certification (CMMC) level requirements in DoD contracts. (27) The CMMC is a framework that measures a contractor's cybersecurity maturity to include the implementation of cybersecurity practices and institutionalization of processes.

The CMMC recommends five basic steps to help companies stay cyber secure. (27)

1. **Educate people on cyber threats.** Most cyber incidents start because of user error. Educate people about the importance of setting strong passwords, recognizing malicious links, and promptly installing the latest security patches.
2. **Implement access controls.** Limit access to information systems to authorized users and the specific privileges for the actions that they need to perform to fulfill their job duties.
3. **Authenticate users.** Use multi-factor authentication tools to verify the identities of users, processes, and devices.
4. **Monitor your physical space.** Escort visitors and monitor visitor activity, maintain audit logs, and manage physical devices like USB keys.
5. **Update security protections.** Make sure to promptly download the latest security patches when new releases are available. Always double check to make sure patches and software updates are coming from a trusted source.

Those five basic steps reflect the DoD perspective on the cyber hygiene best practices previously outlined in Section 4.4.

4.6 Air-Gapped Networks

An air-gapped network is an enclave that is isolated from other networks. It is a common method used to create a highly secure environment to protect sensitive data and systems. Sometimes air-gapped networks also exist when there is no need for interconnectivity between computing systems.

An air-gapped network is often used as a method for reducing the attack surface of a system. However, it would be a mistake to believe that unauthorized access or hostile cyber-attacks cannot traverse the air-gap. Frequently, air-gapped networks support data transfers using media which is colloquially referred to as sneaker-net.

“Countdown to Zero Day: Stuxnet and the Launch of the World’s First Digital Weapon,” recommended in Section 2.9, details how a sophisticated cyber-attack was carried out across an air-gapped network.

It should be noted that the same cyber hygiene practices identified throughout Section 4.4 also apply to air-gapped networks. In other words, the absence of persistent network connections is not a valid reason for failing to follow cyber hygiene best practices.

Air-gapped networks are a great technique for reducing attack surface and preventing access by malicious threat actors. However, even air-gapped networks are not completely impervious to attacks. The owners and operators of those systems must be cognizant of those risks.

4.7 Incident Response

No matter how well an organization’s computing resources are protected, highly skilled and well-funded nation state adversaries will occasionally have success penetrating computing systems. Consequently, it is important that all organizations have a process for incident response in place prior to experiencing a cyber event.

All organizations should have systems in place to detect adversarial cyber activity within their computing assets. Section 4.4.12 described the configuration of auditing and forensic data as well as periodic reviews of that information.

When an incident is suspected or confirmed, there must be a process and strategy in place for containing the impacts. That may involve isolating subnetworks or systems to prevent further spread of the attack. It can also involve disabling accounts or functions within the system.

Once the cyber incursion is contained, the next step is to remove the adversary from the network. This can be challenging as most sophisticated threat agents will take steps to install backdoors or other footholds into the system to maintain persistent access. There are many documented instances where an organization thought that an adversary had been expelled from the system only to have them quickly return.

Once the adversary is eliminated from the system, the recovery process can begin. Section 4.4.11 discussed the importance of automated backup systems. This highlights another reason why forensic analysis is important because understanding when and how the adversary initially gained a foothold is critical to ensuring that their access isn't restored along with the backup data.

All organizations in the aviation supply chain should have an incident response plan. It is also important to exercise that plan on a periodic basis to ensure that it works and that people know how to implement it under pressure.

4.8 Summary

Designing secure aviation platforms relies on reliable and secure development systems. Consequently, the infrastructure of the development organizations within the supply chain is subject to adversarial cyber-attack from advanced nation state threat actors. The network connected enterprise IT systems operated by every participant in the aviation ecosystem represents exploitable attack surface to the adversary.

Additionally, any industrial control systems (ICS) or cyber-physical systems (CPS) in our business and supplier networks potentially represents access vectors for highly skilled cyber adversaries.

There is no one-size-fits-all solution for identifying, managing, and mitigating cyber threats and risks to business and development systems. Organizations will have unique characteristics and practices which can impact the resilience and operation of their private networks.

However, there are steps that can be taken to protect development systems and, by extension, the aviation industrial base and supply chain. It is incumbent on all stakeholders to do exactly that.

Part 2 – Secure Architecture and Design for Air Systems



Two MC-130Js for U.S. Air Force Special Operations Command nearing completion on the C-130J production line at the Lockheed Martin facility in Marietta, Ga.

Chapter 5

The Foundations of Secure Architecture and Design of Air Systems

Modern aviation platforms must be architected, designed, and operated to satisfy high standards of safety, reliability, and performance. The culture of safety was instilled in the aviation industry during the earliest days of aircraft design and development. However, the primary focus of aircraft safety initiatives has been prevention of accidents resulting from systemic failures and human error. The aviation platforms intended for military purposes were also built for survivability against kinetic weapons. Consequently, the industrial base that creates aircraft and parts for both military and commercial use have long been subjected to more rigorous scrutiny for safety, reliability, and performance when compared to general enterprise IT systems.

Unfortunately, it's not enough. As aviation platforms increasingly depend on cyber enabled operation, the very technology that creates performance advantages also opens new avenues of attack by highly skilled and determined cyber adversaries. All aircraft operates in cyber-contested airspace that requires extension of the existing high standards imposed by safety on those platforms. Rigorous evaluation must now consider actions initiated by malicious cyber adversaries as well as incidental cyber events.

All aviation systems must be architected, engineered, implemented, operated, and sustained for high levels of cyber resilience and survivability. Part 2 of this FSAD guidebook extends the foundational concepts described in Part 1 to include additional considerations for the design of cyber-enabled aviation parts and subsystems. These are the baseline concepts that underpin acquisition, development, deployment, and operation of aircraft that are cyber resilient and survivable against both malicious and incidental cyber events.

5.1 Cyber Survivability, Cyber Resiliency, and Cybersecurity

The United States DoD now requires many aviation acquisition programs to include cyber survivability as part of the mandatory System Survivability Key Performance Parameter (SS/KPP). (28) Consequently, cyber survivability is increasingly reflected in top-level program requirements and performance objectives for military platforms.

In that context, cyber survivability means that aviation systems are designed to prevent, mitigate, and recover from sophisticated cyber-attacks and cyber events. How each of the cyber survivability objectives are achieved depends on the unique implementation details for each aircraft platform. The capabilities, mission objectives, and underlying architecture all play a role that guides the selection of cyber survivability objectives.

While non-military aviation systems have generally not been levied cyber survivability requirements, that is likely to change in the future. Commercial aircraft systems rely on the same

types of interconnected technologies and digital infrastructure as military platforms. Additionally, the cyber-attack surface and direct physical access is generally more exposed to hostile cyber actors.

The concepts of cyber survivability, cyber resiliency, and cybersecurity are used repeatedly throughout this FSAD guidebook. Each of these related terms carries a subtly distinct definition.

Cyber survivability is the ability of the aircraft to prevent, mitigate, recover from and adapt to adverse cyber-events that negatively impact mission performance. Cyber survivability also includes preservation of human life and the platform. For military aircraft, cyber survivability is sometimes said to be the ability to live to fight another day. In the commercial context, this might be stated as the ability to fly another day.

Cyber resiliency is the ability of the aircraft to anticipate, withstand, recover from, and adapt to adverse conditions, stresses, attacks, or compromises throughout the full lifecycle of the platform. Cyber resiliency starts during acquisition and extends throughout the operational and sustainment phases by continuously evolving and adapting to dynamic threats within the environment.

Though these terms are sometimes seemingly used interchangeably there is a subtle difference. Cyber survivable systems are the result of designing for cyber resiliency.

The term **cybersecurity** has fallen somewhat out of fashion in favor of cyber resiliency and survivability. In the past, cybersecurity was the dominant term during the nascent phases of security evaluation for military platforms. During that era, preservation of confidentiality was emphasized above other security aspects such as integrity, and availability. Consequently, the term cybersecurity is sometimes misconstrued to be limited only to the protection of data. That narrow definition is never the intended usage within this FSAD guidebook.

Cybersecurity remains fundamental to the expanded emphasis on the specific aspects of resiliency and survivability. In essence, the three terms are complementary to each other, each providing a vital aspect of a holistic strategy to address the evolving and persistent nature of cyber threats. The terminology reflects a maturation in understanding that a multi-faceted, adaptive approach is necessary to safeguard against an increasingly sophisticated and diverse array of cyber threats.

5.2 The Confidentiality-Integrity-Availability (C-I-A) Triad

The Confidentiality-Integrity-Availability (C-I-A) triad is a useful mental model for conceptualizing information security. While the precise origins are lost to history, the cybersecurity domain is filled with references to the “classic” C-I-A triad. Regardless of the origins, it is a guiding model for understanding conceptual threats against all systems, including aviation platforms.




Figure 8. The C-I-A Triad

Figure 8 illustrates the three crucial components of the C-I-A triad. The information security roots of the model are reflected in the original definitions of these core concepts:

Confidentiality: Sensitive information is not disclosed to unauthorized people or systems.

Integrity: Assurance that information has not been improperly altered or destroyed.

Availability: Timely and reliable access to information by authorized people or systems.

These definitions are very much rooted in information security and assurance. DoD platforms with stringent requirements for protection of classified and sensitive data, frequently led to an overemphasis on protecting confidentiality while neglecting integrity and availability. For a time, it was not uncommon for classified military platforms to shut down processing or communications if breaches of confidentiality were suspected. In other words, availability was routinely neglected in the interest of maintaining confidentiality. That is not consistent with the core objectives of cyber resiliency.

Based on past misuse, some cybersecurity engineers are reluctant to leverage or reference the C-I-A triad. However, it remains a valuable foundational model on which cyber resiliency and cyber survivability is built. Balanced consideration of the three core components of confidentiality, integrity, and availability are all equally important for cyber resilient systems.

As more advanced topics in cyber resiliency and survivability are introduced, it is useful to remember and apply the C-I-A triad when evaluating solutions. It is a necessary foundation for systems to achieve cyber resiliency.

5.3 Intrinsic Cyber Resilience

The word "intrinsic" describes qualities or characteristics that are inherent, essential, or integral to the nature of something. Intrinsic cyber resilience is the foundational ability of a system to withstand and recover from cyber threats or attacks. Systems that are intrinsically secure build cyber resilience directly into the architecture and design rather than relying on external security controls or reactive measures.

The acquisition agencies and customers who purchase aircraft systems fundamentally want products that are intrinsically cyber resilient. While the same can also be said of the buyers of traditional IT networked systems, that quality is particularly important for aircraft due to the unique performance requirements and architectural constraints.

Aviation platforms require embedded specialty purpose cyber physical operating systems. Real time deterministic performance is an imperative for flight control, navigation, communication, and capabilities. Aviation platforms require deterministic performance including guaranteed response times, limited latency, and low rates of jitter.

The traditional body of knowledge and best practices for cybersecurity were developed primarily within the context of enterprise IT systems that lack the rigid real time operating constraints required for aircraft. In the business IT environment, it is acceptable for operating systems and communication protocols to use flexible and configurable security functionality that is not well suited for deterministic systems. Additionally, using add-on products and services such as Intrusion Detection Systems (IDS), Anti-Virus (AV), and security service applications are not yet proven to be practical or effective for aircraft platforms.

While some companies and researchers are investing in technology and techniques to enable add-on security solutions for legacy aviation platforms, the best approach starts from an understanding that the most effective security for aircraft will emerge from building the attributes of intrinsic cyber resilience directly into the system from the earliest stages of development.

Along those same lines, it is also important to understand that characteristics of intrinsic cyber resilience cannot be bestowed by external actions. For example, it is not possible to prove that an aircraft platform is cyber resilient by performing cyber test and evaluation. While testing is essential, it can only prove the existence of vulnerabilities. Testing can never provide definitive proof of the absence of issues or problems.

5.4 Build Systems Better / Build Better Systems

The cyber industry is filled with motivational slogans intended to inspire the workforce to develop cyber resilient systems. However, directives and exhortations to “Build Security In” are mere slogans and not effective without the knowledge of how to actually do it.

The development best practices that support and enhance the cyber resiliency posture of a platform fall under two basic categories. Each represents a vital perspective of a system’s cyber resiliency.

- **“Build Systems Better”** refers to the processes and procedures that support the systematic identification and implementation of intrinsic cyber resiliency attributes as the system is developed. This category encapsulates “how” to develop cyber resilient systems.
- **“Build Better Systems”** describes the intrinsic properties of cyber resilience that result from the activities described in the previous category. It includes the architectural

features and qualitative aspects of the system that is developed. This category captures “what” is ultimately reflected in the implemented system.

The subsequent chapters of Part 2 of this FSAD Guidebook describes specific techniques which are identified as best practices to “Build Systems Better.” That includes methods for eliciting resiliency objectives and requirements for the system, and methods of analysis and evaluation that can be used to quantify the effectiveness of those efforts.

This part of the FSAD Guidebook also introduces concepts of “Build Better Systems” through reusable cyber resiliency archetypes. Those design patterns are potential solutions to meeting cyber resiliency objectives and requirements.

Program management plans should include explicit deliverables and verification activities for ensuring that cyber resiliency is considered at every phase of development. “Build Systems Better” and “Build Better Systems” require consistent application of engineering best practices.

5.5 Systems Engineering

Systems Engineering (SE) is an interdisciplinary field of engineering and management that focuses on designing and managing complex systems over the entire life cycle. It involves applying engineering principles and methods to analyze, design, develop, and manage systems, which may include hardware, software, people, processes, information, and facilities.

This FSAD guidebook assumes a baseline level of familiarity with the fundamental concepts of systems engineering. In fact, the chapters in Part 2 are structured in alignment with the SE “V” model. The foundational concepts of systems engineering underpin engineering for cyber resiliency.

Figure 9 is a visual representation of the SE “V” model. The left side represents the high-level tasks that are performed during system development. That includes the System Requirements Definition, Architecture, Design, and Implementation. The right side of the “V” contains the activities that are performed to verify and validate the system that was developed.




Figure 9. Systems Engineering “V” Model

Engineering for cyber resilience doesn’t revolutionize traditional development engineering practices. In fact, performing the foundational systems engineering for all aspects of the aircraft is of paramount importance. Instead, engineering for cyber resilience augments the systems engineering required for system development with additional focus areas and analysis to ensure that the aircraft achieves its cyber resiliency objectives in the face of adversity.

5.6 Security Systems Engineering / Cyber Resiliency Systems Engineering

Security Systems Engineering (SSE) / Cyber Resiliency Systems Engineering (CRSE) is an emerging specialty engineering field that augments the traditional systems engineering tasks and activities with cyber security domain specific extensions.

This guidebook uses the following definition for SSE/CRSE. It is adapted from the U.S. government’s Committee on National Security Systems:

SSE/CRSE is an interdisciplinary approach and means to enable the realization of secure systems. It focuses on defining customer needs, security protection requirements, and required functionality early in the systems development lifecycle, documenting requirements, and then proceeding with design, synthesis, and system validation while considering the complete problem. (29)

Cyber resiliency is a system-level qualitative property that requires holistic system-level analysis. Building aircraft systems that behave in a secure manner requires first developing a clear and unambiguous concept of what constitutes secure behavior at the system level. Only then can subsystem and component level engineers select and create solutions that deliver the desired secure behavior. The identification and development of the suitable set of components takes place across the entirety of the development lifecycle.

When the need for SSE/CRSE first emerged as a specialized discipline of SE, cybersecurity was initially thought to be the exclusive domain of a handful of specialists. However, as the collective understanding of cyber threats and risks has matured, it has become apparent that all stakeholders participating in system development require a great deal of competence and knowledge in this domain. As emphasized in Chapter 3, cybersecurity is everybody's job.

NIST describes SSE as the complementary engineering capabilities that delivers the concept of trustworthiness and secure systems. (30) Trustworthiness starts with demonstrably meeting a set of requirements that are complete, consistent, and correct. A trustworthy system includes security requirements in addition to those for the functions and capabilities of the system. The discipline of SSE/CRSE focuses on the creation of trustworthy secure systems that limit and prevent the effects of cyber adversity from both malicious and non-malicious sources.

Designing and building secure aviation platforms and weapon systems requires a coordinated effort from the entire organization. That starts with senior leadership teams and extends down to program managers, designers, engineers, builders, supply chain experts, and security professionals.

While SE focuses on ensuring that systems effectively and efficiently realize functional objectives, SSE/CRSE focuses on system performance in disruptive and cyber contested environments.

5.7 Lifecycles, Methodologies, and Model Based Systems Engineering

Multiple lifecycle models can be used to perform the systems engineering for complex platforms. There is no singular "right" lifecycle model for the development, deployment, and maintenance of cyber resilient aircraft. While this FSAD Guidebook is expressed against the "V" lifecycle model, that is merely a convenient framework for structuring the best practices that are executed regardless of the lifecycle model used.

Consequently, this guidebook can be used in conjunction with multiple methods. That includes the waterfall model, the agile model, and the spiral model. Each approach requires the same systems engineering focal areas. What differs between each one is the frequency and duration of each iteration. The "best" lifecycle model for any given effort may vary from system to system and from team to team. Systems engineers should have the flexibility to select the most suitable framework based on their unique system goals, team culture, complexity, and stakeholder needs.

A systems engineering lifecycle model is distinctly different from methodologies for creating architecture and design of systems. While a lifecycle model includes the entire lifespan of a system, a development methodology focuses specifically on the methods, techniques, and practices used during development. A methodology dictates how the system is built, including the processes, tools, and best practices employed by the development team.

Model Based Systems Engineering (MBSE) is a methodology that leverages digital models as the primary means of creating, storing, and exchanging information about a system. MBSE was established as an INCOSE initiative in 2006, though the various aspects of the techniques were

in use throughout the 1990s. Recently, MBSE has been gaining traction as a mechanism for engineering aviation systems for cyber resiliency within the DoD.

MBSE is an excellent methodology for the development of complex cyber physical systems required for aircraft. However, simply using an MBSE approach or tool will not automatically result in more cyber resilient systems. If MBSE is used and the best practices identified in this guidebook are ignored, then the system isn't likely to achieve its resiliency objectives. Similarly, a project that uses the techniques outlined in this FSAD Guidebook using any alternate methodology can also deliver an intrinsically secure system.

Organizations that embrace MBSE principles and practices can leverage digital engineering tools and techniques to improve the efficiency, quality, and maintainability of their system development processes. Regardless of the methodology used, attention must be paid to the SSE /CRSE principles outlined in this FSAD Guidebook to create intrinsically cyber resilient systems.

5.8 Airworthiness

Airworthiness is the property of an aircraft to safely attain, sustain, and terminate flight in accordance with the approved usage limits. Commercial and military systems use separate policies and standards for airworthiness certification. In general, both require a repeatable process for verification that an aircraft can be safely maintained and operated within its prescribed flight envelope.

In the United States, commercial aircraft is subject to airworthiness certification by the Federal Aviation Administration (FAA). The need for cyber resiliency in aircraft has recently been recognized as an emergent property necessary for aircraft safety. The FAA has chartered an Aviation Cyber Initiative (ACI) to “reduce cybersecurity risks and improve cyber resilience to support safe, secure, and efficient operations of the Nation's Aviation Ecosystem.” (31) However, most of the efforts to date have focused on enterprise IT ground systems rather than the actual aircraft.

The primary resource for DoD airworthiness certification is MIL-HDBK-516C, "Airworthiness Certification Criteria." (32) It establishes the certification standards and methods of compliance to be used in the determination of airworthiness of all manned and unmanned, fixed, and rotary wing aircraft systems.

The Airworthiness Certification Criteria contains limited language on security functionality. However, much of the context is limited to securing command and control communications. It also contains directives indicating that security techniques must be implemented safely. The handbook also specifies that security requirements have been applied to the processing architecture to protect safety critical functions.

To better address DoD policy that cybersecurity must be considered throughout the lifecycle of the air system, the United States Navy Airworthiness and Cybersafe Office (ACO) created and released a Cyber Security Supplement (33) to the Airworthiness Certification Criteria. (32) Though that resource is no longer publicly available, it provided a tailored set of air system

cybersecurity considerations based on the airworthiness certification criteria, standards and methods of compliance contained in the parent handbook. It also described associated evidence and data artifacts intended to support air system cybersecurity certifications with tasks and activities that are applied throughout the development lifecycle.

The practices specified in this FSAD Guidebook are consistent with the methods of how to structure and document method of compliance used in the Airworthiness Certification Criteria Cyber Security Supplement. For example, the general recommendation to perform Cybersecurity Risk Assessment (CRA) throughout the system lifecycle is described in detail in Chapter 9. It also describes a Defense-In-Depth functional architecture which is described in section 7.4.3.

At first glance, the (now deprecated) Cyber Security Supplement appeared to be onerous based solely on the length alone. However, there was a lot of overlap with the parent Airworthiness Certification Criteria. Essentially, the supplement augmented the baseline aircraft design criterion and standard to determine if a component has cybersecurity impacts. If it does, then a method of compliance was specified.

For example, the criterion outlined in Section 6.1.1.5 of the Cyber Security Supplement which describes “Modeling, simulation, analysis tools and databases” is marked as applicable for cybersecurity. Consequently, that section specifies a cybersecurity specific method of compliance as illustrated by the following excerpt:

“Cybersecurity: Examination: Evaluate models, simulations and tools against the approved list of Navy software allowed for use. Evaluate any commercial software products used to ensure they are the current software version and that the latest patch levels are installed.”

That recommendation is consistent with practices described in Chapter 4 of Part 1 of this FSAD Guidebook.

Both commercial and military aviation platforms are subject to airworthiness certifications. Cyber resiliency is an emerging consideration in both contexts. This FSAD Guidebook is not intended as a replacement for any applicable standard. It is the responsibility of the manufacturer of each part to ensure compliance as is appropriate.

However, this FSAD Guidebook is an excellent resource for understanding the principles and practices that underpin the cybersecurity recommendations currently in airworthiness standards.

5.9 System Assurance

The National Institutes of Standards and Technology (NIST) defines assurance as the “measure of confidence that the system functionality is implemented correctly, operating as intended, and producing the desired outcome with respect to meeting the security and privacy requirements for the system – thus possessing the capability to accurately mediate and enforce established security and privacy policies.” (34)

System assurance is critically important for aviation systems due to the inherently high consequences associated with risks. Aviation systems, including aircraft, air traffic control systems, and associated infrastructure, must operate safely and reliably to ensure the well-being of passengers, crew, and the general public.

System assurance is establishing confidence that the aircraft meets its intended functionality and performance requirements. It involves identifying and mitigating risks associated with development, deployment, and operation to ensure reliability, security, safety, and quality.

Security Assurance is a specialty domain within System Assurance. It ensures that security threats and risks are identified and mitigated as a part of system development. Part 3 of this FSAD guidebook delves deeper into methods of achieving system assurance for software and hardware.

5.10 Additional Reading and Resources

This FSAD Guidebook is written specifically for the manufacturers and designers of aviation systems. It is currently believed to be the only resource that is written directly to that niche audience.

However, there are additional excellent resources that can further illuminate the principles described in this part of the FSAD Guidebook. The following additional reading is recommended for people who wish to go into greater depth on these topics.

- **Systems Engineering Handbook: A Guide for System Life Cycle Processes and Activities**, International Council on Systems Engineering (INCOSE), INCOSE-TP-2003-002-05-2023, Fifth Edition. (35)

This book describes the key process activities performed by systems engineers and other engineering professionals throughout the system lifecycle. It is the authoritative source of a wide range of fundamental system concepts that broaden the thinking of the systems engineering practitioner, such as system thinking, system science, life cycle management, specialty engineering, and system of systems. It covers all the major lifecycle models including waterfall, V model, agile, and iterative methods. This book is ideal for anyone who has an interest in systems engineering practices.

- **Engineering Trustworthy Systems: Get Cybersecurity Design Right the First Time**, O. Sami Saydjari, McGraw-Hill Education, 2018. (36)

This book offers comprehensive examples, objectives, and best practices for designing and deploying highly secure systems. It outlines how to think strategically, identify the highest priority risks, and apply advanced countermeasures that address the entire attack space. It is a practical guide for leveraging the timeless engineering principles resulting in trustworthy systems.

- **Engineering Trustworthy Secure Systems**, Ron Ross et al., NIST Special Publication, NIST SP 800-160v1r1, November 2022. (30)

A detailed framework for incorporating security into the systems engineering lifecycle, emphasizing the importance of integrating security measures from the initial design phase through deployment and maintenance. It outlines best practices, principles, and methodologies for ensuring that systems can withstand, adapt to, and recover from cyber threats. The publication aims to aid organizations in developing robust security strategies, enhancing resilience, and ensuring compliance with security standards and regulations.

Chapter 6

Program Inception and Requirements

The critical nature of aviation systems requires close attention to cyber resiliency starting in the earliest stages of the development process. Cyber resiliency as an intrinsic attribute of the underlying system is best achieved by intentional direction and effort to “build security in” from the onset. The way to make that happen is through establishing clear requirements and objectives at program inception.

It is very difficult and extremely costly to mitigate risks in an aircraft that is suddenly found to lack cyber resiliency during the final stages of the development lifecycle. The deficit is even worse after it has already been fielded for operational use. Sometimes establishing an acceptable cybersecurity posture isn’t possible without significant architectural and design updates. Systems that end up in that predicament almost always suffer from an absence of cyber resiliency objectives integrated within the project scope, plans, and schedules. In the absence of cybersecurity requirements, architects and designers are unlikely to add security functions and features necessary to adequately secure the platform.

This chapter introduces the idea that project plans, schedules, and resource allocations must explicitly identify and include cybersecurity analysis at every step of the development lifecycle. Additionally, it describes how to elicit, identify, and document cyber resiliency requirements and objectives.

6.1 Cyber Resiliency at Program Inception

All stakeholders including program managers, project planners, and systems engineers must understand the significant investment required to develop effective cyber resiliency requirements and objectives as each project first gets underway. This goes beyond platitudes and high-level statement of intent. Detailed analysis that is unique to each aircraft based on the purpose and mission is essential. While specific aviation platforms are likely to face different types of adversity, the need to intentionally plan for how the aircraft will achieve an appropriate cybersecurity posture is critical from the onset.

Both the commercial and military sectors of the aviation industry are currently inundated with new cyber policy and directives. Many of those mandate some semblance of cyber resiliency in the acquisition and development process. Unfortunately, a General Accountability Office (GAO) report issued in 2021 found that three out of every five DoD programs did not include cybersecurity requirements, acceptance criteria, or any processes for verification. (37)

Some program managers and project engineers mistakenly choose to believe that if cyber resiliency objectives are not explicitly identified in program scope, it is a clear indication that the customer does not want it. That is never a valid assumption for any aircraft system.

In addition to cyber resiliency objectives and requirements to be implemented on the platform, program scope should also include analysis of the efficacy of the design and implementation against cyber threats. Scope for cyber risk assessment is necessary as a continuous process throughout the development lifecycle. Each of these activities must be explicitly reflected as tasks and milestones in project planning artifacts.

Aircraft program managers and developers must closely examine cyber resiliency objectives and requirements as projects are initiated. Technical assessment of the adequacy of any requirements that are levied should be performed. In some cases, the development organization may identify potential shortfalls or gaps in cyber resiliency objectives provided by the customer. If that transpires, compensating work to bring cyber resiliency requirements and objectives up to acceptable fidelity may be required in many instances.

6.2 Technical Program Reviews

Successful projects and programs periodically perform technical evaluation and audits throughout the system development lifecycle. These technical reviews explicitly must include status and evaluation of the cyber resiliency objectives for the program. Additionally, these evaluations must include continuous Cyber Risk Assessment (CRA) against the known and anticipated threats within the cyber contested airspace, maintenance, and sustainment environments.

Technical program review allows management, systems engineers, architects, designers, and developers to ensure alignment in the overall direction of the program. That collaboration supports knowledge-based milestone decisions that drive program execution and performance. Cyber resiliency is a key part of that evaluation.

A well-managed program has a detailed roadmap reflecting technical program reviews tied to the key milestones of the project. Each event must have detailed entry conditions with rigorous criteria that ensures that development has achieved an acceptable level of maturity for the review to be conducted. Similarly, exit criteria will define if the review has been successfully completed. The entry and exit criteria for each milestone must include specific items related to cybersecurity.

It should be noted that the need for technical program reviews is methodology independent. Using the waterfall, spiral, iterative, or model-based lifecycle approaches does not eliminate or materially alter the need for these events. Similarly, many DoD acquisition programs have their own defined set of technical program reviews. While that may alter the nomenclature, the fundamental need for technical review remains an imperative.

The technical risk management on the project or program must explicitly look for cyber resiliency risks, record any that are identified, and follow appropriate tracking and mitigation procedures. Cyber technical risks are a key part of the information assessed at technical reviews.

That includes potential cost and schedule implications of mitigating those risks as early as possible in the development lifecycle.

6.3 Systems Requirements Engineering

Well managed projects have a clear baseline of requirements at the onset of the development process. In addition to the features and functionality of the system, the cybersecurity objectives must also be clearly specified. Otherwise, the project risks implementing exactly what the customer wanted, absent the cyber resiliency that is essential to meet mission objectives.

Requirements that establish the functional objectives of a system are critical for the success of the project. It is the foundation of the plans for what ultimately needs to be achieved.

Cybersecurity requirements and performance objectives are similarly essential and must be explicitly accounted for by project management. Clear objectives are a necessity for creating a system that is intrinsically cyber resilient.

Addressing cybersecurity in requirements is necessary for scoping and planning the development effort. These requirements will determine not only the security posture of the system, but also the means and methods of performing periodic assessments against those objectives.

Disciplined and well-structured requirements engineering is a fundamental aspect of project risk management. System requirements are a key source of data for the identification of potential risks and challenges that may arise in project execution. Similarly, cyber resiliency objectives establish the foundation of Cyber Risk Assessment (CRA) that must be performed throughout the development lifecycle.

6.3.1 Requirements Documentation

There are many valid and effective methods for creating system requirements documentation. It is more important to have an authoritative and well-managed repository for that data, rather than using any one particular methodology. Additionally, the approach and tools used to document system requirements will frequently be tailored to the unique needs of each project.

Requirements management tools, computer aided systems engineering environments, and Model-Based Systems Engineering (MBSE) approaches can streamline the requirements documentation process and enhance collaboration. However, selection and use of any tool or methodology is never a substitute for rigorously performing the systems requirements engineering for a project. It is crucial to remember that technology is not a panacea for poor requirements management. Additionally, there is no such thing as a magic tool that will automatically inject cyber resiliency requirements or objectives into any system.

The traditional approach for documenting requirements involves gathering information through interviews, surveys, and workshops with stakeholders. The collected requirements are then documented in detail, often using techniques like use case diagrams and textual descriptions. When using traditional methods, it is important to ensure that the stakeholders possess a baseline level of knowledge and skill in working with cyber resiliency objectives. In many cases,

performance may be improved by augmenting the analysis team with personnel who have direct experience working with cybersecurity requirements in the aviation domain.

An agile approach to requirements management emphasizes collaboration and iterative development with requirements that evolve over time through constant communication between developers and users. Agile documentation typically consists of user stories, prioritized backlogs, and acceptance criteria.

The MBSE approach employs formal modeling languages such as the Systems Modeling Language (SysML) or the Unified Modeling Language (UML) to specify system requirements. MBSE models frequently are expressed visually which is believed to provide better understanding and communication about requirements among stakeholders.

Solid system requirements at project inception are essential for setting the direction, scope, and expectations of the project. It also drives and enables effective planning, communication, risk management, and quality assurance throughout the project lifecycle. While many different approaches to documenting requirements can produce successful outcomes, each has distinct advantages and challenges. The selection of requirements methodology and support will vary from project to project and be influenced by the size, complexity, previous team experience, and organizational culture.

Aircraft systems engineering has a long history of embracing rigorous system requirements management due to the need for compliance with safety regulations. Well managed requirements ensure that flight systems operate as intended and achieve high safety standards. Requirements for cyber resiliency and security form the foundation of that intrinsic qualitative attribute that is similar and closely related to safety. A platform that lacks an appropriate degree of cyber resiliency is also likely not meeting safety objectives in the face of malicious cyber actors and adversity.

6.3.2 Requirements Allocation, Decomposition, and Traceability

The systems engineering process for cyber resiliency requirements and objectives does not stop at program inception. The requirements management process is performed iteratively throughout the system development lifecycle. As the aviation system architecture is developed, the high-level requirements are allocated to subsystems and lower-level elements. As that occurs, the need for new cyber resiliency requirements may emerge as the initial architecture of the system matures.

Cyber resiliency requirements must be allocated and decomposed, just as is done for functional requirements. That systematic mapping of the high-level requirements down to the subsystem responsible for implementation defines what each element ultimately needs to accomplish. Disciplined requirements decomposition and allocation plays a critical role when designing intrinsically cyber resilient systems.

The objectives for preservation of mission critical functions must be directly expressed in the underlying system architecture. For example, if the requirements identify a mission critical

function, mechanisms such as redundancy or network segmentation may be reflected in the architecture to achieve that objective.

Requirements traceability is particularly crucial for cybersecurity requirements in aviation systems. Airworthiness certification relies on implementation of cyber resiliency mechanisms to protect sensitive systems and data from unauthorized access and abuse. When a traceability matrix documents how each requirement is allocated and decomposed, security testing and evaluation can be thoroughly performed. That increases confidence in that aspect of verification testing.

Cyber resiliency requirements are a vital part of a rigorous requirements management process. That means those requirements are subjected to the same allocation, decomposition, and traceability analysis that is performed for functional and capability requirements for the platform.

6.4 Elicitation of Cyber Resiliency Objectives and Requirements

As suggested by the GAO report referenced in Section 6.1, projects and programs frequently struggle to identify good cybersecurity and cyber resiliency requirements. That problem is exacerbated by the absence of exemplars of good cyber resiliency requirements for aircraft that can be leveraged as a turnkey starting point.

The Cyber Survivability Endorsement (CSE) Implementation Guide is one of the more mature resources supporting cyber resiliency system requirements elicitation. (38) That document is the result of a collaborative effort of DoD subject matter experts to define cyber survivability requirements. While the CSE Implementation Guide was created for DoD weapon systems acquisition projects, the content is applicable to both commercial and military aviation platforms.

Another potential source of cyber resiliency requirements comes from compliance with certification processes. For example, the Risk Management Framework (RMF) contains over 800 security controls which are sometimes directly levied as cybersecurity requirements for DoD aircraft systems. However, the majority of those controls address confidentiality, integrity, and availability rather than intrinsic cyber resiliency. (39) RMF and the Assessment and Authorization (A&A) process for aviation systems is discussed in much greater detail in Chapter 12.

At a minimum, system requirements must capture the spirit and intent of cybersecurity related compliance controls. A system that does not effectively implement compliance controls is most certainly not cyber resilient. However, a system that exclusively relies on compliance controls for cyber resiliency is also most certainly not achieving that objective. Cybersecurity compliance controls, such as the ones found in RMF, are necessary to support cyber resiliency but are insufficient in the absence of additional mechanisms.

The CSE Implementation Guide does not specifically identify cybersecurity requirements. Rather, it defines a process to help program and platform stakeholders understand the mission critical functions and cyber risks resulting from the loss of those capabilities. When those requirements are identified early in the development process, they inform engineering decisions that drive intrinsic cyber resiliency into the system architecture and implementation.

The considerations described in the CSE Implementation Guide are based on four pillars of system survivability (38), which are adapted here with language more directly applicable to aviation systems:

- **Prevent** – Design requirements to identify, protect and harden aircraft from adversarial cybersecurity threats.
- **Mitigate** - Design requirements to detect and respond to cyber-events and enabling cyber operational resiliency of the aircraft and mission objectives.
- **Recover** - Design requirements to recover to a good operational state after a cyber-event, preserving safety of flight and mission performance.
- **Adapt** – Enables sustainment functions for the aircraft to adapt to highly dynamic and evolutionary threats against it.

The CSE Implementation Guide identifies ten core attributes of cyber survivability and resiliency. Each of those maps into one of the pillars of system survivability. A summary of those attributes and mappings is provided in Table 3.

System Survivability Pillar	Cyber Survivable Attributes (CSAs)
Prevent	CSA 01 – Control Access
	CSA 02 – Reduce System’s Cyber Detectability
	CSA 03 – Secure Transmissions and Communications
	CSA 04 – Protect System Information from Exploitation
	CSA 05 – Partition and Ensure Critical Functions
	CSA 06 – Minimize and Harden Cyber Attack Surfaces
Mitigate	CSA 07 – Baseline & Monitor Systems and Detect Anomalies
	CSA 08 – Manage System Performance and Cyber Defense
Recover	CSA 09 – Recover System Capabilities
Adapt	CSA 10 – Actively Manage System Configuration to maintain an appropriate cyber survivability risk posture.

Table 3. Cyber Survivable Attributes (CSAs) (38)

The CSE Implementation Guide emphasizes that each of those attributes should ultimately directly translate into characteristics and features in the system architecture. That resource goes on to provide exemplars that can be used as a starting point for developing aircraft specific requirements. In some cases, engineering analysis may reveal that it is appropriate to tailor the attributes or potentially eliminate them entirely.

As a note of caution, it is important to emphasize that the CSAs are not a subset of RMF controls. Despite a strong resemblance in structure and nomenclature, CSAs are design objectives rather than controls. Consequently, the CSAs should never be interpreted as minimal security control implementation. Rather, those attributes are a starting point for identifying and tailoring appropriate cyber survivability and resiliency requirements.

On the commercial side of the aviation industry, the Federal Aviation Administration (FAA) has created the Aviation Cyber Initiative that includes a Cybersecurity Engineering team chartered with the development of requirements. However, those efforts are still in the nascent stages of

development at the time of publication of this guidebook. Consequently, while DoD standards are not compulsory or binding for those platforms, they are nevertheless a good place to start.

6.5 Mission Impact Assessment

The identification of cyber resiliency requirements hinges on insight into the mission critical functions of the aviation system. Sometimes security systems engineers are tempted to assume that the stakeholders intuitively know what the mission critical functions are, as well as the relative order of priority. That is not always a valid assumption. An incorrect understanding of the mission critical functions can have a detrimental effect on the cyber resiliency of the developed platform.

There is value in segregating the identification of mission critical functions into a dedicated analysis activity. Lockheed Martin created the Mission Impact Assessment (MIA) Guidebook that defines a repeatable process for the identification and prioritization of missions and mission essential functions for a system. (40) MIA is one of the critical first steps to gather and produce analysis that is foundational to understanding and developing cyber resiliency objectives for both legacy and new aircraft systems.

The MIA Guidebook provides an introduction to that process including the benefits of performing the analysis. Additionally, the steps required to execute the process are provided. It also includes a description of the artifacts produced by MIA, templates for those artifacts, and instructions for how to execute the process.

One of the primary benefits of using the MIA methodology is independence from the architectural implementation. Consequently, it focuses engineering analysis on mission performance rather than perceived vulnerabilities of the assumed or imposed underlying architecture. MIA produces exactly the kind of information that is necessary to select architectural and design mechanisms that support the resiliency objectives of the system.

Figure 10 presents a framework that is useful to consider when identifying mission critical functions for an aircraft platform. The highest priority missions can be found at the foundation of the pyramid, and all the other mission essential functions build from there. One of the earliest lessons of flight training is recognizable in the diagram. In an emergency, the highest priority for the pilot is to aviate, the next most important thing is to navigate, and the third priority is to communicate.

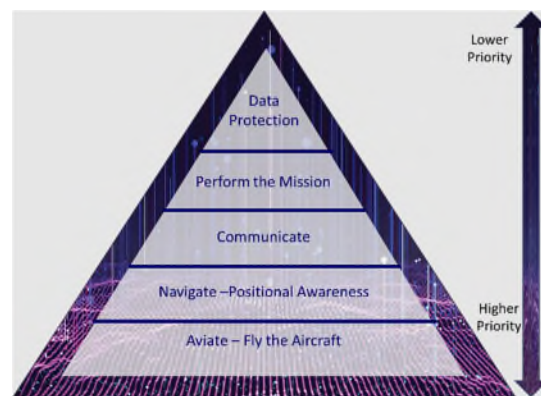


Figure 10. Aviation Platform Critical Mission Framework

The detailed mission essential functions will vary from aircraft to aircraft. Similarly, there are likely to be differences in prioritization of those functions based on the overarching mission that

the platform is intended to perform. Figure 10 is provided as a suggested starting point for that thought exercise.

MIA is an excellent method to elicit cyber resiliency and system survivability objectives for a system. It is also a valuable additional source of information for driving and validating architecture and design decisions.

6.6 Use and Abuse Cases

Use cases are a technique for defining and describing interactions or scenarios within a system. A use case will typically include a description of the interaction along with its goal. It will also specify the “actors” who participate in the scenario. The actor can be a human operator or other systems or subsystems.

A use case typically outlines a series of steps or actions performed by the actors and the system to accomplish a goal. These steps might include inputs from the actor, system processing, and outputs or responses. Depending on the precise methodology used, a use case may be visually expressed by a graphical diagram.

Use cases are important tools in systems requirements engineering because they help stakeholders understand how the system will be used and what functionality it should provide. By capturing various scenarios through use cases, requirements analysts can ensure that the system's design and implementation address the needs and expectations of its users effectively. Use cases are an effective mechanism to identify the foundational requirements necessary to drive subsequent phases of the system development lifecycle.

In contrast, an “abuse case” can be used to identify potential adversarial use of the system to identify cyber resiliency requirements. These scenarios capture potential situations a system may encounter by malicious insiders or within cyber contested airspace. Abuse cases consider how cyber threat agents and conditions might interact with the system to cause harm or to achieve adversarial objectives. Abuse cases are particularly useful when analyzing system requirements for survivability and cyber resiliency. This analysis is one of the most effective methods for identification and mitigation of potential security risks and threats during the earliest stages of the system development lifecycle.

Abuse cases play a vital role when ensuring that objectives for cyber resiliency in the face of intentional malicious actions are adequately identified during the requirements engineering phase of system development. It provides a structured approach to thinking about cybersecurity risks and helps drive system architectures and designs that are intrinsically resilient against malicious activities and unauthorized use.

Abuse cases are particularly important for aviation systems, as the consequences and impacts of successful cyber-attack can be catastrophic. Both use and abuse cases are essential for holistic systems engineering in aircraft design to ensure safety, reliability, and security throughout the lifecycle of the aircraft. It enables engineers to anticipate and address a wide range of operational and security challenges, thereby enhancing the overall cyber survivability, safety, and performance of aircraft systems.

6.7 Requirements for Security Sustainment

When developing an aviation platform, a Security Sustainment Concept of Operations (ConOps) should be created no later than the requirements engineering phase of the system development lifecycle. That way requirements for cybersecurity patching, continuous monitoring, and adaptive measures against emerging threats can drive architectural mechanisms to support that in the delivered system. In the face of a highly dynamic threat environment, it isn't realistic to neglect how the platform will be adapted and evolved to maintain an acceptable posture of cyber resiliency.

For aviation systems, the Security Sustainment ConOps should describe the anticipated operational context for the aircraft. That information can be used to detect when mission execution has diverged from what was originally intended. That may signal a need to revisit or at least understand assumptions that were made when the aircraft was developed. This is particularly important for aircraft systems and components that are reused from platform to platform. While a part that was developed for commercial aviation use might appear to be appropriate for military aircraft, it may not have been developed to the degree of cyber resiliency required for highly contested airspaces.

The strategy for updating software and configuration data should also be documented in the Security Sustainment ConOps. That should include mechanisms for changing cryptographic keys and any digital certificates used. Unless a part is fully attritable or disposable, mechanisms for quickly and securely making updates to it is an absolute necessity. Otherwise, an aircraft platform could be stuck with a part that has known vulnerabilities and no mechanism of mitigation.

Continuous monitoring of emerging threats and events in the operational environment is important for aviation systems. Data collection and analysis is essential for identifying emerging adversarial actions against specific airborne platforms. Adapting to the dynamic threat environment requires source material from which indicators of attempted cyber-attacks and compromises are derived. Onboard collection, coupled with continuous monitoring, is an essential source of that information. Consequently, provisions for meeting that need should be included in the Security Sustainment ConOps.

Security sustainment is a recurring necessity over the service life of every aircraft. Those essential functions must be included in the program requirements. That allows engineering and development personnel to include the critical mechanisms for establishing and sustaining an acceptable cyber resiliency posture for the platform.

Creating a Security Sustainment ConOps for aviation systems is a key enabler for the identification of requirements that help ensure that the aircraft can be continuously adaptable to the highly dynamic cyber threat environment.

6.8 Requirements for Cyber Active Defense and Response

Organizations with a mature understanding of cyber resiliency understand that skilled and well-resourced cyber threat actors will enjoy success from time to time. It is an imperative to

anticipate that systems may one day fall victim to a cyber incursion. Consequently, systems architects and designers must include mechanisms for cyber defense and response.

Cyber active defense is fundamentally more challenging for aviation systems than it is for enterprise IT networks. Business computing platforms and operating systems enjoy a large volume of data that illuminates known attack paths and patterns from a variety of threat agents. The amount of available information supports identification of statistically significant detection mechanisms of cyber-attacks. Additionally, commercial devices and services are available that provides some measure of protection against known attacks.

When a traditional enterprise IT system suffers a confirmed or suspected cyber event, cyber sophisticated organizations rely on human-centric technical expertise for defending the system. Incident response protocols and disaster recovery plans are typically well documented and practiced. A combination of human effort, augmented by security tools, work in concert to minimize damage, restore normal operations, and strengthen defenses against similar threats in the future.

Legacy aircraft do not typically have the same degree of cyber active defense and response mechanisms, processes, and procedures that are commonplace in traditional enterprise IT based systems. There is significantly less information and data on cyber-attack against aircraft from which to draw broad conclusions about what indicators to look for. Additionally, it is rare for an aviation platform to have tools and procedures for data collection and forensic analysis of confirmed or suspected cyber-attack.

Requirements for active cyber defense and response mechanisms should be considered at program inception. A minimum feature to consider for every aircraft and subsystem is a mechanism for forensic data collection in the event of a suspected cyber incident. That foundation of information is essential to build the understanding needed to accurately characterize indicators of cyber-attack against each aircraft.

During flight or ground operations, it would be reasonable to include a mechanism for a pilot or other aircraft operator to flag unusual aircraft behavior that might be indicative of activity within a cyber contested airspace. The ability to mark and snapshot those precise moments are invaluable for forensic analysis. Similarly, systems should have requirements for fast reboot and recovery mechanisms that restore the aircraft to a known good state both during flight and on the ground.

Ideally, aircraft will be intrinsically secure against aviation cyber threats. However, in the face of highly skilled and well-resourced nation state threat actors, it is inevitable that cyber adversity will eventually occur. Requirements anticipating malicious cyber-attack should be developed to position the platform to collect the data needed to inform future adaptive response.

6.9 Cyber Resiliency Adjacent Requirements

This chapter of the FSAD Guidebook focuses on eliciting system requirements that underpin cyber resilient aircraft systems. Under the general umbrella of cybersecurity requirements there are three additional areas that must be considered. These are general domains that frequently

appear on cybersecurity compliance checklists for DoD systems, though these are valid considerations for commercial aviation as well.

Encryption plays a vital role in protecting aircraft systems from cyber threats and supporting the confidentiality of critical data. Robust cryptographic measures enhance the security posture and mitigates risks associated with unauthorized access. Encryption is frequently mandated by regulation and standards for many systems.

With requirements for encryption comes the need for key management practices for generating, storing, distributing, and revoking cryptographic keys used in aircraft systems. Consequently, it is an imperative to include requirements for each of those functions in the system. It is not acceptable to hard code encryption keys based on the assumption that update or modification will never be needed.

Anti-Tamper (AT) is a systems engineering discipline focused on the prevention or delay in exposure or exploitation of Critical Program Information (CPI) for DoD weapons systems. Like cyber resiliency, AT must be considered over the entire life cycle of system development. AT requirements typically include deterrence mechanisms against reverse-engineering, exploitation, and protection of intellectual property.

While commercial aviation platforms are not ordinarily subjected to DoD AT requirements, it is still a best practice to make it more difficult for a cyber adversary to learn aircraft implementation details. That knowledge and insight can be a precursor to developing malicious cyber-attacks. Additionally, AT techniques can protect a manufacturer's intellectual property.

Sophisticated engineering analysis is needed for aircraft systems that have both AT and Cyber Resiliency requirements. Tradeoffs are frequently necessary between the objectives of those two domains which can bring conflicting objectives to the platform. For example, ceasing processing under some conditions might be a mechanism selected to satisfy an AT confidentiality requirement. However, that might be the opposite of a cyber resiliency requirement to continue processing to execute the mission no matter what. Such tradeoffs are the essence of engineering.

TEMPEST is a National Security Agency (NSA) term that refers to the susceptibility of cyber enabled systems to expose information through emission of electromagnetic radiation. Requirements to protect against such leaks are frequently levied against aircraft systems that process DoD classified data. TEMPEST requirements might include separation, shielding, filtering, and masking.

In the pursuit of cyber resiliency, requirements from these cybersecurity adjacent domains cannot be neglected. However, it is also critical to understand that implementing some of those requirements will likely create a need for trade-offs within the system. DoD aviation platforms are much more likely to have these requirements imposed than commercial aircraft. However, the concepts are reasonable considerations for both domains.

6.10 Aviation Cyber Threat Intelligence

Policy directives and guidelines for military weapons systems consistently claim that cyber resiliency requirements must be derived from threat intelligence. That source information is even called out as a key process step in the CSE Implementation Guide. The identification of system-specific threats, as well as the capability of the anticipated cyber adversary, are cited as key pieces of necessary information for eliciting system requirements.

The unfortunate reality is that aviation specific cyber intelligence is typically somewhere between scant and nonexistent. Even when it exists, classification and operational security frequently prevent that information from being freely shared. Insight provided by threat intelligence is ideal. However, the absence of that data can never be used as a justification to discount requirements for cyber resiliency against malicious activity.

Chapter 2 built a case for understanding that the threat against our aviation platforms is real. Applying the insight from that chapter specifically to cyber threat intelligence for aircraft distills down into three fundamental truths:

- 1. All aviation platforms operate in or will be operated in cyber contested environments.**
- 2. The capability of aviation cyber adversaries constantly improves.**
- 3. The ease of exploit against aviation vulnerabilities gets consistently easier.**

The United States military recognizes cyberspace as a warfighting domain. Consequently, during times of conflict, our adversaries can be expected to deploy offensive cyber-attacks against military aircraft intended to generate a decisive advantage in the battlespace. Additionally attacks against civilian critical infrastructure is also likely in event of military conflict. The Transportation Security Agency (TSA) checkpoints at commercial airports are a testament to the belief that non-military aircraft may be targeted by adversarial activity.

It would be naive to believe that the capability of any cyber adversary will remain static over time. The offensive capabilities of the most sophisticated nation state threat actors are always increasing. Those entities invest in technological advancement and continuously add to their skills and knowledge.

At the same time, vulnerability disclosures, security patches, advisories, and published information from security researchers continually lowers the bar for exploiting vulnerabilities that may have initially required extreme skill. Today's nation state capabilities can become accessible by the lowest skilled threat agents in the future. In fact, that occurs on a frequent basis for enterprise IT systems.

Additionally, cyber-attacks may be used as a precursor to escalation of traditional kinetic warfare. The cyber domain can be leveraged for information and intelligence gathering well in advance of declared conflict. Cyber intrusion is already known to be in widespread use for espionage operations.

When considering cyber resiliency for aircraft, it is important to remember that these platforms have extremely long service timelines that exceeds what is normally associated with traditional IT systems. It is not unusual in the commercial and military environments for aircraft to be used for decades. That prolonged service life pushes understanding the current capabilities of cyber adversaries toward irrelevance. The developers of aviation systems must understand that the platform must be architected and designed for intrinsic cyber resiliency against the absolute best cyber adversaries that will only get better and better over time.

When available, both general and specific threat intelligence for aircraft should be used as a basis for determining how to implement cyber resiliency on each platform. That should inform the Concept of Operations, Capability Descriptions, and program requirements. Those artifacts should drive cyber resilient architectures and designs throughout the service life of each platform. However, all those things must be done regardless of whether or not actionable intelligence exists.

Additionally, when threat intelligence is available, it should be regularly reviewed to determine if corrective action or mitigation is required due to the continuously evolving threat landscape. Those assessments should consider the operational mission impacts so that all stakeholders understand the evolutionary nature of cyber risk.

6.11 The Iterative Nature of Cyber Resiliency Requirements

Developing mature cyber resiliency requirements is a continuous process that permeates every stage of the development lifecycle. That is particularly the case when using agile or iterative development approaches. However, it is equally true for development using the traditional waterfall model.

This chapter has focused primarily on requirements elicitation, analysis, and management that occurs during program inception and the requirements phase of development. That is when stakeholders collaborate to identify and define critical cyber resiliency needs based on the aircraft's mission, operational context, and potential threats within cyber contested airspace.

As development progresses through the design, implementation, and testing phases, continuous analysis occurs to refine and validate these requirements in alignment with evolving technologies and emerging cyber threats. Additionally, selected architecture and design implementations should be continuously scrutinized for additional cyber resiliency objectives that may be introduced through those decisions.

These feedback loops are established to incorporate lessons learned from previous phases, ensuring that the cyber resiliency measures remain effective and adaptable as the aircraft design matures. This iterative approach fosters a dynamic and responsive development environment where cyber resiliency requirements are a critical aspect of safeguarding systems against an ever-changing threat landscape.

Additionally, cyber resiliency requirements must be re-evaluated and assessed as new capabilities are added to the platform and as parts and subsystems are upgraded. Every change to a system introduces potential for new vulnerabilities or attack surface that impacts the cyber

resiliency posture of the aircraft. It also presents an opportunity for making improvements to the cyber resiliency requirements and objectives.

6.12 Additional Reading and Resources

The following additional reading is recommended for people who wish to go into greater depth on cyber resiliency requirements at program inception.

- **Cyber Survivability Endorsement (CSE) Implementation Guide, Joint Staff J6, Deputy Director for Information Warfare, Requirements Division, Version 3.0, July 2022.**

The Cyber Survivability Endorsement (CSE) is a System Survivability Key Performance Parameter (SS KPP) of the Joint Capabilities Integration and Development System (JCIDS) Manual. CSE targets the predictable failure of the cybersecurity processes to build-in sufficiently robust cyber capabilities to prevent (resist/anticipate), mitigate (absorb/withstand), recover from, and adapt to the full spectrum of mission assurance cyber-events in plain language requirements for program management.

- **Systems Engineering Handbook: A Guide for System Life Cycle Processes and Activities, International Council on Systems Engineering (INCOSE), INCOSE-TP-2003-002-05-2023, Fifth Edition.**

This book describes the key process activities performed by systems engineers and other engineering professionals throughout the system lifecycle. It is the authoritative source of a wide range of fundamental system concepts that broaden the thinking of the systems engineering practitioner, such as system thinking, system science, life cycle management, specialty engineering, and system of systems. It covers all the major lifecycle models including waterfall, V model, agile, and iterative methods. This book is ideal for anyone who has a need or interest in applying systems engineering practices.

- **INCOSE Guide to Writing Requirements, International Council on Systems Engineering (INCOSE), INCOSE-TP-2010-006-4, 1 July 2023.**

This document describes how to express need statements (needs) and requirement statements (requirements) clearly and precisely in textual form to support analysis and implementation, independent of any systems engineering (SE) tool that may be used to capture and manage those needs and requirements throughout the system lifecycle. It supports the development of clear, concise requirements.

- **Mission Impact Assessment Guidebook, Lockheed Martin, PIRA #AER202007005, Version 4.3, August 18, 2020.**

The Mission Impact Assessment (MIA) Guidebook describes a process developed and used by Lockheed Martin (LM) to perform identification and prioritization of missions and mission essential functions for a system. MIA is one of the critical first steps to gather and produce analysis that is foundational to understanding and developing cyber

resiliency objectives for both legacy and new development programs. Effective Cyber Risk Assessment (CRA) relies on the analysis produced by MIA to optimize, guide, and inform those efforts. This document provides an introduction to the MIA process including the benefits of performing the analysis.

Chapter 7

Cyber Resilient Architecture and Design

The intrinsic property of cyber resilience in any system is best achieved with intentional efforts during the architecture and design phases of development. That is especially true for aviation systems with the potential for catastrophic consequences from cyber incidents. Aircraft manufacturers must design systems that can successfully operate in the face of cyber threats against the platform.

This chapter describes concepts for integrating and prioritizing cyber resilience when architecting and designing aviation systems. The techniques include ways to deliver functions and capabilities that are intrinsically cyber resilient. That is done through methods to minimize vulnerabilities introduced by the architecture and design. In many cases, a cyber resilient system will also include architectural features that exist solely to support the cyber resiliency objectives of the aircraft.

7.1 Introduction to System Architecture and Design

The Systems Engineering Body of Knowledge (SEBoK) describes a system architecture as a comprehensive solution to fulfill a purpose. (41) The activities that create a system architecture are based on timeless principles and concepts of good design. A mature system architecture has features, properties, and characteristics which satisfy its set of system requirements and lifecycle objectives.

The system architecture for an aircraft is a high-level conceptual structure of how it will ultimately achieve its mission objectives. It encompasses the subsystems, relationships, and principles that guide how each component will be designed and operated. It also includes how various parts of the system interact with each other and the data and control flows that are necessary to make it all work.

A well-defined system architecture is crucial for designing, implementing, and maintaining complex systems to ensure functional requirements, performance objectives, and quality standards are achieved. That is particularly important for aviation systems. Effective system architecture promotes modularity, flexibility, scalability, and maintainability of the aircraft throughout its lifecycle.

While system architecture and design are closely related concepts, they refer to different aspects of the system design process. One of the key differences is the level of scope and abstraction. Architecture focuses on the structure at the highest level. It is the overall conceptual framework for the complete system. In contrast, design is the translation of the architectural blueprint into

detailed specifications for system implementation. It is a much more detailed view of the technical features that elaborates on the architectural concepts.

The detailed design specifies how the system is to be implemented by identifying components, interfaces, algorithms and data structures. It breaks down the architecture into smaller modules each with a defined purpose and set of system responsibilities. This is where decisions are made to address technical challenges, performance optimization, and ensuring the aircraft works as a holistic system.

Detailed attention to the architecture and design is essential for creating robust, scalable, and maintainable aviation systems.

7.2 Architectural Allocation of Cyber Resiliency Requirements

In aviation systems, allocation and decomposition of cyber resiliency requirements should essentially follow the same process used for non-cyber related requirements. That includes the same level of detail and rigor as is done for capabilities and functional requirements. The process typically involves breaking down high-level cyber resiliency requirements into more specific, detailed requirements that can be allocated to low-level system components within the architecture. This allows for a systematic and consistent approach to addressing cyber resiliency throughout the system development lifecycle.

This architectural allocation will pay careful attention to how the aircraft preserves the critical missions and manages access and availability of computing resources. This can be achieved through various architectural techniques for protecting high priority mission critical system functions, even under adverse conditions.

Following the same process for both cybersecurity and functional requirements helps to ensure that cyber resiliency is integrated and considered in all aspects of the system, rather than being treated as an afterthought or separate concern. By following a consistent process, organizations can ensure that cyber resiliency is effectively addressed and built into the system from the beginning, resulting in a more secure and resilient system overall.

7.3 The Foundational Objectives of Cyber Resiliency

A robust architecture and design for any aviation system necessarily includes support for all features and capabilities that the platform is expected to provide. While those functions should be constructed to eliminate or minimize vulnerabilities to the greatest extent possible, that is usually not enough. A highly cyber resilient architecture will also include mechanisms that only exist to support the cyber resiliency requirements and objectives for the system.

Systems architects and designers must apply fundamental best principles that support cyber resiliency. These foundational concepts described in this section are a composite view of the recommendations from the best current policy guidance, augmented with experience from Lockheed Martin aviation development engineering programs. In particular, the Cyber Survivability Endorsement (CSE) Implementation Guide (38), the NIST Cybersecurity Framework (42), and NIST SP 800-160 Volume 2, “Developing Cyber Resilient Systems: A

Systems Security Engineering Approach” (43) each provide excellent information for architecting and designing cyber resilient platforms.

The first fundamental principle is to create systems that **prevent** successful cyber-attacks or incidents from ever occurring in the first place. While that is a necessary and noble objective, it is also understood that highly capable nation state cyber adversaries will experience some success from time to time. The cybersecurity industry has yet to produce a system that is completely impervious to cyber-attack.

Consequently, the second fundamental principle of cyber resiliency is to create systems that can **mitigate** the effects of adversarial cyber adversity when that inevitably occurs. That leads directly into the third and fourth principles. The system must be designed to include methods necessary to **recover** to a known good state. Finally, mechanisms and procedures must be developed so the system can **adapt** to be less susceptible or to prevent known attacks from succeeding in the future. The knowledge and insight gained closes the loop by informing future efforts to **prevent** cyber adversity.

These principles of cyber resiliency are most effectively implemented when they are built directly into the architecture and design of the system. Prioritizing these principles when synthesizing the system architecture ensures that cyber resiliency is a fundamental intrinsic characteristic, rather than an afterthought.

7.3.1 Prevent

It is of paramount importance to architect and design systems that protect the aircraft’s mission essential functions from the threat of cyber adversarial action. Quite simply, a cyber attacker cannot successfully exploit a vulnerability that does not exist. Minimizing vulnerabilities requires intentional analysis as the aircraft is conceptually synthesized. Architects and designers must be cognizant that decisions during this phase of development could have profound consequences on the operational resiliency of the aircraft.

It is essential to **identify** the mission critical functions that must be preserved along with the assets and data that must be defended. That helps architects and designers understand where more stringent protection is needed within the system.

Additionally, it is important to **anticipate** that the aircraft will operate in cyber-contested airspace, as well as the types of adversity impactful to the platform. Deferring consideration of cyber threats until evidence of malicious activity occurs significantly increases the risk to the aircraft. Waiting for evidence of cyber-attack is too late to instill the characteristics of intrinsic cyber resiliency.

Finally, the architecture and design must **protect** the system resources to the greatest extent possible. Multiple security measures can be used to increase the resiliency posture of the platform. Even if one mechanism is defeated, others in place may prevent the cyber adversary from achieving success.

In 1736, Benjamin Franklin advised fire-threatened people in Philadelphia that “an ounce of prevention is worth a pound of cure.” (44) The same principle applies to engineering intrinsically cyber resilient systems. The best way to prevent successful cyber-attack is to identify what needs defending, anticipate adversarial action, and to protect what is mission critical within the system.

7.3.2 Mitigate

A common dilemma of the cybersecurity domain states that to be successful, system defenders must defeat every attack every time while the attackers only have to prevail once. That asymmetrical truth underpins the understanding that highly capable well-funded nation state adversaries will experience success from time to time, even when attacking systems with a secure underlying architecture. While striving to build systems that prevent all cyber-attacks is ideal, it is equally important to recognize that mechanisms are required to mitigate the effects of successful attacks in the event they occur.

Aircraft operators and designers must architect systems that mitigate the effects of successful cyber-attacks. That is typically accomplished by implementing design principles that enable the mission essential functions to survive attacks and preserve mission execution.

One of the most challenging aspects of system design for aircraft is to **detect** a cyber-attack as it occurs. In fact, if it is known what any specific attack against the platform looks like, it may be easier to eliminate the susceptibility to the attack rather than implementing a detection mechanism in many instances.

However, there is tremendous value of identifying and collecting data of suspected adversarial cyber activity, regardless of whether it produces a cyber effect on the platform. It is incumbent on architects and designers to collect the information needed for subsequent analysis of potential cyber-attack attempts, even when the end state objective against the aircraft failed.

Another key principle of resilient design for aviation systems is to **respond** by taking appropriate action when a cyber event occurs. The term “event” was intentionally selected rather than “attack” in the preceding sentence. For aircraft platforms, observed cyber effects may result from actions initiated by malicious threat actors. However, it is also likely that those effects may have resulted from an error or unusual transient condition on the platform. Cyber resiliency means that mission critical processing continues, regardless of the source of disruption.

Consequently, responding to cyber events on aviation platforms closely resembles fault tree analysis traditionally performed for safety. That analysis is augmented by potential malicious actions which may have previously been excluded from consideration on legacy systems. Every aircraft should be architected and designed with consideration of all things that can go wrong along with the actions that should be taken to restore full mission capability. By doing that, platforms can be designed to **withstand** many types of cyber events.

7.3.3 Recover

The ability to **recover** from a cyber-attack and restore the system to full operational capacity is a necessary design principle for aviation systems. In many instances, the recovery mechanisms

may closely resemble those identified for responding to or withstanding attacks described in the previous section. However, there is typically a subtle difference between the actions used during operational use to preserve mission execution and those that are subsequently taken to restore the platform to full capacity. Recovery includes both operational and ground-based actions that fully restore resilience and confidence in the platform for future mission execution.

7.3.4 Adapt

Engineering for cyber resiliency is a continuous process that never ends throughout the lifecycle of each platform. Data collected through operational use must be continuously analyzed to **adapt** the aircraft system for cyber resilience against new and emerging threats.

A full lifecycle approach will collect data from potential cyber events that were detected to determine if it was a real or potential cyber-attack. That powers the insight which identifies necessary updates to the system to prevent, mitigate, respond, and recover from future recurrence. Aviation platforms must be architected and designed to adapt to things which may occur in the environment. That includes mechanisms for updating software, changing configuration settings, and potentially implementing new response mechanisms.

7.4 Additional Objectives of Cyber Resilient Architecture and Design

This section shares important aspects of the philosophy that enhances the foundational principles identified in Section 7.3. These additional underlying concepts should guide and shape the creation and development of aviation systems. That includes the approach, decisions, and priorities that are made throughout the aircraft design process. At its core, these objectives reflect a commitment to holistic cyber resiliency on the platform.

7.4.1 Minimizing Cyber Attack Surface

In cybersecurity, the concept of attack surface is the totality of all possible points, avenues, and vulnerabilities through which a malicious actor can attempt to compromise a system. It includes both electronic cyber-enabled vectors as well as human error, social engineering, and physical access. Minimizing the attack surface is a fundamental objective during architecture and design of any system. While it is usually not feasible to eliminate all potential means of access and attack, great care should be taken to avoid creating attack vectors that do not serve a necessary function within the system.

Understanding and managing the attack surface is also critical during operational usage. In many instances, systems architects and designers should anticipate that operators of their systems might need methods for shutting down functionality of a subsystem, or other response mechanisms, when a cyber-attack is confirmed or suspected. As the state of the practice in offensive cyber evolves, it is likely that the attack surface on every targeted system will expand. That means that an ongoing sustainment effort is necessary to maintain a robust security posture.

Minimizing the cyber-attack surface is especially crucial for aircraft due to the significant safety implications involved. Aircraft systems are reliant on digital technologies for navigation,

communication, and flight controls. While external interfaces to each of those functions are typically required for platform operation, the architecture and design should consider the potential for malicious activity. All information and signals received by an aircraft should be regarded as possibly hostile.

Exposure to potential cyber-attack can occur both during flight and maintenance ground operations. Each access point could be exploited by malicious actors to disrupt operations, compromise safety, or gain unauthorized access to critical controls. That could endanger human life and compromise mission execution.

Reducing the attack surface of aircraft systems through rigorous cybersecurity engineering is paramount to ensuring the safety and integrity of air transportation. Doing that successfully can significantly decrease the overall risk exposure. Those efforts start during architecture and design.

7.4.2 Design Anticipating Nation State Cyber Adversarial Activity

Each of the published Cyber Risk Assessment (CRA) methodologies described in Chapter 9 require identification of likely threat actors against the system as well as an estimate of their capabilities. For enterprise IT systems, that is a key source of information that can help identify the most likely and impactful attacks against the platform. That insight can help program managers optimize decisions on resource allocation to mitigate the highest priority risks.

For aircraft, all stakeholders must understand that these platforms are subject to adversarial cyber-attack from well-resourced and highly capable nation state threat actors. Both military and commercial aviation systems are a vital part of national critical infrastructure. Aviation is a key part of the transportation sector essential to both the physical and economic security of the nation. Military aircraft are crucial to project power and defend national interests.

Hostile cyber-based attacks against both commercial and military aviation systems are a significant threat to national security. Consequently, aircraft architects and designers must account for motivated, well-funded, and highly capable attackers. In addition, it is important to remember that those actors will only get better over time. That is particularly important when building aircraft due to the long service lifespans particularly when compared with traditional enterprise IT systems.

It is never acceptable to wait for intelligence or for confirmed attacks to emerge. By the time that happens, disaster has already occurred. Aircraft systems must be architected and designed with an understanding that nation state level attacks are most certainly forthcoming.

7.4.3 Implement Defense in Depth and Coordinated Protection

There is no single cyber defense mechanism or control that can prevent a determined adversary from occasionally experiencing success. Consequently, it is important to implement a strategic integration of multiple defense mechanisms to safeguard systems. In other words, architects and designers should anticipate the failure of individual cybersecurity mechanisms, yet still build systems that can accomplish the mission critical functions despite that.

Defense in depth can be thought of as the layers of an onion, each adding an additional level of protection. That requires attackers to defeat more than one protection mechanism to compromise the system. It reduces the likelihood of a single point of failure that puts the entire aircraft at risk.

Coordinated protection augments the concepts of defense in depth with an understanding that each protection mechanism implemented within a system increases the possibility of unanticipated side effects that impact cyber resiliency. For example, a protection mechanism that automatically shuts down an external communications interface when unusual messages are detected could create a self-induced denial of service that impacts mission execution. Consequently, the protection mechanisms implemented to defend the platform must be implemented in close coordination with each other and always with an eye toward cyber resiliency.

7.5 Techniques for Cyber Resilient Aviation System Architecture and Design

An archetype is a pattern, model, or template that embodies a solution to a recurring design problem or requirement. Archetypes can be thought of as reusable building blocks that help architects and designers streamline the development of complex systems by providing proven solutions to common challenges.

Leveraging archetypes can promote using best practices and potentially reduce the time required to create architectures and designs. These patterns can also promote consistency, scalability, and maintainability in systems, reduce development time and costs, and facilitate knowledge sharing and collaboration among architects, designers, and developers. Archetypes also help ensure that systems are designed and implemented in alignment with standards and architectural principles.

Many security controls recommended or required for the Assessment and Authorization (A&A) processes described in Chapter 12 can be thought of as candidate archetypes for a system. However, the definition of those controls is heavily influenced by usage within enterprise IT networks. In many cases those security mechanisms will not be directly applicable to aircraft platforms. The architects and designers must approach traditional security controls with the understanding that modifications are frequently needed for real-time aviation requirements and constraints.

One of the best resources of archetypes for aviation systems is “Developing Cyber-Resilient Systems: A Systems Security Engineering Approach.” (1) That NIST special publication identifies cyber resiliency techniques expressed as goals, objectives, approaches, and design principles. Those descriptions were leveraged as a basis for some of the recommendations within this section of the FSAD.

Sometimes the archetypes of cyber resilient design are collected into a reference architecture. That is a curated set of standardized components, relationships, and interactions that represent an idealized solution for a specific domain or application. A reference architecture can be thought of as a recipe book used for system and subsystem design.

The reference architectures specific to the aviation domain are typically closely held as proprietary information. In the absence of a public well-established standard, manufacturers and

suppliers should curate their own best practices into a reference architecture. In some cases, suppliers may be able to gain access to a useful and applicable reference architecture from the prime contractor of their system.

The remainder of this section describes the high-level archetypes that can be leveraged within aviation systems. These are loosely organized under the fundamental principles of cyber resiliency previously defined in Section 7.3, which are to prevent, mitigate, recover, and adapt to cyber adversity. However, it should be noted that this organization is less than perfect because some of these archetype techniques can be used to satisfy more than one of those principles.

Consequently, it is the responsibility of aviation system architects and designers to be familiar with the foundational and fundamental objectives, as well as the techniques that can be used to achieve them.

Additionally, these subsequent sections describe the best practices as understood when this FSAD guidebook was published. It is both possible and desirable that the body of knowledge for the techniques will continue to grow. The most skilled architects and designers should strive to contribute to the state of the practice in the industry.

7.5.1 Techniques for Preventing Cyber Adversity

Ideally, aviation systems will be impervious to cyber-attacks. However, it is not possible to create platforms guaranteed to resist currently unknown attacks that may emerge in the future. Nevertheless, platform architecture and design should implement “Prevent” security mechanisms that strive for the objective of preventing successful cyber-attacks from occurring.

7.5.1.1 Privilege Restriction

Privilege restriction is the practice of limiting the access rights and permissions granted to users, processes, or components within the system. Those limitations can be based on roles, responsibilities, trust levels, or some combination of those factors. Each user, process, or component within a system should only be granted the minimum access rights and permissions required to achieve legitimate mission objectives. Least privilege is one of the most important archetypes to implement on an aviation system.

The principle of least privilege has not consistently been implemented on legacy aircraft. That omission was based on the faulty assumption that all users, communication interfaces, software, and hardware components could be trusted to not intentionally commit malicious cyber actions. Our modern understanding of cyber contested environments and supply chain risks have dispelled that notion.

Privilege restriction limits the attack surface available to an adversary with access to a particular part of the system. For example, compromise of a subsystem through supply chain interdiction limits actions and effects from that component to only authorized actions. At the access points where users interact with the aircraft, privilege restriction can prevent the specific threat of insiders taking actions that exceed their authority.

Privilege restriction is a fundamental security measure that should be implemented in all aviation systems to limit the attack surface available to adversaries and ensure the safe and secure operation of aircraft.

7.5.1.2 Segmentation and Partitioning

Segmentation and partitioning are effective techniques for protecting aircraft systems from cyber adversity by creating isolated and compartmentalized subsystems. To the greatest extent possible, aircraft systems should be isolated from each other based on mission criticality, data sensitivity, or functional role within the system. This isolation prevents cyber adversity in one part of the system from propagating to other areas.

Examples of segmentation and partitioning is reflected on many aviation platforms as flight controls are isolated from less safety critical systems. That historic best practice allowed aircraft architects and designers to optimize performance for parts of the system that were less critical to safety. Similarly, this archetype can be used to support and enhance cyber resiliency by isolating mission critical functions and resources from other parts of the system with lower priority.

While this design pattern might lead to the idea that all mission critical functions should be isolated from each other, that isn't always feasible within an aircraft. Implementing multiple physical networks and processing elements will have an impact on the weight and power requirements on the platform. Design tradeoffs will frequently be encountered between segmentation and performance.

7.5.1.3 Unpredictability

Unpredictability can be an effective cyber defensive mechanism because indeterminant behavior increases the difficulty of successful attack. It makes it harder for threat actors to identify vulnerabilities and determine how to exploit them. Unpredictability complicates offensive cyber operations by increasing the time required to develop attacks, while simultaneously decreasing the reliability of exploits developed.

Many best practices in enterprise IT systems introduce some measure of unpredictability into systems. Simple techniques such as changing passwords and varying the time that security actions are performed makes it difficult for an attacker to reuse or bypass protection mechanisms.

Unfortunately, unpredictability is a generally undesirable attribute for aircraft systems due to safety concerns and airworthiness certification requirements. An aircraft is a complex system that relies on precise control and determinate behavior. One of the non-cyber design principles for aviation systems is to minimize unpredictability.

Consequently, the standard methods and approaches to implement unpredictability in enterprise IT should be approached with caution. Tradeoffs will be required between the desired security posture of the system balanced out against safety and performance requirements.

The definition for the unpredictability archetype described in NIST 180-160 Volume 1 (30) describes it as making changes randomly or unpredictably to create and maintain an environment of uncertainty for the adversary. It recommends techniques such as reauthentication at random intervals or performing routines at different times of the day.

At the user level, random periodic reauthentication is not likely to be accepted by the airworthiness or pilot communities. A pilot should never be prompted to reenter a password when performing intellectually intensive flight activities such as takeoff and landing. For military aircraft, the heat of battle would be a similarly inopportune time.

While unpredictability can be an effective cyber defensive mechanism, it may not be a suitable approach for aircraft systems. It is important to carefully consider the tradeoffs and any potential impacts on system performance and safety. As a result, it is essential to carefully evaluate the potential benefits and drawbacks of unpredictability before committing to use of this archetype in this domain.

7.5.1.4 Dynamic Positioning

Dynamic positioning distributes and relocates system resources to create a moving target for cyber attackers. Continuously adjusting the configuration and behavior of system resources makes it more difficult for a hostile cyber actor to successfully carry out an attack.

The benefits of dynamic positioning are similar to those of unpredictability. It can make it more time-consuming and costly for an attacker to carry out an attack. However, dynamic positioning does not necessarily have to be unpredictable or indeterminate. An aviation system architecture might be synthesized with several defined modes of operation that distributes cyber enabled resources across processing elements within the platform. Each of those modes can be independently tested for determinate behavior. It is important to carefully consider the tradeoffs and potential impacts on system performance and usability when implementing dynamic positioning. While that increases the cost and complexity of integration testing, it provides a possible avenue to satisfy airworthiness constraints.

Another approach to implement dynamic positioning is distributed processing across multiple components. That makes it difficult for a cyber adversary to completely deny a capability because multiple coordinated cyber effects would be required to succeed against all processing locations.

The archetype of dynamic positioning must be selected during the development of the architecture and design of the aircraft system. Both the efficacy and feasibility drop precipitously after that phase of development is completed.

Many aviation systems have a history of implementing degraded modes of operation within the architecture and design of platforms. Historically, that was done for safety in response to a malfunction or failure. However, degraded modes of operation using dynamic positioning can also be an effective mitigation to offset cyber adversarial action on the platform.

When using dynamic positioning and predefined degraded modes, it is important to carefully consider the tradeoffs and potential impacts not only on system performance, but also on

integration and testing complexity. To be fully effective, this archetype is best selected as the architecture and design of the system is conceptualized.

7.5.1.5 Substantiated Integrity

Substantiated integrity is the use of cryptographic checksums to provide assurance that data, systems, and processes, have not been modified. In the C-I-A triad introduced in section 5.2, integrity is one of the core properties alongside confidentiality and availability. Integrity ensures that data remains accurate, complete, and unaltered throughout its lifecycle.

Though they are related concepts, substantiated integrity is considerably more robust than non-substantiated integrity techniques such as parity checks and checksums. While those methods can be used to assess the integrity of data, those only protect against accidents and errors. An adversary with the sophistication to tamper with software or data on an aviation system will also be able to easily recalculate parity and checksum values.

Substantiated integrity uses keyed cryptographic methods to verify the authenticity and integrity of data. An attacker without the appropriate cryptographic key cannot generate a digital signature that matches the information. Substantiated integrity provides a much higher level of confidence in the accuracy and authenticity of data, as it is resistant to tampering, modification, and other types of attacks.

Due to the data intensive nature of aviation systems, the substantiated integrity archetype should be leveraged extensively. Candidates for use include all software and configuration data consumed by the aircraft platforms. Whenever possible, external communication interfaces should be protected using this technique to ensure the authenticity and integrity of data to protect the aircraft from external spoofed messaging attempts. Along those same lines, substantiated integrity should also be considered for internal messaging on the aircraft. That protects subsystems from attacks should another subsystem be compromised.

Substantiated integrity is closely related to the concept of a "root of trust," which is a fundamental aspect of secure systems described in section 14.6.3. A root of trust is a foundation of trust that is built into a system, providing a secure foundation for all other security measures. It is typically implemented in hardware or firmware and is used to establish a chain of trust for the system's security functions.

Substantiated integrity is built on the strength of the root of trust on the platform, which provides a secure foundation for implementing the substantiated integrity checks. This is how the aircraft knows that the keying material used for verification of substantiated integrity is valid.

Substantiated integrity provides a high degree of trustworthiness that data, software, or a system hardware component has not been illicitly modified. When those integrity checks pass, the platform can operate with high confidence that it is processing good software and data. When substantiated integrity failures occur, it is a possible indication of cyber activity.

7.5.1.6 Encryption

Encryption is an effective technique to protect sensitive data and communications from unauthorized access. Encrypting aviation software can prevent adversaries from leveraging inappropriately leaked or disclosed information to identify potential vulnerabilities and to devise exploits. When data communications are encrypted, it is much more difficult to craft malicious messages or to spoof data sent to a platform. Consequently, preserving the confidentiality of sensitive data, can also prevent some forms of cyber-attack.

Unfortunately, encryption can cause availability issues if not properly implemented, or keying mismatches occur during operation. Additionally, encryption can introduce additional processing overhead, which can slow down the performance of the system and cause delays in the exchange of data. This can be a particular concern in real-time systems, such as those used in aviation, where even small delays can have significant consequences. In those cases, encryption sometimes can have a negative impact on cyber resiliency.

Encryption can be implemented using either symmetric or asymmetric keying material. In symmetric encryption, the same key is used for both encrypting and decrypting the data. This means that the sender and the receiver must both have a copy of the same key, which requires a secure method of key distribution. Symmetric encryption is generally faster and more efficient than asymmetric encryption but comes with considerably more overhead and risk associated with preserving the confidentiality of the keys themselves.

In asymmetric encryption two different keys are used for encrypting and decrypting the data. This allows the sender and the receiver to use different keys and eliminates the need to securely transmit the key between them. Asymmetric encryption is generally slower and less efficient than symmetric encryption, but it provides a higher level of security, as the private key is not shared and is thus more difficult for an attacker to compromise.

Both symmetric and asymmetric encryption are commonly used in aviation systems to protect sensitive data and communications. When used, architects and designers must consider how key distribution and key management will be performed in the operational environment. Additionally, every aviation system that implements encryption must include a mechanism for quickly changing crypto keys in the case of compromise.

Encryption should be considered for aviation systems, as it can help protect sensitive data and communication from unauthorized access. That is crucial in cyber contested environments. However, encryption is not without its challenges. This archetype may create unacceptable processing loads or latency in the system. Additionally, architects and designers must account for key distribution and management in the overall aviation system.

7.5.1.7 Redundancy

Redundancy is the duplication of critical components or functions to ensure system availability and reliability. Aviation systems have a rich history of using redundancy due to the safety critical nature of flight with potentially catastrophic consequences. The use of redundancy allows critical functions to continue to operate even in the event of a failure by moving processing to an

alternative device. The technique reduces risk by ensuring that aircraft can continue to fly safely and efficiently.

The redundancy archetype duplicates critical components or functions within the system. The elements selected are those that are necessary to continue a determined threshold of performance should a subsystem or piece of hardware become unavailable. Aircraft platforms are frequently built with some degree of redundancy to support safety of flight. Increasingly, this technique is used to preserve components required for mission performance, particularly on DoD aircraft.

Another approach for redundancy on aviation platforms is correlation of data from diverse sensors or instruments. For example, some aircraft have multiple sources of navigation data that should be aligned. When those multiple sources are available, the system should be architected to determine which one is the most trustworthy.

Redundancy should be implemented strategically, targeting critical subsystems or components where failure is most impactful. On aviation platforms, the tradeoff for implementing redundancy includes cost, weight, and power considerations. The decision on whether and where to implement this archetype is complex.

7.5.1.8 Diversity

Diversity within a cyber-enabled system refers to the intentional variation of components to enhance system resilience, security, and robustness. In the context of cybersecurity, diversity aims to mitigate risks associated with single points of failure, common vulnerabilities, and targeted attacks. This archetype can manifest in several forms, such as employing assorted hardware and software components, utilizing varied communication protocols, implementing multiple authentication mechanisms, and varying data storage and processing methods. By diversifying these aspects of a system, it becomes more difficult for adversaries to exploit common weaknesses or launch widespread attacks, thereby bolstering the system's overall security posture. Additionally, diversity can also enhance system adaptability and flexibility, enabling it to withstand unexpected challenges or changes in the operating environment.

The diversity archetype can be used to enhance the effectivity of redundancy implemented in the system. For example, having two identical processing subsystems as a mitigation against cyber-attack is not as effective if both are susceptible to the same vulnerability. Using a different operating system or communication protocol are examples of diversity techniques that might reduce that risk. At the same time, that will certainly impact the cost of procurement, development, and integration testing.

Another way to support diversity within system architecture and design is to create isolated segmented networks and processing environments. While that decreases the opportunity for data sharing and collaboration between those systems, the risk of cross contamination from adversarial cyber actions is reduced. This is yet another example where engineering analysis is required to determine the appropriate trade-offs.

Diversity can be created in software using techniques such as address space layout randomization. That creates executable code located at unpredictable and ever-changing

positions in memory, negating the viability of certain types of cyber-attacks. Some compilers also support randomizing elements of processing layout. However, those techniques should be used with caution as indeterminate behavior can negatively impact airworthiness certification.

Diversity can significantly enhance system resilience, security, and robustness. Intentionally varying components and implementing diverse mechanisms in hardware, software, communication protocols, and processing methods, makes it more difficult for adversaries to exploit known weaknesses or to launch effective widespread attacks. Implementing cyber diversity within a system will require tradeoffs, but it is a useful technique to prevent and withstand cyber adversity in the operating environment.

7.5.2 Techniques for Mitigating Cyber Adversity

This section provides an overview of the architectural techniques that aircraft system designers can use to mitigate and withstand the effects of cyber adversity. It is crucial to design systems that not only prevent cyber-attacks but also mitigate the effects when the adversary inevitably experiences some degree of success. These archetypes embody the design principles that enable mission-essential functions to continue despite successful attacks. By designing aircraft systems to consider and withstand various types of cyber events, system architects can enhance the cyber resiliency posture of the aircraft.

7.5.2.1 Analytic Monitoring

Analytic monitoring consists of examining a wide range of operational properties and behaviors on an ongoing basis coordinated across system resources. This archetype is used to enable real-time detection of potentially adverse conditions, stresses, or attacks. In short, it provides the operator of the system with cyber situational awareness.

In a sense, aviation systems already have examples of the analytic monitoring archetype implemented. For example, every aircraft has gauges and sensors that provide real time operational status to the pilot. Those instruments detect and warn of events such as exceedances in operating temperature or loss of a navigational signal.

Abnormal conditions on an aircraft could be indicative of a cyber event, which is why the analytic monitoring archetype frequently appears as a cybersecurity best practice. However, most aviation systems have experienced anomalies resulting from other factors such as mechanical failure, environmental events, or human error. Consequently, implementing analytic monitoring for that domain should start with data analysis to develop a broad baseline of normal operation before flagging anomalies as potential cyber events.

From an operational standpoint, the pilots and support crew of aircraft need contextual awareness of the state of flight and mission critical systems. That is a separate archetype explored in section 7.5.2.2.

Analytic monitoring is an important mechanism to collect valuable data essential to sustain aviation platforms to an appropriate degree of cyber resiliency over the entire service life of each platform. This functionality was likely limited on some legacy platforms due to storage and latency concerns. Additionally, many aircraft were not originally designed with any requirement

for the type of data collection and analysis defined in the analytic monitoring archetype. Even in cases where analysis isn't routinely performed, there is still value in the data collection should there be an incident.

Analytic monitoring involves continuously examining a variety of aircraft properties and behaviors for real-time detection of adverse conditions, stresses, or attacks. That data can be used to provide cyber situational awareness for the pilot and support crew. This archetype is a critical mechanism for both the current and long-term cyber resiliency posture of the aircraft.

7.5.2.2 Contextual Awareness

Contextual awareness provides an understanding of the surrounding environment, situation, or context in which the aviation system is operating. In aircraft, contextual awareness refers to the pilot and support crew's understanding of the state of flight and mission-critical systems. This archetype works in concert with the analytic monitoring described in Section 7.5.2.1. It is how the analysis of data is presented to the pilot and other aircraft operators.

Knowing when cyber enabled resources are operating in a degraded mode or are otherwise disrupted or denied is critical situational awareness for the pilot. That information enables appropriate and timely actions in response to potential cyber events. The implementation will vary widely based on the system that is impacted along with any recovery mechanisms available on the platform.

It is essential that aviation systems provide real-time awareness of the status and performance of cyber-enabled resources that are flight and mission critical. That enables the pilot to make informed decisions and take appropriate action to mitigate ongoing risks or issues. By detecting and responding to degraded or disrupted cyber-enabled resources in a timely manner, the pilot can maintain safe and efficient aircraft operation, safeguard the wellbeing of passengers and crew, and preserve mission execution.

7.5.2.3 Deception

Deception is an architectural technique that constructs systems with the intent to mislead, confuse, or hide critical assets from adversarial cyber threat actors. This can cause attackers to waste time and resources on targets that have no tactical value. Deception can also increase the opportunity to detect an adversarial presence on the system which might expose their tactics, techniques, and protocols.

The use of deception has not previously been widely adopted in aircraft systems due to the potential impact on system performance and efficiency. One way this archetype can be implemented is to create decoy subsystems or data intended to attract adversarial attention. However, the aviation industry has traditionally been conservative in allowing unused or unnecessary functionality onboard aircraft. Even minor errors or failures can have catastrophic consequences.

However, there is a strong argument for at least considering the use of deceptive techniques in aircraft platforms. For example, honeypot messages or subsystems might entice an adversary

with a foothold on the platform to attempt execution or access. Unlike anomalies that can arise from normal errors during operations, that would be a clear indicator of adversarial cyber activity.

In any system, when the deception technique is used, it should not interfere with flight or mission critical operations. The presence of deceptive techniques should be implemented in such a way to be transparent to the pilot and the system operators.

The use of deception in aircraft systems can provide an additional layer of security and enhance the overall resilience of the system. Creating decoy or fake targets may mislead attackers into believing they have discovered a way to create or trigger a desirable cyber effect only to expose their position. When a deceptive element is triggered, it can also provide valuable insights into the attacker's tactics, techniques, and protocols which can be extraordinarily valuable for detecting and preventing future attacks.

7.5.3 Techniques for Recovering from Cyber Adversity

This section provides an overview of the architectural techniques that support recovery following an adverse cyber event. It is important to prioritize and implement recovery mechanisms into aviation system design from the outset, as it is late to need once cyber adversity occurs. The capability to restore aviation systems quickly and efficiently to full operational capacity is a critical design objective for ensuring mission resilience and safety. That enables aircraft operators to better respond to cyber events which minimizes the impact on operational use and mission execution. Implementation of these techniques not only restores functionality but also rebuilds confidence in continued operations.

Recovery reconstructs mission capability after an adversary has reduced or eliminated that capability through a cyberattack. In aviation systems, the effectiveness is heavily dependent on adding support for recovery capabilities into the architecture and design of the system. Additionally, system recovery also relies on awareness and preparedness in the operational environment for quick execution when the need arises.

7.5.3.1 Non-Persistence

Non-persistence is a cyber resiliency archetype that creates software and systems that initialize each session from read-only data memory. For example, aircraft software that instantiates itself from unmodifiable memory is less susceptible to recurring cyber-attacks from an adversary who successfully injects code or backdoors to attempt to create a persistent foothold. Since the run time software and data is discarded once the system is shut down, cyber adversaries who experience success only gain temporary access into a non-persistent system.

Combining non-persistence with the substantiated integrity techniques described in section 7.5.1.5 provides even better protection. Substantiated integrity ensures the software and data written to persistent memory has not been modified by a threat agent with physical access to the hardware either in the operational environment or through supply chain injection.

Non-persistence is a crucial cyber resiliency design principle for aviation systems. It also supports fast reinitialization to a known good state from software and data stored in non-volatile read-only memory.

7.5.3.2 Rapid Restart (Even in Aviation Systems)

Whenever users reach out to an Enterprise IT system help desk, they are almost always first directed to perform a restart. That is because many common issues can be resolved by clearing memory and reinitializing the software. Though aviation system architects and designers strive for better stability and reliability than that, sometimes a reset can also resolve transient issues on the platform.

Consequently, all components within an aviation system should have provisions for rapid restart both on the ground and in flight. Fortunately, the requirement for that capability is typically imposed to meet safety and airworthiness standards.

Subsystem restarts should not be overlooked as a recovery mechanism for any cyber enabled system that is exhibiting abnormal behavior. Additionally, architects and designers should consider mechanisms for restart that explicitly consider reinitialization of non-volatile data and software that may have been modified by a malicious adversary. In some circumstances it may make sense to extend the restart mechanisms to include a reduced functionality “safe mode” like some enterprise IT operating systems.

7.5.3.3 Software/Firmware Update Mechanisms

It is an inescapable reality that the software on all aviation systems will have to be updated or reloaded at some point in time. The need to upgrade an aircraft to include new capabilities is one reason why updated software is needed. Additionally, sometimes updates are required to fix errors or to mitigate emerging vulnerabilities against the platform. Consequently, it is essential to include mechanisms for software updates into the architecture and design of all aircraft components.

In addition to the high-level capabilities implemented by the software, embedded systems rely on firmware that interacts with hardware controlling low-level cyber physical systems. In general, firmware updates are usually required less frequently than for application software. However, the need for update is still a requirement and must be accounted for in the architecture of the system.

Systems architects and designers must understand that software and firmware update capabilities are a viable attack vector for malicious cyber threat agents. For example, an attacker with access to the software supply chain could deliver a software update that includes vulnerabilities or triggerable cyber effects. It is also possible for the firmware to be targeted using that same mechanism.

Consequently, some degree of protection against loading unauthorized software is strongly advised for all aviation systems. Using the substantiated integrity archetype described in 7.5.1.5 is one approach for at least partially mitigating that risk. Similarly, firmware updates should be protected by manual interlocks to protect against accidental or malicious update by a cyber threat

agent. Due to safety considerations, update of firmware and software in flight is usually not advised.

In the event of suspected or confirmed cyber-attack against a subsystem, aircraft operators need assurances that software reloads will eradicate any footholds, backdoors, or vulnerabilities the attacker might have implanted. Software loading mechanisms should be designed to provide the highest possible assurance that no remnant of cyber incursion is left resident on the system. Considering that use case once a suspected or confirmed cyber-attack has occurred is late to need.

7.5.3.4 Rollback Versions

In enterprise IT software development, a rollback is the process of reverting a system or application to a previous version. This is usually done in response to a problem or bug in a new release of the software. Rolling the system back to a known good state is more common in enterprise IT systems where the consequences of software release failures are relatively low.

The concept of rollbacks is rare for aviation systems. New software updates go through extensive flight testing prior to deployment. However, there is merit to the idea of designing and deploying systems that can work with previous versions of software should a new release prove to be unstable or vulnerable.

The ability to swap to previous software versions either in flight or on the ground is a potential response mechanism in the event that new vulnerabilities or issues emerge. However, integration of the various permutations of rollback configurations would require thorough testing and evaluation to ensure the system can seamlessly switch between software versions without encountering compatibility issues.

This added complexity in testing and integration should be carefully weighed against the potential benefits. Another alternative is to have a rollback version for all subsystems on an aviation platform that have all been verified to work together.

7.5.3.5 Establish a Backup Process

The creation of offline system backups has long been recognized as a best practice for enterprise IT systems. However, that idea has not propagated to aviation platforms. While aircraft systems do not have the same degree of transient user and business information that is important to preserve in the event of failure, there are some types of data that should be considered for backup and recovery. In addition to the operational flight programs containing executable software for the platform, some configuration files may bear calibration data that is specific to a part or a unique installation on the aircraft. Additionally, some aviation systems are loaded with maps and navigation information that is specific to the current region or theater of operation.

Having a backup process and system restoration procedures enables aircraft to be quickly wiped and restored in the case of a suspected or confirmed cyber event. That is likely to be faster than rebuilding or recalibrating the configuration files that are unique to each aircraft. In the event of

emergency, the ability to quickly restore a platform to a provably airworthy configuration could be critical.

7.5.3.6 Rapid Software Update, Test, and Deployment

There is increasing desire to develop aviation systems that support agile rapid software updates, test, and deployment capability. The increased velocity is one of the projected benefits of the move toward digital transformation in the industry. Additionally, in the case of perceived or confirmed cyber-attack or vulnerability, mechanisms for quick updates of systems are believed to be essential.

Rapid software update, test, and deployment is a genuinely good idea that is challenged and constrained by the traditionally long development and authorization timelines within the aviation industry. In fact, safety and security threats are frequently cited as both a benefit and a drawback of this increased velocity.

Architects and designers should consider the emergence of rapid software updates, test, and deployment when designing new systems. Building agility into both the development pipeline and mechanisms for rapid updates on platform can enhance the capacity for rapid response to address emerging threats and vulnerabilities. In addition to a robust and secure software development and deployment process, thorough testing mechanisms are needed for software updates to ensure they do not negatively impact system performance or the security posture.

Rapid software update, test, and deployment also creates the need for a robust change management process. Deployment of software updates can only be performed under tight configuration management and processes that ensure that only correct and authorized software is loaded onto each system.

7.5.4 Techniques for Adapting Against Future Cyber Adversity

Aviation systems must evolve and adapt as information about the highly dynamic cyber threat environment changes. That insight can come from threat intelligence or actual adversity experienced by the platform. It isn't enough to simply know and understand the types of threats facing aircraft. Those systems must be architected and designed with the knowledge that the platform will inevitably require timely update or modification in response to cyber-attacks.

Adapting aviation systems for cyber resilience is an ongoing process that requires data collection through continuous monitoring and analysis to identify when updates are needed. This full-lifecycle proactive approach requires mechanisms for updating software, changing configuration settings, and implementing new response mechanisms to prevent, mitigate, respond, and recover from future cyber-attacks. It enables the platform to adapt with agility.

This section explores the architectural techniques and design principles that support adapting aviation systems for sustaining cyber resilience.

7.5.4.1 Adaptive Response

Adaptive response is defined in NIST 800-160 Volume 2 (1) as implementing agile courses of action to manage emerging risks against a platform. That includes the ability to evolve the aircraft to respond in a timely and appropriate manner to emerging adverse conditions, stresses, and attacks. It also includes evaluation of events which may be indicators or probes that portend future incursion attempts.

Aircraft systems require a high degree of reliability and predictability, and unexpected behavior can have catastrophic consequences. Automated threat detection in enterprise IT systems is often based on large dataset analysis of network traffic and user behavior. It also leverages a rich set of threat intelligence and indicators of cyber-attacks that have occurred in the past.

In the future, automated response may be possible for aviation systems, but at this time even relatively benign self-healing actions such as automatic restarts or shutdowns should be approached with caution. This area is an example of why it is essential to carefully evaluate and adapt traditional IT security techniques to ensure the mechanisms are appropriate for safety of flight. While the exact techniques may not be appropriate or desirable for the real-time operating constraints of aircraft platforms, nevertheless the principles behind those mechanisms should still be considered.

Predefined degraded modes of operation such as those defined in Section 7.5.1.4 could potentially be selected automatically or manually in response to operating conditions which may be indicative of cyber adversity. Similarly, a pilot or other aircraft operator should have the ability to restart or shut down systems or datalinks that are problematic. Mechanisms to perform those actions should be considered as a fundamental part of the system architecture and design.

It's important not to view adaptive response solely as human reaction in an aircraft system. At the same time, it is dangerous to assume all mechanisms must be automated and self-healing. Overemphasizing human reaction may lead to neglecting necessary support mechanisms, while assuming automation may be unnecessary or undesirable in certain situations. Striking a balance is crucial for effective adaptive response.

The architects and designers of systems should leverage the results of threat modeling analysis and cyber risk assessment to identify operational responses that are appropriate for successful mitigation of specific threats and risks against the platform. Those responses should have support built into the architecture and design.

7.5.4.2 Forensics Capability

In aviation systems a “black box” is a device that records data from an aircraft's flight instruments and cockpit conversations during a flight. The purpose of the black box is to provide investigators with information about the flight's conditions, actions, and events leading up to an accident or incident.

Paradoxically, many aviation platforms do not have built in mechanisms for performing forensic data collection in the event of a suspected cyber-attack. In contrast, forensic data collection would typically include things like log files, system configuration, data files, memory dumps,

and user activity for an enterprise IT system. These same types of information would be useful for analysis of a cyber incident on an aircraft.

Without forensic data collection, the analysis required to understand exactly what transpired cannot be performed to ascertain what actually happened in the system. The absence of that information makes it difficult to understand how the aircraft should be updated or modified to prevent recurrence of similar cyber events in the future. Without that information, both the developers and operators of aircraft systems are flying blind.

Conducting forensics on aviation platforms is extremely difficult without built-in tools to collect the data to perform meaningful analysis. Consequently, architects and designers should look for opportunities to create cyber forensic data collection capabilities on the platforms they create.

7.6 Pitfalls and Fallacies of Cyber Resilient Design for Aviation Systems

Achieving and maintaining resilience against ever-evolving threats is paramount for aircraft platforms. Unfortunately, the aviation industrial base is plagued by some fallacies which can undermine efforts to create cyber resilient platforms. Flawed logic and reasoning lead to misconceptions about the nature of cyber threats and mischaracterization of cyber risks. This section describes some of the prevailing ideas that can contribute to decisions that compromise the cyber resiliency posture of the aircraft.

7.6.1 Assuming Aircraft are Protected By an Air Gap

It is such a tempting idea to assume that aircraft platforms are protected from adversarial cyber activity by a literal “air gap.” During operational use, these platforms are not tethered to a network cable that ties them to public networks. The disconnected nature of aircraft systems when compared to enterprise IT systems is frequently used as a justification for discounting the likelihood of adversarial cyber-attack against aircraft.

The reality is that even during flight, most aircraft are connected to offboard systems. That can come through voice and data communication links with other aircraft, satellites, and ground stations. Both commercial and military platforms sometimes have connections to traditional IT networks which includes the public internet and passenger entertainment systems.

The physical isolation of a cyber-enabled network was once considered to be an effective security measure to protect critical systems from cyber threats. However, examples of sophisticated cyber-attacks that traverse air gap boundaries such as the attack described in “Countdown to Zero Day” in Section 2.9 have shown that defense through air gaps is never absolute.

Aviation platforms are potentially exposed to the public internet via connections to maintenance and mission/flight planning devices. Those systems are typically implemented using enterprise IT networks with data links that traverse public or private networks. An adversary who gains access to a maintenance system may be able to use it as a foothold to pivot onto the platform itself.

The existence of a literal air gap should never be used as a justification for not implementing support for cyber resiliency on aircraft platforms. Similarly, the concept must be used with caution when used as a factor for estimating the likelihood of adversarial cyber risks to the platform.

7.6.2 Overvaluing Physical Access Controls

Time and time again, physical access security controls, such as locks and guards, have proven to be ineffective against determined adversaries. Commercial aviation systems have regulatory requirements for protection imposed by the FAA and the Transportation Security Agency (TSA). Theoretically, aircraft are stored in secure locations where people who are not authorized do not have access. Additionally, personnel who have authorized access are typically subjected to background checks and in many cases security clearances.

Physical security measures can be costly and difficult to manage, especially for aircraft that are operated and stored in the relatively large spaces of commercial airports. In many locations it may be difficult to restrict physical access or to monitor who is coming and going. Similar challenges can be experienced when recording and auditing physical access events.

Maintenance and sustainment that requires direct access to the aircraft can be a blind spot even when personnel are well vetted and trusted. Cyber-enabled maintenance equipment that can provide an access vector into systems are frequently connected to enterprise IT networks where exposure to cyber adversaries is higher. Similarly, software and configuration data frequently traverse public networks where malicious actors may have gained access. Even replacement parts and subsystems were likely to have been distributed using public transportation and warehouse systems where physical access is less controlled than on the flightline of an airport or a military installation.

It is critically important to maintain a balance between physical security and cybersecurity for aviation systems. While controlling physical access is a vital security measure for aircraft, it is never a justification for assuming that cyber risks are low for any attack that is believed to require some degree of direct access. History has shown us that it is frequently easier than we think.

7.6.3 “Insider” Attacks Against Aviation Systems

Historically, cyber defenses for aviation platforms and weapon systems have been primarily focused on keeping attackers out. Threats from insiders with legitimate access to the aircraft throughout both the development and operational lifecycles have been historically neglected. Insider cyber-attacks pose a significant threat to aircraft. These attacks can be difficult to detect and prevent and can cause catastrophic damage to the platform.

Historically, the aviation industrial base has operated under a cultural assumption that aircraft internal subsystems and busses can be trusted to not perform malicious actions. While it is true that most aircraft were developed with a significant emphasis on safety and reliability, it is incorrect to suggest that the potential for internal threats against the platform do not exist.

Therefore, a blanket assumption that internal subsystems can be trusted ignores the risk that a cyber adversary may gain a foothold on one subsystem and use that point of access to create cyber adversity against another. Additionally, the increasing adoption of plug and play open system architectures means that an attacker may also gain platform access through components that are loaded to the aircraft.

Consequently, it is essential for aviation system architectures to be developed with an attitude of vigilance and skepticism, rather than blind trust of internal subsystems. Onboard subsystems must protect themselves against both people and other aircraft components doing malicious or unexpected things.

The aircraft must maintain safety of flight and mission performance despite those challenges.

7.6.4 The Paradox of “Hardened” Systems

In cybersecurity, the term "hardened" is typically used to describe a system or application that has been fortified or made more resistant to attack. While hardening is a widely accepted best-practice in cybersecurity, a decision was made during the editorial process that this FSAD guidebook would intentionally avoid using the term. The security controls, actions, and mitigations typically associated with hardening systems is fundamentally reactive. Hardening security measures are made in response to known threats or vulnerabilities rather than the proactive approach advocated for by this resource.

Waiting until after detection is not an effective strategy for engineering cyber resilient systems, which is the objective of this FSAD guidebook. Aircraft systems must be proactively built to prevent cyber-attacks long before they occur and even before the nature of new and emerging threats and vulnerabilities are known.

However, hardening aviation systems against cyber-attacks is still an important part of a comprehensive cybersecurity strategy for any platform. During architecture and design, aircraft must be constructed to support ongoing security sustainment. When an aircraft is deployed for operational use, that security sustainment must be diligently performed. At a minimum, that includes deployment of software patches and updates to configuration settings.

An aviation system that is only hardened retroactively against known threats and vulnerabilities is most certainly better than one that hasn't taken those steps. At the same time, a platform that has only been hardened will almost always exhibit less cyber resiliency than an aircraft that was architected and designed specifically for that intrinsic characteristic.

Legacy aviation systems that were developed and deployed before the current understanding of cyber resiliency engineering should be hardened to the best extent possible. Sadly, many of those platforms were not designed to support those kinds of updates. Any development performed to upgrade such aircraft should look for opportunities to add those capabilities.

Hardening systems remains an important part of a comprehensive cybersecurity strategy. However, it should be regarded as one of many techniques that should be used in concert to protect aviation platforms. Hardening alone will not create cyber resilient systems.

7.6.5 Inherent Flaws/Zero Days

The best way to deliver intrinsically secure aviation systems is to avoid inherent vulnerabilities or flaws in the first place. A “Zero Day” is a latent vulnerability that has yet to be detected. Sometimes the absence of publicly disclosed Zero Day vulnerabilities for aviation systems is falsely interpreted to mean that aircraft platforms are already more secure than other cyber-enabled systems. That can sometimes lead to the idea that embracing the techniques and best practices for engineering cyber resilient systems isn’t necessary. That isn’t an acceptable attitude.

A Zero Day that has not been discovered or disclosed does not mean that it does not exist. In fact, many latent vulnerabilities go undetected for years, and some may never be discovered at all. However, aviation systems are targeted by well-resourced and highly capable nation-state adversaries that are dedicated to searching for and exploiting previously unreported vulnerabilities in those systems. As a result, it is an imperative for acquisition agencies and development organizations to insist on and implement the techniques and best practices that eliminate those latent vulnerabilities to the greatest extent.

At the same time, while the information presented in this chapter is focused on using best practices to minimize the presence, scope, and impact of any flaws, it is important to understand that perfection is unlikely. Consequently, the system must be architected with mechanisms that support timely application of security updates and configuration settings required to keep the aircraft secure.

Sometimes developers are tempted to overlook the probable existence of inherent flaws and zero days in software that is custom developed for purpose specific systems such as aircraft. Intrinsic cyber resiliency comes from exercising caution when assigning trust. An absence of publicly disclosed vulnerabilities is never definitive evidence that inherent flaws do not exist.

7.6.6 Security by Obscurity Doesn’t Work

Security by obscurity is a fundamentally flawed idea that keeping the design and implementation of a system or application out of the public eye makes it more difficult to attack. The thought that underpins that erroneous belief is that without information and insight about targeted systems, cyber threat actors will be unable to find and exploit vulnerabilities. That bad logic can contribute to resistance to applying the techniques and best practices described in this chapter.

It is never valid to assume that an absence of implementation details about a system will prevent adversarial cyber events. Nation-state adversaries are known to have enjoyed access to business networks of suppliers who are a part of both the commercial and defense industrial base. Consequently, it is not a safe assumption that information that may have been gleaned from access to the supply chain is not available to cyber adversaries. Additionally, any hardware or software that has been obtained by malicious agents can be reverse engineered which could potentially expose system vulnerabilities.

When designing aviation systems, architects and designers should assume that the adversary will one day have access to all the development artifacts and data ever produced about it. In fact, our enemies may come to understand the platform better than the original developers or operators.

Aircraft must be designed to perform safely, securely, and exhibit the characteristics of intrinsic cyber resiliency. It is a key part of developing systems that are resistant to cyber adversity that can only come from detailed insight into the implementation.

7.7 Advanced Topics in Aviation Systems Cyber Resiliency

The field of aviation systems cyber resiliency is highly dynamic, and there are several techniques that should be considered when addressing the challenges and risks when safety and performance are also key considerations. This section offers advanced insights and guidance for those looking to optimize cyber resiliency for aircraft.

7.7.1 Exercising Operational Security for Aviation Systems

The concept of "stealth" or "low observability" for military platforms refers to a set of design features and technologies to minimize aircraft detectability by radar, infrared sensors, and even visual observation. The goal is to reduce the aircraft's detection range and enhance survivability and effectiveness in combat environments. In a kinetic contested environment, the adversary cannot shoot down what they don't see. In a cyber contested environment, hostile actors cannot exploit aspects of the platform unless they have some degree of visibility.

The abstract concept of cyber "low observability" is an important design element for aviation platforms. Publicly disclosed or leaked design and implementation characteristics of an aircraft can be leveraged by determined adversaries to craft future effective attacks. Consequently, architects and designers must take care to protect critical information that could provide cyber adversaries with an advantage.

Design specification and implementation details must be protected at all times. That starts during system development. Chapter 4 previously discussed the importance of securing the development and production infrastructure against malicious cyber activity. That includes preventing design and implementation details from exfiltration to unauthorized sources.

Additionally, the platform should be architected and designed to not leak information about the underlying characteristics when in operational use. That includes techniques such as encryption, data protection, and even considering how the system responds to unexpected conditions and messages. Exercising the concepts of cyber operational stealth is a key design philosophy to remember when designing aviation systems.

7.7.2 War Reserve and Combat Mode

In military applications, "War Reserve" is used to describe a stockpile of equipment, supplies, or ammunition that is set aside for use in wartime or other emergency situations. The practice ensures sufficient resources to support combat operations and other mission-critical activities during times of conflict or crisis.

For military equipment, such as aircraft systems, the conceptual equivalent of war reserve is known as "Combat Mode." That refers to specific operational capabilities that are only used for

actual warfighting situations. When combat mode is enabled, the equipment is typically configured for maximum performance and survivability.

For example, military aircraft may have a combat mode that enables the use of advanced weapons systems and sensors, as well as enhanced navigation and communication capabilities. The specific features and capabilities of combat mode varies from platform to platform.

The strategy behind combat mode is preventing adversaries and combatants from observing systems in routine operational use and leveraging that insight for developing countermeasures against the capabilities. From a cyber perspective, combat mode may also prevent an adversary from developing and testing attacks against aircraft functionality that is not routinely used.

However, most of the exposed attack surface on any given aviation system is unlikely to be altered when combat mode is enabled. Consequently, if an attacker has found a general method to disrupt performance on the platform, invoking the increased capabilities of combat mode may not provide additional protection.

Some strategists have suggested that multiple versions of software could be used as a virtualized war reserve mode. For example, using alternate layouts of executable code in memory may be a useful technique to mitigate attacks that rely on knowing the location. Unfortunately, if that approach was selected it would add considerable cost and testing time to the overall development schedule. A careful analysis of the trade-offs between the cost impacts and risk reduction should be carefully considered if this approach is selected.

When "War Reserve" or "Combat Mode" is considered for military aircraft as a mitigation against cyber-attack, it is important to consider the potential limitations and costs. A thorough analysis of these trade-offs is necessary to determine the most effective approach for mitigating risks while maximizing the probability of mission success.

7.7.3 Modular Open Systems Architecture (MOSA)

The DoD has recently started in initiative to implement a Modular Open Systems Approach (MOSA) as an acquisition strategy for designing more affordable and adaptable systems. (45) MOSA can be defined as a technical and business strategy for designing an affordable and adaptable system. The approach is required by United States law for all major defense acquisition programs. (46)

MOSA compliant systems employ a modular design for interfaces between systems and components. It is intended to ensure that modular system interfaces comply with consensus-based standards or other mechanisms that support "plug and play" operation. The objective of the MOSA initiative is to acquire systems that use architectures that allow major system components to be independently added, removed, or replaced throughout the life cycle of the platform. It is intended to be a driver both of enhanced competition and innovation.

The DoD is anticipating significant cost saving, more agile deployment of new capabilities, schedule reduction, and increased interoperability as a result of the MOSA initiative. Those benefits are projected to benefit acquisition, development, and sustainment of systems.

The use of MOSA within aviation platforms has an impact on the cyber resiliency posture. Open architectures means that potential attackers have greater visibility into the system interfaces and inner workings of each subsystem. That can make it easier to identify vulnerabilities and develop exploits.

Additionally, the “plug and play” approach lacks centralized control over software and hardware components. It can make use of techniques such as substantiated integrity more challenging and increase the need for zero trust architectures. Decentralization can also lead to inconsistencies in security practices and configurations between system elements.

Aviation system architects and designers of MOSA compliant systems must understand that the approach increases the attack surface of the overall system and system components. That increased risk should focus prioritization of attention and resources toward the architectural techniques that mitigate that increased exposure.

7.7.4 Vulnerability Management for Aviation Systems

Sustainment is an ongoing activity for aviation systems due to environmental and operational stressors that lead to wear and tear, degradation, and failure over time. Additionally, aviation systems are complex and constantly evolving, with new technologies, regulations, and mission requirements emerging on a regular basis. Consequently, there is a general understanding that long term sustainment of aircraft is required to maintain operational readiness.

Cybersecurity is a critical aspect of aircraft sustainment. That typically involves proactively identifying and addressing vulnerabilities in aircraft systems to determine if additional security measures or updates are needed. For enterprise IT systems, cybersecurity sustainment usually includes scanning systems with commercially available automated tools that search for known vulnerabilities and configuration errors. That approach is problematic for aviation systems because there are no commercial tools to perform that automated scanning. Even if there were, the body of knowledge of vulnerabilities specific to aviation cyber physical systems is well below the threshold of statistical significance.

That does not negate the need for a robust ongoing vulnerability management program for aircraft. However, it requires some degree of human capital to implement it in the current absence of automated tools. One way to approach this challenge is to charter an ongoing engineering sustainment effort to proactively monitor emerging vulnerabilities across the entire cyber domain for patterns or instances that may impact the aviation system.

Neglecting cybersecurity sustainment can leave aircraft vulnerable to attacks that can result in significant downtime, financial losses, and damage to the reputation of both the manufacturer and operator. Therefore, it is essential that cybersecurity sustainment that includes a vulnerability management program be identified as a top priority as systems are architected and maintained. Passively waiting for a cyber-attack to occur is not an effective method to stay ahead of emerging threats to the platforms.

7.7.5 Aviation Cyber Kill Chain

In 2011, Lockheed Martin published a landmark whitepaper that took the established concept of a “kill chain” and applied an adaptation of those principles to the defense of enterprise IT systems. (47) That paper outlined an intelligence driven approach in which indicators of adversarial activity are used to identify and contain intrusions before the cyber adversary can accomplish their objectives.

The concept that underpins a “Cyber Kill Chain” is based on the idea that there are typically multiple opportunities to detect and stop a cyber-attack. If that happens and the campaign is thwarted, it is a success for system defenders even though the attackers experienced some degree of success before being caught. The Cyber Kill Chain concept popularized a useful framework for a layered defense in depth approach that maximizes the opportunity to defeat attacks.

While that original paper was focused primarily on enterprise IT systems, the Cyber Kill Chain can be conceptually adapted for use in an aviation context. One potential mapping of the stages is presented in Table 4.

2011 Cyber Kill Chain (Enterprise IT)	Aviation Systems Cyber Kill Chain
<u>Reconnaissance</u> Research, identification, and selection of targets, often represented as crawling Internet websites such as conference proceedings and mailing lists for email addresses, social relationships, or information on specific technologies.	<u>Reconnaissance</u> Target selection of the platform based on nation state strategic and tactical objectives. Enterprise IT systems likely targeted as a part of reconnaissance for weapons systems. Initial research tactics often not materially different as for enterprise IT systems.
<u>Weaponization</u> Coupling a remote access trojan with an exploit into a deliverable payload, typically by means of a weaponized automated tool. Increasingly, client application data files such as Adobe Portable Document Format (PDF) or Microsoft Office documents serve as the payload.	<u>Cyber Effect Conceptualization</u> Adversary conceives attack vectors based on the technical contents and implementation of the target system. Attack vectors are framed in terms of desired cyber effects asserted in the target system. For example, a desired cyber effect might be to intercept a communications link or deny usage of navigation systems.
<u>Delivery</u> Transmission of the weapon to the targeted environment. The 2011 Cyber Kill Chain paper identified the three most prevalent delivery vectors of weaponized payloads by APT actors as email attachments, websites, and USB removable media.	<u>Injection</u> Injecting a weaponized payload into a target system will be largely dependent on the aircraft implementation. This can come from supply chain injection, through maintenance support equipment, or insider access. If a latent flaw is detected, adversarial injection may not be necessary.

2011 Cyber Kill Chain (Enterprise IT)	Aviation Systems Cyber Kill Chain
<p><u>Exploitation</u></p> <p>After the weapon is delivered to the victim host, exploitation triggers the intruders' code. Most often, exploitation targets an application or operating system vulnerability, but it could also more simply exploit the users themselves or leverage an operating system feature that auto-executes code.</p>	<p><u>Activation</u></p> <p>Once the weaponized payload is resident on the system, in some instances it must be activated to make it available to subsequent triggering. If the payload is a part of executable code for the operational flight program, activation may not be necessary.</p>
<p><u>Installation</u></p> <p>Installation of a remote access trojan or backdoor on the victim system allows the adversary to maintain persistence inside the environment.</p>	<p><u>Pivot/Persistence</u></p> <p>Payload software may execute to propagate and pivot through the system and/or to create persistence.</p>
<p><u>Command and Control</u></p> <p>Typically, compromised hosts must beacon outbound to an Internet controller server to establish a Command and Control (C2) channel. Once the C2 channel is established, intruders have "hands on the keyboard" access inside the target environment.</p>	<p><u>Triggering</u></p> <p>Within an aircraft system, triggering is one of the most critical and challenging aspects of cyber exploitation. The adversary needs a reliable way to trigger the weaponized payload or inherent vulnerability to produce the desired cyber effect at exactly the point in time when it is desired or needed. For example, it is most beneficial to trigger during combat rather than during training or testing.</p>
<p><u>Action on Objectives</u></p> <p>Only now, after progressing through the other phases, can intruders take actions to achieve their original objectives. Typically, this objective is data exfiltration which involves collecting, encrypting, and extracting information from the victim environment. Violations of data integrity or availability are potential objectives as well. Alternatively, the intruders may only desire access to the initial victim box for use as a hop point to compromise additional systems and move laterally inside the network.</p>	<p><u>Cyber Effect</u></p> <p>The range of adversarial cyber effects will largely be influenced by the access vectors that the adversary can ultimately identify and exploit as well as trigger in some reliable manner. The effects can include loss of life, loss of aircraft, denial of mission critical functions, and annoyances. This will vary widely based on what is possible on the aviation system as well as the objectives of the cyber adversary.</p>

Table 4. Mapping the Cyber Kill Chain for Aviation Systems

Adversarial access to aviation platforms is most likely through supply chain injection, removable flight/mission planning media, and maintenance support equipment used to service the platform. This is fundamentally different from how cyber threat agents typically access traditional enterprise IT systems.

Table 4 also introduces the concept of triggering latent vulnerabilities or implanted exploits. In an enterprise IT system, triggering an exploit is usually a function of a Command and Control

(C2) channel. If C2 is implemented in an aviation system, it is likely to be via extremely low latency methods. Consequently, it is postulated that the most likely triggering is more likely to be some external stimulus to the system that may or may not include a feedback mechanism to the attacker. Alternatively, potential triggering could also be performed on the basis of location or time. Viability of triggering approaches will vary significantly from platform to platform.

7.7.6 Zero Trust Concepts for Aviation Systems

In 2020, NIST issued a special publication that described a new model for cybersecurity known as “Zero Trust” (ZT). (48) That concept was developed specifically for traditional enterprise IT networks and was created to address the problem that perimeter-based network security had been disproved as fully sufficient for defending unclassified business systems.

The security model of ZT assumes that an attacker might have a foothold in any given system. Furthermore, ZT asserts that no internal or external entity can be assumed to be trustworthy. It creates a new paradigm that places a requirement on systems to continuously analyze trust and risk to mission execution.

ZT is not a single architecture but rather a set of guiding principles for system design and operational use. While the concept was developed for enterprise IT systems, the following abstractions of the core ZT principles can be applied to any context, including aviation systems. (49)

1. **Universal Authentication.** All devices, components, users, and communication that interacts with and within the system are authenticated.
2. **Access Segmentation.** Access to any system resource is segmented into the smallest piece possible.
3. **Minimal Trust Authorization.** Access to the system by devices, components, users, and communication operates with the least privilege needed to perform its part of the mission.
4. **Encryption Everywhere.** All communication or data storage is assumed to be monitored by an untrusted adversary.
5. **Continuous Monitoring and Adjustment.** All transactions with the system should be logged and monitored for signs of compromise.

ZT architectures can be challenging to implement in embedded aircraft systems. Timely processing of data is critical for both safety and performance of flight platforms. Introducing ZT structural concepts such as encryption or authentication, may introduce latency or other delays that could negatively impact the responsiveness of the system.

Additionally, aircraft typically have constrained resources in terms of processing power, memory, and storage. A mix of legacy and modern technologies are often in use on aviation platforms. To compound matters, legacy systems may not have been designed with security in mind, making them difficult to integrate into a ZT architecture without significant modifications or updates. Consequently, implementing the principles of ZT and still ensuring system performance demanded for safety of flight can be challenging.

Nevertheless, the less an aircraft relies on trust from both internal and external systems, the less opportunity there is for compromise through violating that trust. The interest in application of ZT concepts to aviation systems is high. Implementing the core concepts of ZT to the greatest extent possible should be pursued during the development of aircraft systems.

7.7.7 Virtualization

Virtualization is a technology that allows multiple operating systems and applications to run on the same physical processing system at the same time. It does this by abstracting the hardware resources of a computer, such as the CPU, memory, and storage, and presenting them as virtual resources to the operating systems and applications. This allows multiple virtual machines (VMs) to be created on a single physical machine, each with its own operating system and applications.

Aviation systems can benefit from the use of containerized virtualization for several reasons. By isolating the file system, network stack, and other resources, it is possible to establish and enforce data separation for systems that handle information at varying levels of sensitivity. Containerized sandboxes can also serve as an isolation mechanism for segmented processing, which helps protect against malicious cyber activity spreading from one compromised application to another. Using a network with properly configured virtual local area networks (VLANs) or subnets, along with hardware-enforced virtualization, can create a system with multiple security layers. That approach limits the potential for threats to proliferate.

Virtualization is also an attractive approach for implementing system diversity as described in Section 7.5.1.8. It is an economic way to employ diverse hardware and software components. That makes it more difficult for adversaries to exploit common weaknesses or launch widespread attacks, thereby bolstering the system's overall cyber resiliency posture. Virtualization also supports air system adaptability architecture and design objectives.

Due to size weight and power restrictions, physical isolation is frequently not practical or feasible on an aviation system. Consequently, virtualization is an attractive alternative that has been gaining acceptance for flight platforms. The use of virtualization in aviation systems is expected to continue to grow as organizations seek to take advantage of its many benefits.

7.7.8 Architecting and Designing for Worst-Case Survivability

Aviation system architects and designers play a crucial role in ensuring the safety and reliability of aircraft in the face of worst-case scenarios. This requires a holistic and proactive approach to the cyber threats that might be catastrophic for the platform. Rather than believing that disastrous events will never happen, aircraft designers must consider mechanisms that might be essential if it does.

Andy Greenberg shares a story in his book “Sandworm: A New Era of Cyberwar and the Hunt for the Kremlin’s Most Dangerous Hackers” that serves as a cautionary story for architects and designers of modern highly cyber enabled systems. (16) In that case the system under attack was

a power grid rather than an aviation system, but one stunning observation is highly relevant to aircraft.

Modern power grid management systems have become so dependent on cyber capabilities that many are being manufactured without interface mechanisms for manual operation. For example, if an attacker opened a circuit breaker that should be closed, there is no way for the operators to manually flip a switch back to the correct position.

For aviation systems, mechanisms and methods for restarting or even shutting down subsystems in the case of a suspected cyber-attack should be available to the pilot and the support crew. In modern platforms, this is increasingly done through cyber-enabled flight controls. In many instances it may make sense for some components to be connected to a physical mechanism to shut down processors or to activate degraded modes of operation.

As aircraft systems incorporate more and more powerful cyber-enabled technology, it is important for architects and designers to remember this also introduces new risks and vulnerabilities. The superior capabilities enabled by cyber also creates a dependence. In the event of a cyber-attack or other disruption, this could result in catastrophic loss of all capabilities.

By carefully considering the risks and vulnerabilities associated with cyber-enabled technology, and by planning for worst-case scenarios, designers of aircraft systems can help ensure the safety and reliability of the platform in the face of potential threats. It is essential to remember that relying solely on cyber capabilities can leave a system vulnerable, and that designers must consider and plan for the possibility of a cyber-denied environment.

7.8 Summary and Additional Resources

Achieving cyber resilience in aviation systems is a critical goal that requires intentional efforts in architecture and design. Given the potential for catastrophic consequences from cyber incidents, it is essential for aircraft manufacturers to design systems that can successfully operate in the face of cyber threats.

This chapter presented knowledge and techniques for integrating and prioritizing cyber resilience in the architecture and design of aviation systems. By using methods to deliver functions and capabilities that are intrinsically cyber resilient, and minimizing inherent vulnerabilities introduced by the architecture and design, manufacturers can create systems that are better able to withstand and recover from cyber threats.

The following additional reading is recommended for people who wish to go into greater depth on cyber resiliency during architecture and design.

- **Systems Engineering Handbook: A Guide for System Life Cycle Processes and Activities, International Council on Systems Engineering (INCOSE), INCOSE-TP-2003-002-05-2023, Fifth Edition.**

This book describes the key process activities performed by systems engineers and other engineering professionals throughout the system lifecycle. It is the authoritative source of

a wide range of fundamental system concepts that broaden the thinking of the systems engineering practitioner, such as system thinking, system science, life cycle management, specialty engineering, and system of systems. It covers all the major lifecycle models including waterfall, V model, agile, and iterative methods. This book is ideal for anyone who has an interest or need to apply systems engineering practices.

- **Cyber Survivability Endorsement (CSE) Implementation Guide, Joint Staff J6, Deputy Director for Information Warfare, Requirements Division, Version 3.0, July 2022.**

The Cyber Survivability Endorsement (CSE) is a System Survivability Key Performance Parameter (SS KPP) of the Joint Capabilities Integration and Development System (JCIDS) Manual. CSE targets the predictable failure of the cybersecurity processes to build-in sufficiently robust cyber capabilities to prevent (resist/anticipate), mitigate (absorb/withstand), recover from, and adapt to the full spectrum of mission assurance cyber-events in plain language requirements for program management.

- **The NIST Cybersecurity Framework (CSF), The National Institute of Standards and Technology, Version 2.0, February 26, 2024.**

The NIST Cybersecurity Framework (CSF) 2.0 is a comprehensive guide designed to help organizations manage and reduce cybersecurity risks. It provides a flexible taxonomy of high-level cybersecurity outcomes suitable for organizations of any size, sector, or maturity level. This flexibility is particularly beneficial for those architecting and designing intrinsically cyber-resilient systems, as the framework does not prescribe specific solutions but instead focuses on desired outcomes and links to a wide array of resources and best practices.

Chapter 8

Threat Modeling and Assessing Cyber Resiliency

This chapter introduces how to use threat modeling and similar techniques for assessing cyber resiliency. Threat modeling is a proven and effective technique for identifying, quantifying, and addressing potential negative impacts on systems. It is a proactive approach to security that helps organizations anticipate and mitigate cyber adversity before it occurs.

By identifying threats and vulnerabilities early in the development process, aviation system manufacturers can implement security measures to reduce the risk of a successful attack. Threat modeling is a continuous process that should be revisited and updated regularly to ensure the system remains resilient as new cyber threats emerge.

This FSAD Guidebook has a recurring theme that the best way to develop intrinsically cyber resilient systems is to integrate cyber resiliency best practices throughout the development lifecycle. Threat modeling is yet another vital aspect of a comprehensive strategy to do just that.

8.1 Threat Modeling

At the fundamental level, threat modeling is a technique of abstraction and decomposition of a system to find cybersecurity issues. As a general term, modeling is a representation of a system created to assess and answer specific questions about it. It is a proven and effective method for assessing the quality and fidelity of the development artifacts. Modeling is a powerful tool for identifying opportunities to build better systems.

There is a general initiative in the aviation industry to move toward Model Based Systems Engineering (MBSE) techniques for new development programs. The structured and visual approach is well suited to managing the ever-increasing complexity of modern aviation systems. It facilitates collaboration and can lead to a more efficient, effective, and collaborative development engineering process.

While threat modeling has been used long before the current push toward MBSE approaches emerged, there is tremendous synergy between the two techniques. Threat modeling is an effective mechanism to ensure that security considerations are integrated throughout the system development lifecycle. The proactive identification of potential threats and vulnerabilities, leads to more cyber resilient system designs and improved risk identification and management.

Whether threat modeling is integrated with MBSE development or initiated as a standalone analysis activity, it must address four key questions (50):

1. What is being built?
2. What can go wrong from a cyber resiliency/security perspective?

3. What should be done about the things that can go wrong?
4. Is the analysis thorough and complete?

Additionally, threat modeling does not have to be confined to the development lifecycle. It is equally applicable for use to refine operational defense in deployed systems. Continuous performance of threat modeling into both the engineering and operational lifecycles enables proactive identification and mitigation of potential adversity. By embedding threat modeling into both development and operational phases of the system lifecycle, organizations can better manage cyber risks and enhance the defensive posture of the system against evolving threats.

(51) That perspective is particularly important given the long service lifespans of aviation systems.

Threat modeling is a proven technique for the discovery and elimination of vulnerabilities and flaws inherent to the architecture and design of systems. It is a vital best practice of cyber resiliency engineering.

8.2 Threat Modeling in the Systems Engineering Lifecycle

Threat modeling should be used in the earliest phases of aviation system architecture and design as an effective means to identify and mitigate cybersecurity risks. That maximizes the opportunity for systems architects and designers to create systems that are intrinsically cyber resilient. The proactive identification of security concerns is necessary for ensuring that cyber resiliency is considered from the onset of the design process. The analysis produced by threat modeling can directly contribute to more secure and resilient aviation systems.

A threat is anything that can cause a cyber effect. Threats can arise through malicious exploit of a vulnerability, failures, or accidents. A comprehensive threat modeling initiative will consider the risk of potential damage or loss from all those sources.

At a minimum, it is recommended that a threat model be created, analyzed, and reviewed as exit criteria for architecture validation and design completion. Threat modeling can also add value at later phases in the development lifecycle as a mechanism to create test plans and to assess the impact of issues that emerge during cybersecurity testing. However, it must be emphasized that it is generally more expensive to make architectural and design changes later in the system development lifecycle. Consequently, threat modeling should be performed as early as possible in addition to continuously throughout the full lifecycle of a system.

8.3 Trust Boundaries and Attack Surface In Threat Modeling

Threat modeling effectively identifies the system architecture's "trust boundaries" and "attack surface," making it a powerful tool for understanding and mitigating security risks. This identification is a compelling reason to perform threat modeling early in the development process. The insight is a necessary perspective for understanding how to best build aviation systems that are intrinsically cyber resilient.

A trust boundary refers to places where threats can traverse an interface between systems or subsystems that have varying degrees of assurance. For example, an external message arriving to

an aircraft comes from outside the boundary of trust for that platform. Trust boundaries can also be present within the aircraft between subsystems.

Threat Modeling is also an effective means for identifying and analyzing the “attack surface” of a system. Attack surface is the sum of all the points where a cyber adversary can attempt to enter, exploit, or damage it. On an aircraft, it includes things like communications interfaces, onboard systems, network connections, and software that could be exploited by malicious actors. A large attack surface increases the likelihood of a successful attack. Sometimes that size is unavoidable due to the complexity of cyber aviation systems. However, identifying the attack surface inherent to the architecture and design of the system is a key first step for proactively minimizing it.

Understanding how a cyber adversary can access and influence a system is a critical perspective necessary to identify potential risks and vulnerabilities. That insight is essential to implement appropriate security measures. Identifying trust boundaries and attack surface through threat modeling is a key viewpoint when developing and assessing cyber resilience.

8.4 A Threat Relationship Model for Assets

At a fundamental level, cyber resiliency threats against a system are expressed against assets. In an aviation system that includes things like flight control, communication networks, and navigation. Engineering for intrinsic cyber resiliency requires an understanding of which assets are vital to the safe, secure, and efficient operation of the aircraft.

Consequently, it makes sense to take an asset-based approach when threat modeling to ensure that significant risks to the most important parts of the aircraft are addressed. That approach improves the overall security posture by identifying where allocation of resources and targeted measures are most needed. Figure 11 presents an abstraction of a relationship model that is a good starting point for taking an asset-based approach.

This abstraction describes the relationship between threats, assets, and controls and highlights how threat actors or other cyber events interact with the system. Threats are asserted against attack vectors and vulnerabilities that impact a system component. Security controls or architectural mitigations are identified and applied to counter or mitigate negative consequences to the assets. This relationship model abstraction supports the necessity of aligning cyber resiliency mechanisms and controls to neutralize threats against system assets.

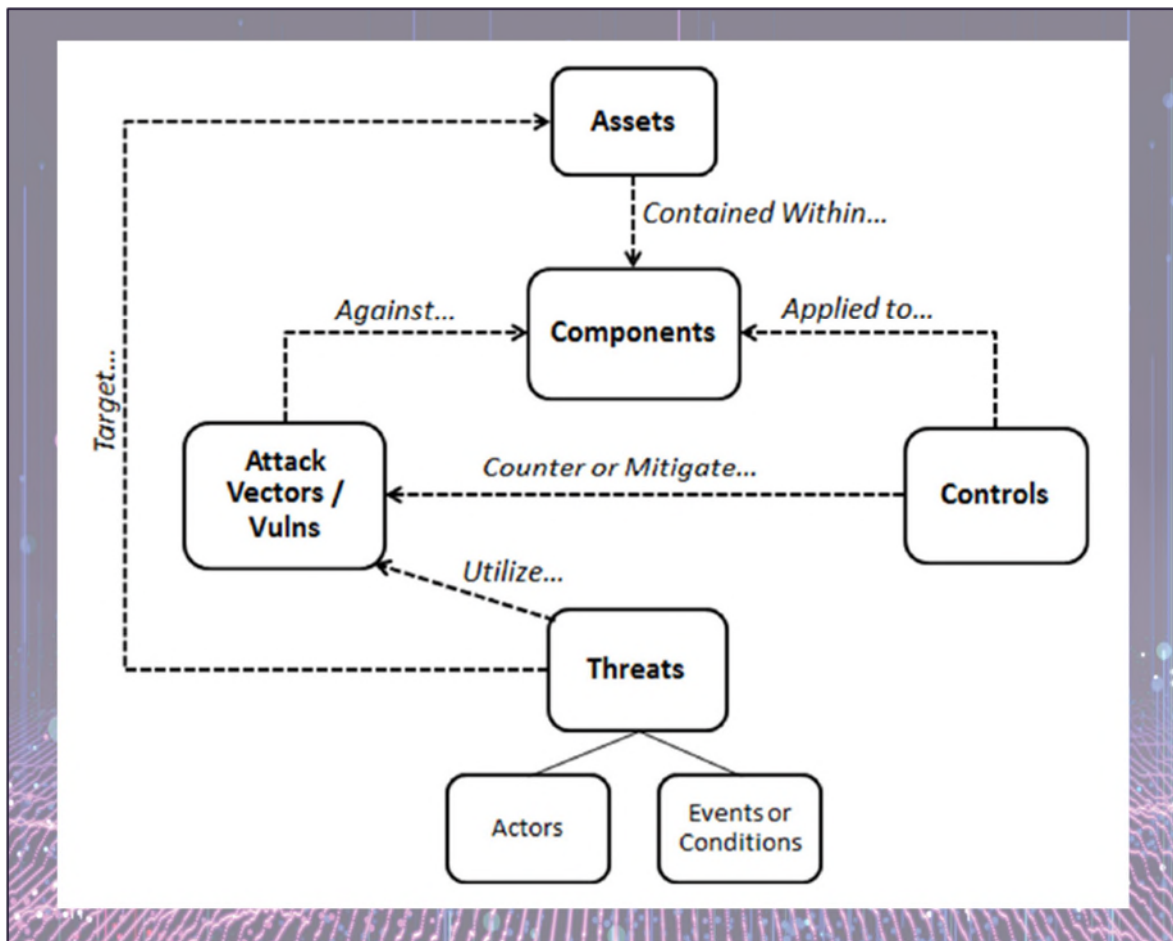


Figure 11. Threat Modeling Relationships (51)

The sequence of a threat modeling effort will typically align with the following activities:

1. Identify the assets such as data, data flows, and processing elements that the mission essential functions depend on.
2. Define the trust boundaries in the system.
3. Analyze the attack surface of the trust boundaries.
4. Identify attack vectors.
5. Analyze the impact of successful attacks.
6. Prioritize the threats based on likelihood and impact.
7. Identify mitigations that would prevent successful attacks or reduce the impact.

Ideally, this process iterates until the residual risk is assessed to have reached an acceptable level.

8.5 Limitations of Threat Modeling

Threat modeling is not a substitute for architecting for cyber resiliency. However, it can be used iteratively to assess and improve the overall system architecture and design. That conceptual understanding is critical to successful threat modeling initiatives.

Threat modeling is a very effective technique for identifying and mitigating vulnerabilities and flaws in systems. In fact, cyber resiliency hinges on effective elimination of any vulnerabilities and flaws detected. However, if threat modeling is the only technique that is used for cyber resiliency analysis during architecture and design, other functions and capabilities needed to maximize the system's support of the mission objectives may escape identification.

It is possible to neglect the mission essential functions when modeling threats. Additionally, threat modeling analysis of an architecture that does not have functions or capabilities that exist only to support the cyber resiliency of mission essential functions may fail to identify that absence. Consequently, it is important to consider techniques described in Chapter 7 in addition to threat modeling efforts.

In other words, threat modeling and architecting for cyber resiliency are both required. It is not an "either/or" situation. Effective cybersecurity in aviation systems necessitates a holistic approach that incorporates both comprehensive threat modeling and robust system architecture. Threat modeling allows for the identification and analysis of potential threats and vulnerabilities specific to critical assets within the system. Concurrently, architecting for cyber resiliency ensures that the system is designed to withstand and quickly recover from cyber adversity.

This dual approach ensures that all potential threats are anticipated and mitigated through strategic design and planning, thereby enhancing the system's overall security posture. By integrating both threat modeling and cyber-resilient architecture, manufacturers can create more robust, reliable, and secure aviation systems capable of defending against and adapting to an ever-evolving threat landscape.

8.6 Threat Identification

An effective threat modeling initiative requires identification of the threats to be considered. This section provides a high-level description of methods that are effective for eliciting threats against a system.

8.6.1 STRIDE

STRIDE is a model that was originally created to support identification of computer security threats by Loren Kohnfelder and Praerit Garg of Microsoft. (52) It is a useful aid to ensure a comprehensive answer to the question of what can go wrong in a cyber-enabled system. While this approach was originally developed for enterprise IT systems, it has proven to be useful in embedded and cyber physical domains. In other words, it is a valuable approach to take for identifying threats against aviation systems.

STRIDE is actually a mnemonic where each letter represents a category of threats against a system.

- **Spoofing.** Impersonating a person or communication mechanisms.
- **Tampering.** Something being modified that is not supposed to be modified.
- **Repudiation.** Denying that a message was sent or that an action was taken.
- **Information Disclosure.** Exposing information to people or systems who are not supposed to see it.
- **Denial of Service.** Preventing a system from providing a service.
- **Elevation of Privilege.** When a person or a program can do something that is not authorized.

The STRIDE model is effective for the unique processing requirements and potential threats specific to aviation systems. For example, spoofing could occur if an attacker impersonated legitimate components onboard the aircraft. That could include things like the aircraft's sensors providing false information to the flight control system. Additionally, data traversing communication interfaces and navigation systems spoofed to provide false information to the aircraft could have catastrophic consequences.

Tampering might involve unauthorized modification of the aircraft's hardware, software, or data. This can include altering the flight control software, modifying sensor data, or interfering with communication protocols.

Within an aviation system repudiation occurs when a person or an entity can effectively deny having performed an action. For example, changes made to flight control software should be coming from a reputable source with evidence supporting that provenance. For aviation systems that operate on data communication channels, it is important to assess the validity of the source of the information as well as the contents.

Aviation systems should eliminate unnecessary information disclosure that exposes sensitive information to unauthorized entities. This could include unauthorized access to flight plans, mission data, or real-time flight telemetry.

Denial of Service threats against aircraft systems can impact the availability of critical functions and capabilities. Example threats might include overloading the communication systems, jamming navigation signals, or disabling critical flight systems.

Elevation of privilege involves gaining unauthorized access to higher-level functions or data within the aircraft system. For example, it could involve a passenger gaining access to flight controls. It could also include a maintainer interacting with functions and capabilities typically reserved for the pilots.

STRIDE can be used to identify the specific threats and vulnerabilities inherent in aviation systems. It is a great technique for getting started in threat modeling within the domain.

8.6.2 Attack Trees

In the context of cyber resiliency, attack trees are an extension of the concept of fault analysis techniques originally developed to support safety analysis of complex systems. Attack trees are created by decomposing threat vectors into the sequence of events that ultimately result in an

impactful cyber incident. While many threat modeling approaches are centered around identifying the bad things that can happen to a system, attack trees are more focused on how those threats can occur.

An attack tree is a structured, hierarchical diagram that breaks down the steps and conditions required for a particular threat to be realized. Each branch and node represent different attack paths and the sequence of events needed for the attack to succeed. This detailed breakdown allows engineers and analysts to thoroughly understand the various attack vectors and the specific conditions necessary for a cyber effect or incident to occur.

Figure 12 presents a high-level notional attack tree against aircraft navigation. In this case, the cyber effect is to create an event where the pilot does not know where the aircraft is currently operating by denying navigation. This partial diagram identifies four ways that could occur and further breaks down example methods for attacking an onboard GPS receiver.



Figure 12. Notional Attack Tree Against Aircraft Navigation

Constructing an attack tree creates nodes and branches that encapsulate how an attack vector can propagate through a system. This example also illustrates challenges of scaling up full attack tree analysis. Even this simple example had to be truncated in the interest of space and time. For highly complex systems it may not be possible to create exhaustive attack tree models.

While attack trees are a valuable tool for assessing the security of aviation systems, it is important to note that there is no one "correct" methodology to follow. Different analysts may have different perspectives and priorities, which can lead to variations in the structure and

content of the attack tree. Additionally, the complexity and unique characteristics of aviation systems mean that a single, universal approach may not be optimized for all systems.

Instead, it is recommended to think of attack trees as a flexible and iterative process that allows for the incorporation of new information and insights as they become available. This may involve revising and refining the attack tree as needed, in order to ensure that it accurately reflects the potential threats to the aviation system.

8.6.3 Attack Libraries

A library of known attacks can also be a useful source to identify threat patterns against systems. There are quite a few public repositories that catalog known types of attacks. It should be noted that most mature libraries are almost exclusively influenced by attacks against traditional IT systems. Consequently, they are not typically directly applicable to aviation platforms. However, the general patterns of the attacks in these libraries can often be used indirectly to determine how the spirit and intent of each entry could be asserted against an aircraft.

The MITRE Corporation maintains a library known as the “Common Attack Pattern Enumeration and Classification (CAPEC).” (53) CAPEC has cataloged well over 500 attack patterns and the list continues to grow. MITRE also publishes an ATT&CK (Adversarial Tactics, Techniques, and Common Knowledge) library (54) which can be thought of as a knowledge base of adversary tactics and techniques.

The MITRE Corporation maintains the Common Weakness Enumeration (CWE) (55) which is a community-developed system for identifying software and hardware weaknesses and vulnerabilities. They also maintain the Common Vulnerabilities and Exposures (CVE) system (56) that provides a reference method for publicly known information-security vulnerabilities and exposures.

Other attack sources include the OWASP (Open Web Application Security Project) Top Ten list (57) which highlights the ten most critical web application security risks. NIST maintains a National Vulnerability Database (NVD) (58) which is a repository of information about hardware and software flaws that could compromise cyber-enabled system security.

In addition, many companies and organizations offer proprietary threat libraries. However, these are typically also enterprise IT centric rather than specialized for the aviation industry. As a result, while these private resources can be beneficial, they lack the specialized information required to comprehensively protect aviation systems from targeted attacks and industry-specific threats.

While threat libraries are a valuable source of insight, it is crucial for aviation security professionals to supplement these general resources with additional, aviation-focused threat analysis to maximize robust and effective protection for aircraft systems.

8.6.4 Criticality Analysis

Criticality Analysis is a process outlined in the DoD Trusted Systems and Networks (TSN) Analysis document. (59) It is a systematic method used to identify and prioritize mission-critical functions and components within a system. This process begins with an end-to-end functional

decomposition, where systems engineers map out the mission-critical functions and the hardware, software, and firmware components that implement them. The analysis assesses the impact of potential failures or compromises on the overall mission, assigning criticality levels based on the severity of the impact.

The criticality levels range from total mission failure to negligible impact, ensuring that each function and component is evaluated for its importance to the mission's success. This detailed mapping and assessment helps identify which components need the most robust protection and which are more resilient. The process also involves identifying suppliers of critical components and documenting these findings. It provides a comprehensive view of the system's vulnerabilities and the necessary countermeasures.

Outcomes of criticality analysis inform subsequent steps such as threat, vulnerability, and risk assessments. It also informs the selection of security measures. By understanding the critical components and their vulnerabilities, aviation system architects and designers can implement appropriate countermeasures to mitigate risks effectively. This continuous evaluation and adaptation ensure that the system remains secure and resilient against evolving threats throughout its lifecycle.

Performing criticality analysis as a form of cyber assessment is essential for supporting the cyber resiliency posture of aviation systems. By mapping these key elements and assigning them levels of criticality, aviation system developers can implement a proactive approach to ensure that the most vital parts of the system are robustly protected against cyber threats.

8.7 Mitigating Threats

The purpose of threat modeling isn't to draw beautiful diagrams or enumerate threats that are known to exist for a system, but rather to identify how threats can be mitigated or eliminated from it.

The outputs of threat modeling should be integrated as appropriate with the current phase of development. For example, threat modeling performed during architecture and design should be used to identify potential modifications that would mitigate those threats ideally before that phase is finalized.

Threat modeling used at the end of the development lifecycle is useful for determining that the implementation of the system did not diverge from the planned architecture or unintentionally create threats that were not previously identified.

Threat modeling is also useful during sustainment to determine whether there are new or emerging threats against a system that require mitigation. In addition to architectural updates, mitigation techniques may include configuration settings, physical security controls, or operational procedures.

8.8 Integrating Threat Modeling with Risk Management

A cyber threat against a system that is not eliminated or effectively mitigated creates a latent cyber resiliency risk. In aviation, where the stakes are incredibly high, unaddressed threats can

compromise not only the integrity of the aircraft but also the safety of passengers and crew. Consequently, it is crucial to have a robust feedback mechanism in place that continuously aligns threat modeling with the broader Cyber Risk Assessment (CRA) and management framework for the organization. This alignment ensures that threats are not only identified and assessed but also effectively managed and mitigated throughout the system's lifecycle.

Without an integrated feedback loop with cyber risk management, the insights gained from threat modeling may not be fully achieved. This feedback mechanism allows for the constant update of threat models and ensures that the system's risk posture is dynamically managed. For aviation systems, this means regularly revisiting and revising the threat models to address new vulnerabilities, evolving threats in the environments, and changes in the operational use of the platform. Such a mechanism ensures that the security measures in place remain current against the threat in cyber contested airspaces.

Integration of the outputs of threat modeling with cyber risk management is essential for maintaining the cyber resiliency posture of aviation systems. This approach ensures that threats are continuously identified, assessed, and mitigated. Given the critical nature of these systems and the high stakes involved, this integrated approach is an imperative.

8.9 Alternative Threat Modeling Approaches

This FSAD Guidebook recommends taking an integrated approach to threat modeling that is based on mission performance and the mission essential functions. This section describes alternative approaches that are sometimes couched as threat modeling efforts.

It is important to emphasize that these alternative approaches are not necessarily recommended to satisfy threat modeling objectives. These descriptions are provided so stakeholders can recognize when a threat modeling effort might be too narrowly focused, along with the associated pitfalls and limitations. While the information produced by these alternative approaches is useful, aviation system architects and designers should be aware that it is likely to result in an incomplete system assessment for that domain.

A **compliance approach** to threat modeling is based on a set of security controls, such as the ones provided by the Risk Management Framework (RMF). While compliance with these controls is sometimes mandated for security authorization of the system, it does not assure cyber resiliency because the mission and mission essential functions are not considered. Furthermore, a compliance approach can result in implementation of controls that do not mitigate an actual threat or risk to the system. A strict compliance approach is not recommended to satisfy the objectives of threat modeling.

Similarly, a **vulnerability approach** is usually indicative of a highly reactive operational environment. Vulnerabilities that are detected by automated security scans exist at a micro level that may or may not be indicative of the larger scale threat against the system. It isn't a good way to identify threats because only known vulnerabilities are identified. Additionally, automated scanning tools and threat intelligence required for meaningful scanning isn't usually available for most aircraft platforms. Consequently, a strict vulnerability approach is not recommended to satisfy the objectives of threat modeling for aviation systems.

The threat modeling approach recommended by this Chapter can be characterized as **integrated threat driven**. In the integrated approach, the missions, mission essential functions, and any other critical assets of the system are assessed against damage that can result from threats that are specific to the environment and aviation platform implementation. These threats to the system assets and resources are the primary driver used to gauge the efficacy of the system architecture, design, and implementation in preventing and defending hostile cyber action.

8.10 Summary and Additional Resources

Threat modeling is a proactive technique that identifies potential risks and vulnerabilities throughout the development and operational lifecycle of a system. It is an effective cyber resiliency analysis approach for aviation systems as threat modeling can identify potential vulnerabilities and threats that may be asserted in cyber contested environments. In an ideal world, threat modeling enables the selection of security measures during architecture, design, and implementation that eliminates risk to the aircraft. However, it isn't always possible to completely mitigate or eliminate every residual threat or risk to the system.

Consequently, the insights gained through threat modeling, including risks that are not fully resolved must be integrated with a cyber risk management system to ensure that threats are continuously tracked and periodically reassessed. In aviation, where the stakes are incredibly high, unmanaged threats can be catastrophic.

Given the critical nature of aviation systems and the high stakes involved, a highly disciplined practice of integrated threat modeling is an imperative for maximizing the cyber resiliency posture.

For more information on the engineering art of threat modeling, the following resource is recommended:

- **Threat Modeling: Designing for Security, Adam Shostack, 1st Edition, 2014. (50)**

This is an industry recognized authoritative resource for anyone who wants deeper insight into threat modeling. This book offers specific, actionable advice on how to incorporate security into the design of complex systems from the outset. It covers various threat modeling approaches, including asset-centric, attacker-centric, and software-centric models. This book also shares proven techniques used at Microsoft and other top companies.

Chapter 9

Cyber Risk Assessment

Cyber Risk Assessment (CRA) is the process of identifying and evaluating the potential risks associated with a given system. This analysis exercise is a critical aspect of any comprehensive cyber resiliency engineering effort, as it assesses how potential threats and vulnerabilities might impact execution of mission essential functions. CRA is also a vital mechanism for prioritizing mitigation of residual risk within the system. In essence, CRA provides a snapshot of cyber risk associated with a system at the time the assessment was performed.

CRA is an integral activity during the development and operation of aviation systems. It is used throughout the Systems Engineering lifecycle to identify and assess potential cybersecurity risks associated with the architecture, design, and implementation. Evaluating the security of the hardware, software, and subsystem components onboard the aircraft in the face of potential threats and vulnerabilities can be used as feedback toward creating a platform that exhibits the properties of intrinsic cyber resiliency. CRA can be thought of as a measuring stick for how effectively cyber resiliency objectives are realized by the system implementation.

An ongoing practice of CRA is also important for aircraft during operation and sustainment. By continuously evaluating the evolving threat landscape and assessing the resilience of existing safeguards, emerging vulnerabilities can be proactively identified, assessed, and mitigated to minimize operational impacts.

Mastering the art of CRA is a critical skill for organizations that develop and manufacture aviation systems. Proficiency performing risk assessments informs the development process with actionable information for how to improve the cyber resiliency posture of the platform. At the same time, it provides valuable feedback on how effective cyber resiliency engineering efforts are being performed.

9.1 Characteristics of Effective Cyber Risk Assessment

One of the most valuable resources that describes how to perform CRA is NIST Special Publication 800-30, “Guide for Conducting Risk Assessments.” (60) Rather than describing a singular best practice for executing CRA, this NIST standard describes the general characteristics of an effective and disciplined CRA practice.

In fact, the special publication outlines several cautionary notes that warn readers that effective CRA initiatives will vary in formality, rigor, and level of detail. Indeed, different degrees of each of those attributes will vary based on what is being evaluated as well as the current stage within the systems engineering and sustainment lifecycles. Organizations have flexibility in the

selection of methodologies, tools, and techniques to be used in their risk assessments. With that discretion comes great responsibility.

The key characteristics of an effective CRA methodology are as follows:

- **Systematic and Structured Approach:** While there is great flexibility in the methods used, it should be a well-defined and systematic process. That includes documented steps and procedures for identifying, analyzing, and evaluating risks. This ensures consistency and repeatability in risk assessments.
- **Context-Specific:** The results of cyber risk assessments are expressed within the specific context of the platform and the unique operational environment. For example, potential vulnerabilities identified by automated scanning is not a substitute for cyber risk assessment in isolation from the context of system operation.
- **Operational Stakeholder Involvement:** For aviation systems it is important that key operational roles have a voice to provide input and feedback during the assessment. For aircraft platforms this will typically include pilots and maintenance personnel at a minimum.
- **Platform Subject Matter Expertise (SMEs):** For highly complex aviation systems, it is important to have personnel who understand how the system and subsystems under evaluation are designed to operate. SMEs can articulate how different parts of the system interact, which is essential for correctly identifying vulnerabilities and assessing the potential impact.
- **Informed Threat Expertise:** The cyber risk assessment includes people who have knowledge or insight into the type of adversarial threats the system will be subjected to. This includes the Tactics, Techniques, and Procedures (TTPs) that the system is expected to face in the operational environment.
- **Tailorable Scoring Rubrics:** A rubric is a scoring tool used to evaluate and assess the methods of characterizing threats and impacts against the system. When performing cyber risk assessment, the rubrics provide criteria for determining quantitative or qualitative scoring. Rubrics may be tailored to the specific system which can improve the granularity and fidelity of the results.
- **Risk Prioritization:** The methodology must include mechanisms for prioritizing identified risks. This helps program management and system operators make informed decisions when allocating resources to mitigate any findings.
- **Clear Results and Reports:** The results of cyber risk assessment should be communicated clearly and effectively to all relevant stakeholders. This includes providing actionable recommendations and strategies for risk mitigation.
- **Provisions for Regular Updates:** Cyber risk assessments produce data that is a snapshot in time. Mature cyber risk assessment initiatives include provisions for periodically reassessing cyber risk within a highly dynamic threat environment.

- **Satisfies Compliance Objectives:** The risk assessment methodology should align with regulatory requirements and compliance standards.

The ultimate objective of CRA is to inform decision makers. That includes project and program managers who are responsible for allocating resources required to mitigate risks. However, the architects, designers, and people implementing the system also make decisions which should be informed by the data and insight emerging from an ongoing practice of CRA.

9.2 Cyber Risk Assessment Process

Effective Cyber Risk Assessment is performed as four distinct phases. These are Preparation, Conducting, Analyzing and Reporting, and Maintaining.

9.2.1 Preparation for Cyber Risk Assessment

Organizations that perform effective CRAs understand the importance of investing in preparation. This phase can be thought of as a series of decisions, each of which guides and informs the exercise. Before anything else can occur, the purpose and the scope of the CRA must be determined. All stakeholders should understand and agree to why the cyber risk assessment is being performed within the context of the system.

Once the purpose of the CRA is understood and agreed to, the scope of the assessment is then determined. That includes which parts of the system are being evaluated as well as the types of threats and risks that will be considered. That goes hand in hand with identifying and documenting any assumptions and constraints that bound the CRA effort.

The decisions of the system under evaluation will guide selection of the cyber risk methodology that is best suited to produce the type of information that the CRA exercise is designed to elicit. Those same data points drive the identification and selection of stakeholders, Subject Matter Experts (SMEs), and the threat expertise that will be required. If the methodology selected defines specific roles and responsibilities during the exercise, each participant will be slotted into one of those and briefed on the expectations.

As the personnel who will be supporting the CRA come on board, a series of briefings on the system under test will be conducted. For larger scope assessments, that information leveling process will likely be conducted iteratively in phases. The people who are identified to support the cyber risk assessment exercise will generate Requests for Information (RFIs) identifying data and documentation needed so that CRA participants go into the exercise with solid understanding of the system under evaluation. Additionally, if there is specific threat intelligence or briefings needed, those meetings would be planned and conducted during the preparation phase.

The preparation phase is also the optimum time for reviewing any scoring rubrics that are used. In many cases, it will make sense to tailor the rubric based on the system under evaluation as well as the scope of the cyber risk assessment exercise.

Logistical arrangements for when the cyber risk assessment exercise is conducted are locked down during the preparation phase. This necessarily includes the location and duration and any data sensitivity constraints.

9.2.2 Conducting the Cyber Risk Assessment

Paradoxically, in calendar duration, the shortest part of performing a CRA is conducting the actual exercise. The ultimate goal is to elicit information about cyber risks to the system under evaluation. That is performed in accordance with the purpose, scope, and methodology selected during the planning phase.

A typical sequence of events is to identify adversarial threat objectives against the system, and postulate what threat events could produce desirable cyber effects toward those goals. The analysis will consider what vulnerabilities or susceptibilities could be leveraged in the pursuit of the cyber adversarial intent against the system.

Most cyber risk assessment methodologies generate estimates that characterize the likelihood or probability that each threat vector is believed to have against the system under evaluation. Once that is determined, the operational impact of the intended cyber effect is separately decided. How each of those data points are elicited is dependent on the methodology used, including the rubrics or tailored rubrics created for the exercise. Additionally, the fidelity of the information collected will be impacted by the skill level and expertise of the personnel playing the key roles in the CRA execution.

Prioritizing data elicitation and collection is crucial because that forms the foundation of accurate and comprehensive analysis and reporting in the subsequent phase. When the key stakeholders participating in the exercise are gathered together, the absolute priority is the data collection.

By focusing on robust data collection, the assessment can ensure that all critical factors are considered, providing a solid basis for subsequent analysis. The writing of the report, which synthesizes and interprets the collected data, is best left to the analysis and reporting phase.

This allows the team to carefully analyze the gathered information, draw meaningful conclusions, and provide well-informed recommendations. That supports creation of a final report that is both insightful and actionable.

9.2.3 Analyzing and Reporting Cyber Risk Assessment Results

Once the CRA exercise is completed, the data collected is analyzed, synthesized into a report, and communicated to the appropriate stakeholders.

The analysis phase of data collected is critical because it transforms raw data into actionable insights. During this phase, the collected information about threats, vulnerabilities, and existing controls is validated, interpreted, and organized to best communicate the outcomes. This analysis potentially enables the identification of patterns, correlations, and emerging trends that may not have been apparent as the exercise was being conducted.

The written report is intended to effectively identify, prioritize, and communicate cyber risks against the system. Ultimately, this phase informs the development of effective risk mitigation strategies, actionable recommendations, and insights.

The report emerging from effective CRA efforts typically includes the following elements:

- Identification of threats to the system.
- Identification of known or potential system vulnerabilities.
- Assessment of the impact of successful exploitation of a vulnerability by a threat.
- Estimation of the likelihood that harm will occur from successful exploitation of a vulnerability by a threat.

This data in aggregate is an expression of cyber risk associated with the system under evaluation. A typical way that information is expressed is within a “Risk Cube.” Figure 13 provides an example of a common format used.

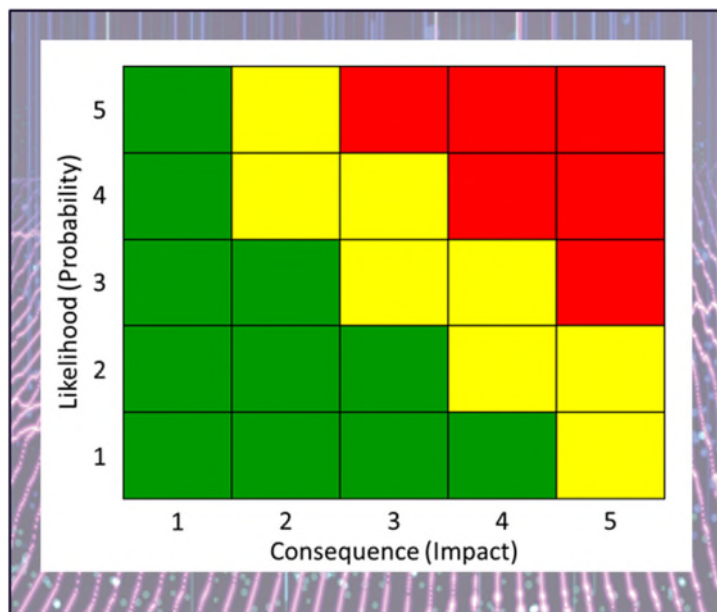


Figure 13. Risk Cube Example

While different risk cubes may have scoring variations, in general the red squares are high risks, yellow are moderate risks, and green is regarded as low risks.

Each threat vector analyzed during the CRA exercise will be reflected as a data point within the risk cube in the final report. The precise location is determined by the data collected during the CRA exercise for the consequence of impact and the likelihood of occurrence. Those numbers are based on the SME evaluations and the scoring rubric used at the event.

A completed risk cube provides a comprehensive, visual representation of the interrelationships between various risks to the system. It illustrates the complete risk landscape including the relative scoring reflected by the likelihood and impact.

The output of the analysis is a written report and appropriate communication mechanisms that deliver cyber risk assessment results to the appropriate stakeholders. That is essential to ensure decision makers have access to the risk data and actionable recommendations. The information in cyber risk assessment reports is critical for resource prioritization and allocation.

9.2.4 Maintaining Cyber Risk Assessment Results

Cyber risk assessment doesn't end when the report is delivered. The risk posture of a system resulting from CRA execution is a snapshot that represents a single point in time. The threat landscape facing aviation systems is highly dynamic. Additionally, knowledge and awareness of the nature of vulnerabilities and adversarial cyber threats is continuously refined.

A system of ongoing review and update is needed to ensure that the CRA data being used to drive decisions remains valid and up to date. Identification and evaluation of new and emerging threats to the system is an ongoing process. Additionally, system events and threat intelligence occurring after the CRA is completed may provide additional insight into cyber risk assessment reports.

In most instances, CRA is not a "one and done" occurrence.

9.3 Cyber Risk Assessment Integrated With Systems Engineering

Some form of CRA should be a recurring and integral part of every phase of the development lifecycle described in Section 5.5. It should be leveraged as a mechanism to determine if requirements, architecture, design, and implementation are achieving the objectives of cyber resiliency established for the system under development.

CRA performed at the earliest stages of development typically result in lower cost and schedule impacts when addressing mitigation opportunities that are identified. On the other hand, CRA executed at the late stages of the development lifecycle and during operation will yield higher fidelity results that closely align with the actual fielded system.

The integration of CRAs throughout the entire system development lifecycle is a crucial continuous investment towards ensuring cyber resiliency. CRAs are best regarded as an ongoing process leveraged to enhance the platform's cyber resiliency posture. It is an effective way to identify opportunities to create and sustain platforms that protect against ever-evolving threats.

9.4 Selection of a Cyber Risk Assessment Methodology

One of the challenges in getting started with CRA is selection of the methodology to be used. The cybersecurity industry has published a plethora of approaches that are compliant with NIST Special Publication 800-30, "Guide for Conducting Risk Assessments." (60) Each candidate approach comes with its own strengths and weaknesses. The decision can seem overwhelming.

This section provides additional information and considerations for making an initial methodology selection.

9.4.1 The Case for Mission Based Cyber Risk Assessment

It may seem counter intuitive, but many published methods of cyber risk assessment were created prior to the emergence of our modern understanding of cyber resiliency. Since that attribute is critically important for aviation systems, it is important to know which established methods are best suited for that objective. Selecting approaches that emphasize the operational context and mission-critical functions of the systems being assessed maximizes the opportunity for CRA to positively impact cyber resiliency.

The primary mission for aviation systems is to ensure safe and reliable flight operations. Overlooking mission impact during cyber risk assessment causes critical risks to be underestimated or ignored. Consequently, the mitigation strategies developed may not adequately protect against scenarios that could lead to catastrophic outcomes, such as loss of aircraft control, communication failures, or navigation disruptions. Therefore, a mission-focused approach is essential to ensure that all relevant risks are comprehensively assessed and appropriately mitigated.

Mission Based Cyber Risk Assessment (MBCRA) is the industry term for methodologies that prioritize the identification and mitigation of cybersecurity risks based on their potential impact on the specific missions or critical functions. Rather than focusing solely on technical vulnerabilities or generic threats, this approach evaluates how cyber risks can affect the ability of the aircraft to achieve its key operational objectives and maintain essential capabilities.

Given the critical importance of cyber resiliency for aviation systems, MBCRA approaches are strongly recommended. That sub-category of CRA is simply better suited for aircraft systems than non-mission based methods.

9.4.2 Cyber Table Tops (CTT)

The Cyber Table Top (CTT) is a MBCRA methodology which was originally created by the Lockheed Martin National Cyber Range, transferred to the Navy, and is now officially maintained by the DoD. (61) A CTT is a collaborative cyber risk assessment approach that is oriented around mission analysis.

A CTT is a lightweight, intellectually intensive exercise that explores the effects of cyber offensive operations on the capability of systems to carry out their missions. It is a wargame-like exercise between two teams with opposing purposes. An operational force is charged with executing a defined set of missions. The opposing force is tasked with causing those missions to fail.

For systems and teams that are just getting started with CRA, CTTs are an excellent starting point. The DoD process is publicly available and well documented. The methodology is straight forward and easy to both follow and implement. Additionally, CTTs are an efficient means of producing the type of insight that might point to other more specialized and focused forms of MBCRA for follow on analysis.

9.4.3 Legacy CRA Methods

Legacy aviation systems that went into production prior to the publication of NIST 800-30 may have previously used a methodology that satisfies the spirit and intent of the criteria defined in that special publication. If the previous CRA approach satisfies that criteria, there may be no compelling reason to update to a more current methodology. Aviation platforms that previously had rigorous CRA performed in accordance with a well-documented process should not assume that CRA must be redone just because NIST published an updated standard.

It is recommended that any CRA approaches previously used on legacy platforms be evaluated against the criteria defined in NIST 800-30 and as described in section 9.1 of this FSAD guidebook. In some cases, continuing to use the prior approach implemented for the system may be a viable option. However, that reconciliation exercise may reveal that the legacy method has gaps that must be closed to better leverage CRA toward advancing the cyber resiliency objectives of the platform.

As another cautionary note, some DoD aircraft programs have reportedly satisfied Risk Management Framework (RMF) CRA requirements by mapping data gleaned from security control assessment scans to risk cubes. While that approach produces valuable information about latent known deficiencies in a system, it is not particularly useful for understanding cyber resiliency. It is strongly recommended that platforms that only have these control/compliance versions of CRA explore upgrading that data with true MBCRA exercises.

It is important to emphasize that just because a security control assessor accepted a particular method of CRA in the past, it does not necessarily mean that the data is sufficient for driving cyber resiliency efforts and decisions. That is a consideration that cannot be neglected for aviation platforms.

9.5 Cyber Risk Assessment is not a Precise Measurement of Risk

The data emerging from CRA exercises should not be regarded as a precise measurement of risk. Instead, it should be viewed as an informed estimate that helps organizations understand potential threats and vulnerabilities within a certain context. This perspective is crucial because cyber risk is inherently dynamic and complex. Additionally, it is influenced by emerging threats, changing technologies, and new vulnerabilities. As such, risk assessment results are not definitive but rather indicative. CRA reports are a snapshot of the risk landscape as understood at a particular moment in time.

NIST SP 800-30 includes a cautionary note warning that risk assessments should be interpreted with caution. The special publication acknowledges that the results are subject to uncertainty due to the difficulty of predicting future threat events and the complexities involved in estimating the impact of what transpires. It highlights that risk assessment data relies on both qualitative and quantitative inputs, which can be influenced by subjective judgments, assumptions, and the quality of available information. Therefore, while the data provides valuable insights for decision-making, it should not be considered as providing absolute or precise measurements of

risk. This recognition helps ensure that stakeholders maintain a realistic understanding of the limitations and uncertainties associated with CRAs.

Understanding the approximate nature of cyber risk assessment data encourages organizations to adopt a more flexible and adaptive approach to risk management. Rather than relying solely on static risk assessment reports, organizations should continuously monitor their risk environment, update their assessments as new information emerges, and be prepared to adjust their risk mitigation strategies accordingly. By acknowledging the inherent uncertainties and limitations of risk assessment data, organizations can make more informed risk management decisions.

9.6 Cyber Risk Assessment and Broader Business Impacts

The data from CRA exercises should not be isolated to a cyber-only context. The information is also impactful when integrated into a business-wide risk management effort which considers a holistic view of the total risk landscape. Cyber risks are interconnected with various other types of risks, including operational, financial, reputational, and compliance. Isolating cyber risk assessment data can lead to a fragmented understanding of how these risks interrelate and impact the broader organization.

Integrating cyber risk assessment data into a system-wide risk management program enhances decision-making and strategic planning. It allows for a comprehensive evaluation of risk interdependencies and the identification of cascading effects that a cyber incident could trigger across various domains. This integrated approach helps prioritize risks based on their overall impact on the aircraft platform, rather than considering cyber risks in isolation. This holistic perspective ensures that the organization is better prepared to respond to complex, multi-faceted risk scenarios.

9.7 Cultural Attributes of Effective Cyber Risk Assessment

Chapter 2 of this FSAD Guidebook built a foundation of understanding that aviation platforms operate in cyber contested environments. While CRA efforts are most effective when the participants and stakeholders participating in the exercise are knowledgeable and aware of threats to the aircraft, the most important success attribute is a fundamental belief that the threat is real.

Providing CRA participants with a cyber threat briefing during the preparation phases of the exercise can build and reinforce the mindset necessary for effective outcomes and analysis. Even when detailed threat information is not available for a specific platform, awareness of past incidents and indications help teams understand that the threat is real. Knowledge of the Tactics, Techniques, and Procedures (TTPs) that are used in this domain illustrates the art of the possible.

Additionally, CRA exercises are most effective when the organization and stakeholders embrace the practice as an effective method for identifying opportunities to improve the cyber resiliency posture of the platform. That contrasts with viewing cyber risk assessment as one last hurdle to clear as a requirement for Approval to Operate (ATO). When the only cyber risk assessment

performed on a platform comes immediately before authorization is needed, participants will feel pressure to show the platform is secure which could reduce the benefits of the exercise.

Building cyber resiliency into aviation systems requires a cultural belief in the reality of the cyber threat and a mindset that every cyber risk assessment is an opportunity to do better.

9.8 Summary and Resources

Cyber Risk Assessment (CRA) is a vital activity for ensuring the cyber resiliency of aviation systems. CRA provides a systematic and structured approach to identifying, analyzing, and evaluating potential cybersecurity risks associated with aviation systems, allowing organizations to understand how various threats and vulnerabilities might impact critical flight operations and safety. By prioritizing mitigation and correction of residual risk, CRA helps organizations allocate resources effectively and make informed decisions regarding risk management. Furthermore, CRA is an ongoing practice essential for proactively identifying, assessing, and mitigating emerging vulnerabilities during system operation and sustainment. By mastering the art of CRA, organizations that develop and manufacture aviation systems can ensure that their products exhibit the properties of intrinsic cyber resiliency, protecting against ever-evolving threats and ensuring safe and reliable flight operations.

For more information on Cyber Risk Assessment, the following resources are recommended:

- **Guide for Conducting Risk Assessments, National Institute of Standards and Technology (NIST) Special Publication 800-30, Revision 1, September 2012. (60)**

A comprehensive, structured approach to evaluating risks to information systems, integrating risk assessments into a broader organizational risk management framework. It outlines a detailed process for preparing, conducting, communicating, and maintaining risk assessments, emphasizing the relationships among threats, vulnerabilities, impact, and likelihood. Applicable across multiple organizational levels, it offers flexibility in using quantitative, qualitative, or semi-quantitative methodologies. The guide supports enhanced decision-making, effective resource allocation, and compliance with national standards, making it an essential resource for thorough and standardized cyber risk assessments.

- **The Department of Defense Cyber Table Top Guide, Version 2.0, 16 September 2021. (61)**

The authoritative resource for conducting CTT exercises, which simulate cyber offensive operations to assess their impact on mission-critical systems. It offers a structured, step-by-step methodology for planning, executing, analyzing, and reporting these exercises, emphasizing mission-based risk assessment. It fosters collaboration among multiple teams and provides practical methods of risk quantification. This guide emphasizes comprehensive, actionable insights into potential vulnerabilities and their operational impacts, supporting informed decision-making and effective cyber defense strategies.

Chapter 10

Cyber Test and Evaluation

The developers of aviation platforms must understand how systems will be evaluated and tested for cyber resiliency throughout the product development and operational lifecycle. In many cases, independent cyber testing is required to achieve authorization and airworthiness certifications. The information presented throughout this FSAD Guidebook describes how aviation systems can be architected, designed, and implemented to achieve success when those critical evaluations are performed.

This chapter describes how Cyber Test and Evaluation (CT&E) is executed to verify the trustworthiness and cyber resiliency of aviation systems. Anticipation and knowledge of what occurs during that phase of testing can help guide decisions that are made throughout the development process that increase not only the results from those testing events, but also the overall cyber resiliency posture of the platform.

CT&E testing can encompass a wide range of verification activities. That includes vulnerability assessments, security control testing, and system survivability evaluation. It also includes various forms of offensive and defensive testing. One success factor that contributes to good outcomes during these evaluations is early and continuous engagement with the testing community throughout the system development lifecycle.

10.1 Integrating Cyber Test and Evaluation in the Systems Development Lifecycle

One of the most authoritative sources for understanding the CT&E process is the DoD Cybersecurity Test and Evaluation Guidebook. (62) That resource provides a roadmap for how the testing community can contribute to the development of cyber resilient systems throughout the acquisition lifecycle. It identifies six phases of CT&E engagement.

1. **Understand Cybersecurity Requirements.** The objective of this phase is to examine the system's cybersecurity, cyber survivability, and operational cyber resilience requirements. Considering cyber testing at this stage ensures that all stakeholders have a clear understanding of the cybersecurity requirements including how they will be evaluated.
2. **Characterize the Attack Surface.** In this phase, testing resources identify vulnerabilities and potential attack paths that adversaries might exploit. One approach to this engagement for engaging cyber testing expertise is to perform cyber risk assessment as identified in Chapter 9. This involves analyzing the system architecture, components, data flows, and mission dependencies to understand the cyber-attack surface. It is an advanced look at testing approaches likely to be used during independent cyber testing.

3. **Cooperative Vulnerability Identification.** This phase invokes early testing to verify cybersecurity and operational cyber resilience. As the name suggests, it also identifies vulnerabilities. The purpose is to inform system designers, developers, and engineers about necessary improvements to enhance system cyber survivability and operational cyber resilience while development is still in progress. This phase should be performed iteratively throughout the development lifecycle.
4. **Adversarial Cybersecurity Developmental Test and Evaluation.** During this phase, test teams use adversarial techniques to test the critical functionality of the system using a mission context. This evaluation focuses on verifying that the system can withstand adversarial cyber activities.
5. **Cooperative Vulnerability and Penetration Assessment.** The purpose of this testing is to characterize the cybersecurity and resilience of a system in an operational context. It is a more mature advanced look at testing approaches likely to be used in independent cyber evaluations.
6. **Adversarial Assessment.** This final phase characterizes the operational mission effects of threat-representative cyber activities against the system. It assesses the effectiveness of defensive capabilities and the overall resilience of the system when faced with realistic cyber threats. The goal is to understand the impact of cyber activities on critical missions and to verify the system's defensive posture.

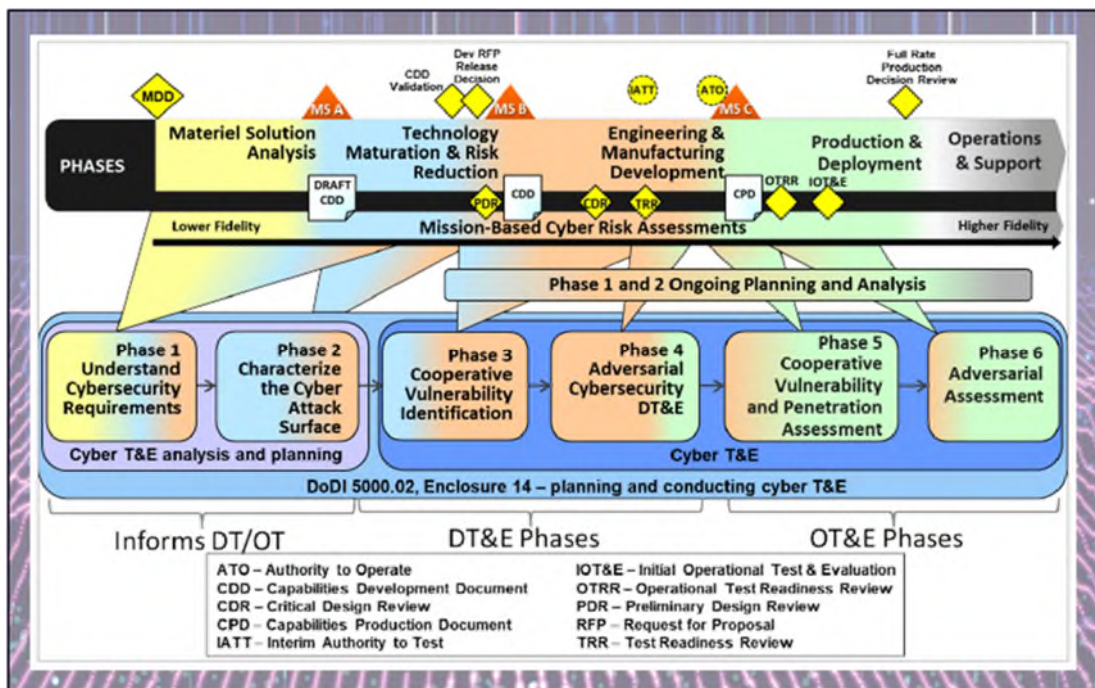


Figure 14. CT&E and the DoD Acquisition Lifecycle (62)

Figure 14 illustrates these six phases mapped against the DoD Acquisition Lifecycle. The iterative nature of these phases allows for continuous refinement and improvement of the

aviation system cybersecurity posture throughout the development lifecycle. It is a best practice for ensuring robust architectures and designs that support cyber resiliency.

The six phases of the CT&E process can also be neatly mapped against the Systems Engineering lifecycle that was introduced in Section 5.5 of this FSAD Guidebook. Those relationships are captured in Figure 15.

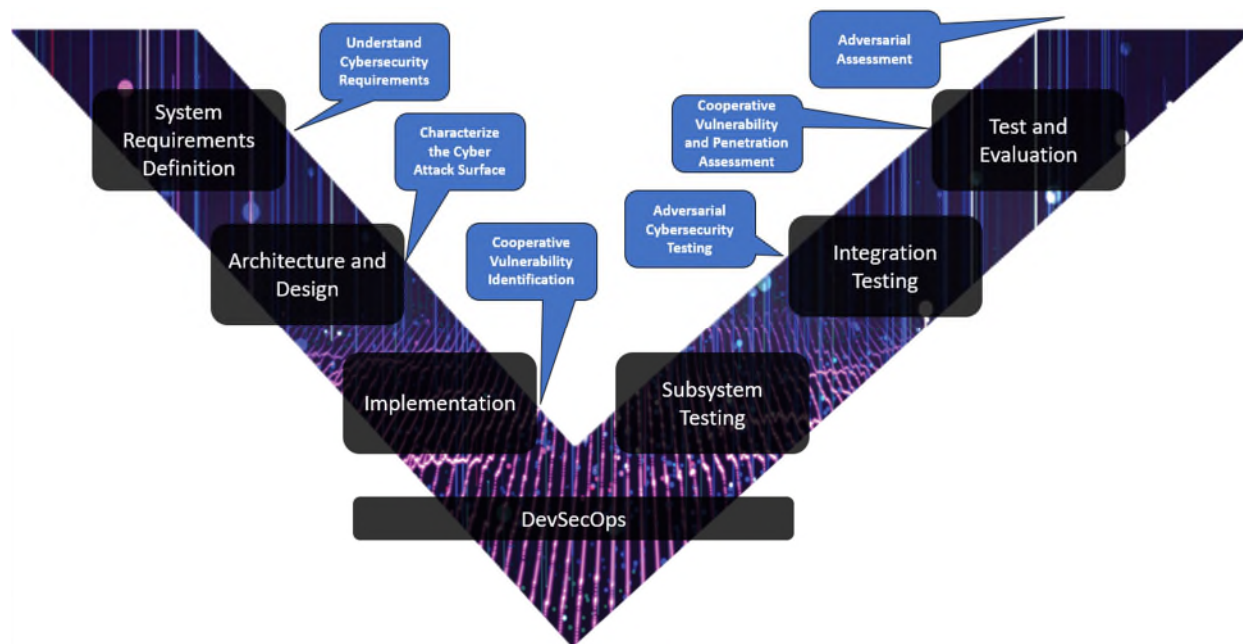


Figure 15. CT&E and the Systems Engineering Lifecycle

On many legacy aviation systems, CT&E may not have been performed at all against the aircraft prior to operational use. In fact, the DoD acquisition agencies and defense industrial base are still grappling with the most effective methods for performing that vital testing on aircraft platforms and related weapons systems. This delinquency leads to late discovery of cybersecurity issues during operational use when vulnerabilities are extremely costly to mitigate.

To achieve robust cybersecurity in aviation, it's imperative to integrate CT&E throughout the entire aircraft development lifecycle. Continuous evaluation is the best way to ensure that the aircraft is fortified against potential cyber threats and that risks are identified, eliminated, or mitigated. Strong collaboration between the designers of aviation systems, cyber resiliency engineers, and skilled cyber testers is a critical best practice for building cyber resilient aircraft.

10.2 Cyber Test and Evaluation and Cyber Resiliency

While CT&E can be used to assess if there are latent cybersecurity concerns with a system, it cannot be used as an effective mechanism to bestow the characteristics of intrinsic cyber resiliency into it. While testing can certainly uncover issues that must be addressed, it is a reactive rather than a proactive means of performing that identification.

Additionally, CT&E can only prove the absence of cyber resiliency in a system. The converse to that statement is not true. CT&E that does not identify cyber security concerns does not prove that the system is cyber resilient, but rather only that no concerns were identified.

The bottom line is that it is impossible to test a system to the point of cyber resiliency. That characteristic only emerges from a foundation of rigorous cybersecurity best practices applied throughout the development engineering lifecycle.

10.3 Security Control Testing/Compliance

For DoD aviation platforms, security control assessment is typically required to receive an Authorization to Operate (ATO). While that is a necessary achievement, it is important to emphasize that security control compliance is insufficient as evidence that a platform will operate as expected in a cyber-contested environment. Security control verification assesses whether the controls are implemented within a system. It does not assess mission effectivity in the face of cyber adversity during operational use.

Executing security control verification is an integral part of each program plan. However, if the only testing performed against an aircraft is evaluation for compliance of security controls, the platform will likely receive an ATO, but with a high degree of uncertainty about the actual cyber resiliency posture.

10.4 Blue Team Testing

“Blue” is a general industry term for a team that focuses on defending an organization’s computing assets against adversarial cyber-attack. A blue team analyzes systems to identify security flaws and to verify the effectiveness of the implemented security measures. Blue Teams work cooperatively with system developers and owners to identify vulnerabilities, develop detection mechanisms, and deploy remediations to improve the security posture.

Blue Teams use many of the same tactics and techniques as the Red Team, which is described in Section 10.5. The difference in the blue team activity is a focus and prioritization on helping development and operational teams understand how cyber activity can be prevented, detected, and mitigated rather than demonstrating specific cyber effects.

Blue team testing is roughly equivalent to the Cooperative Vulnerability and Penetration Assessment phase described in Section 10.1.

10.5 Red Team and Penetration Testing

A Red Team simulates adversarial cyber-attacks by playing the role of a cyber aggressor. The objective of the Red Team is to demonstrate the impacts of successful incursions resulting in cyber effects against the system under evaluation. That knowledge and insight can be leveraged by developers to mitigate the threats and risks demonstrated by the Red Team.

A good Red Team will emulate the TTPs of cyber adversaries without doing permanent harm or damage to the system. Red Teaming is a useful and essential mechanism for understanding how a system will perform under cyber duress.

Red Team testing is roughly equivalent to the Adversarial Cybersecurity Developmental Test and Evaluation phase described in Section 10.1.

10.6 Cyber Verification and Validation

Before a system can be considered complete, all allocated requirements must be verified to ensure that they have been satisfied through implementation. The requirements selected and allocated for the explicit purpose of cyber security and cyber resiliency are no exception. Verification of cyber security requirements is required in every system.

This verification can consist of testing of functional controls. In some cases, verification of a requirement is only possible through analysis of evidence produced. For example, if static code analysis was imposed as a requirement for development, there is no test in the integrated environment that can show it was done. However, evidence collected during static code analysis reporting and dispositions can be used as a verification mechanism.

Validation is based on evidence that the overall system fulfills mission objectives when used in the intended operational environment. Validation of systems security engineering objectives occurs through demonstration of the desired level of trustworthiness, survivability, and cyber resiliency.

Security validation demonstrates the effective operation of the system while under threats. In short, it is a testament that the system can execute the mission in a cyber contested environment.

10.7 Summary and Additional Resources

Cyber Test and Evaluation (CT&E) is a critical process for ensuring the trustworthiness and cyber resiliency of aviation systems. By understanding and integrating CT&E throughout the systems development lifecycle, developers can make informed decisions that enhance cyber resiliency and minimize the risk of vulnerabilities. The DoD Cybersecurity Test and Evaluation Guidebook provides a roadmap for CT&E engagement, with six phases that can be mapped against the acquisition and systems engineering lifecycles.

While CT&E is important for identifying and addressing cybersecurity concerns, it cannot be used as the sole means of achieving cyber resiliency. By utilizing these CT&E processes and engaging with the testing community throughout the development lifecycle, developers can build more cyber resilient aviation systems.

- **Cybersecurity Test and Evaluation Guidebook, Department of Defense, Version 2.0, Change 1, February 10, 2020.**

A comprehensive guide for integrating cybersecurity test and evaluation throughout the development lifecycle. It provides more detailed information on the structured six-phase CT&E process and emphasizes early and continuous involvement of cybersecurity practices, collaboration between various stakeholders, and alignment with overall system engineering and T&E activities. It is a vital resource for program managers, architects, designers, and testers to use in the development of secure and cyber resilient systems.

Chapter 11

Securing the Complete Air System

Aviation security is not just about securing the airplane itself. It also encompasses maintenance devices, support systems, and communication interfaces. Modern aircraft are increasingly connected to a variety of systems, all of which can be potential conduits for cyber-attacks.

Under some circumstances compromised ground systems can be used by sophisticated cyber attackers to disrupt aviation operations, cause system malfunctions, or even take control of the aircraft. Consequently, it is essential to consider the security of all maintenance and support systems when implementing aviation security measures. The entirety of the air system must be secured.

11.1 Enterprise IT and Aviation Cybersecurity

The support and maintenance equipment that connects to aircraft platforms are frequently built on enterprise IT-based platforms. That approach is generally more cost effective than developing custom computing devices. Ground support systems typically don't need capabilities provided by the specialty embedded computing devices that the aircraft requires. Additionally, enterprise IT computers are ubiquitous and thus familiar to more of the general public.

Unfortunately, that support equipment also introduces new security challenges, as these systems can be potential targets for cyber-attacks. Enterprise IT-based systems are often connected to other devices and networks, which can provide attackers with multiple avenues for gaining access to the aircraft platform. Additionally, compromising a standard desktop computer generally does not require the same degree of sophistication as the aircraft itself.

This FSAD Guidebook is targeted toward flight platforms. However, it would be remiss to not observe the significant attack surface and vulnerability that can come via support and maintenance equipment.

Securing aircraft maintenance and support systems that use traditional IT is critical to supporting the safety and security of aircraft platforms. To effectively defend the entire aviation system, it is important to follow industry best practices for IT security. Regular vulnerability assessments and penetration testing can help identify potential weaknesses in support and maintenance equipment. That enables organizations to proactively address any issues before they can be exploited by attackers. Furthermore, implementing intrusion detection and prevention systems can help detect and respond to potential security threats in real-time. By following industry best practices for IT security, organizations can ensure the safety and security of their aircraft maintenance and support systems and protect against potential cyber-attacks.

11.2 Air System Components

The definition of an air system includes the aircraft, mission/flight planning, maintenance devices, and logistics systems. Each has a physical or logical perimeter that must be secured against adversarial cyber activity. Additionally, the cybersecurity and resiliency posture of each system has a bearing on the overall security and safety of the aircraft platform.

11.2.1 Aircraft Boundary

The primary focus of the aviation ecosystem is the Aircraft boundary. That is illustrated in Figure 16. The aircraft is the reason this FSAD Guidebook exists.

Many offboard systems are essential for supporting aircraft operation. Almost every one of those interfaces uses some form of enterprise IT in its implementation. Attacks on systems at the Aircraft boundary could result in denying that subsystem, resulting in a loss of mission capability.

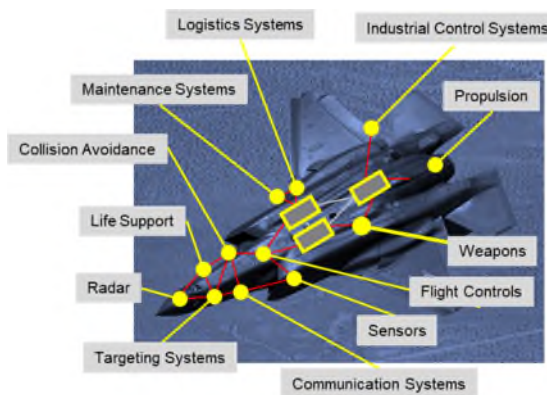


Figure 16. Aircraft Boundary

Successful adversarial cyber-attacks against offboard systems could ultimately cascade to the loss of critical aircraft functions.

11.2.2 Mission/Flight Planning Boundary

Flight planning systems are critical to the safe and efficient operation of aircraft. These systems are used to create flight plans that include information such as routes, altitude, speed, and fuel requirements. Flight planning systems typically consist of a variety of software applications and databases that are used to collect and analyze data from various sources. That includes things like weather forecasts, air traffic control, and aircraft performance data. These flight plans are then used to guide the aircraft during flight, ensuring that it follows a safe and efficient route.

Mission planning for military platforms augments standard flight planning with additional information to support combat operations. In addition to flight routes, altitude, speed, and fuel requirements, mission planning includes information about mission objectives, threats in the environment, tactical sensors, and weaponry. Mission planning systems used in military operations are typically more complex than those used in civilian aviation, as they must take into account a variety of factors including kinetic threats.

Mission planning systems used in military operations often include advanced analytics and modeling capabilities to help military personnel make informed decisions about mission objectives and tactics.

Flight planning and mission planning systems must be designed with security and cyber resiliency in mind, as they can be potential targets for cyber-attacks. Attackers can use these systems to gain access to sensitive information, such as flight routes and passenger data, or to disrupt aviation operations. Therefore, it is essential to implement robust security measures to protect flight planning systems from cyber-attacks. By ensuring the security of flight and mission planning systems, operational units can support rather than detract from the cyber resiliency posture of the aircraft.

11.2.3 Maintenance System Boundaries

Aircraft maintenance systems are critical to the operation of aircraft. These systems include a variety of software applications and databases that are used to perform onboard diagnostics, load software, download maintenance status, and update configuration settings on the aircraft.

Securing aircraft maintenance systems against adversarial cyber-attacks is essential to ensuring the safety and security of aircraft. Attackers can use these systems to disrupt maintenance activities or possibly even create hazards on the platform itself.

Securing maintenance system boundaries can help prevent unauthorized access and maintain the integrity of maintenance activities. Additionally, it is important to regularly monitor and audit maintenance systems to detect and respond to potential security threats.

11.2.4 Logistics Systems Boundary

Logistics systems are critical to ensuring the efficient and effective operation of aircraft. These systems include a variety of software applications and databases that are used to manage and track logistics activities. That includes supply chain management, inventory management, and maintenance actions.

Defending logistics system boundaries against adversarial cyber-attacks is essential to ensuring the efficiency and security of aircraft operations. Attackers can use these systems to gain access to sensitive information, such as supply chain data or transportation schedules. It is also possible for adversarial cyber-attack to disrupt logistics activities. This could potentially lead to operational delays, or even the grounding of aircraft.

By securing logistics system boundaries, organizations can help prevent unauthorized access to sensitive information and maintain the integrity of logistics support. It is important to regularly monitor and audit these systems to detect and respond to potential security threats. By ensuring the security of logistics systems, organizations can maintain the efficiency and effectiveness of aircraft operations by protecting against threats to these offboard systems.

11.2.5 Training Boundaries

Offboard training systems for aircraft are a type of equipment that allow pilots and other personnel to train and practice various scenarios and procedures in a controlled and safe environment, without the need for a physical aircraft. These systems typically use a mix of real and simulated aircraft software and hardware in a simulated operational environment. Training

systems can be used for a variety of training purposes, such as procedural training, emergency procedures training, and mission rehearsal.

It is important to secure the boundaries of offboard training systems against cyber-attacks. If a cyber-attack were to occur, it could potentially compromise the confidentiality, integrity, or availability of the training system, which could have an impact on pilot readiness and availability. For example, an attacker could potentially gain access to sensitive training data, manipulate the training scenarios to provide false or misleading information, or disrupt the training process entirely.

Some advanced aircraft training systems utilize actual aircraft software to provide realistic and effective training experiences for pilots and crew members. However, this integration of real aircraft software into training systems can potentially present a significant cybersecurity risk. These training systems, which are typically connected to a network, can be targeted by cyber attackers for aircraft reconnaissance.

Cyber attackers with access to a training system can gain valuable insight into the architecture, design, and even potential vulnerabilities in the actual aircraft. Furthermore, these training systems can also serve as a testing ground for potential attacks, allowing cyber-attackers to evaluate and refine their Tactics Techniques and Procedures (TTPs) against the real aircraft. As a result, it is essential for organizations to prioritize the cybersecurity of their aviation training systems to protect both training operations as well as the aircraft itself.

To secure the boundaries of offboard training systems, it is important to implement appropriate cybersecurity controls. It is also vital to regularly monitor and assess the security of the training system, and to promptly address any issues that are identified.

11.3 Summary and Additional Resources

Aviation system security is a complex and multifaceted issue that extends beyond just the aircraft itself. Maintenance devices, support systems, and communication interfaces all play a critical role in the safe and efficient operation of aircraft. Consequently, those offboard systems must also be engineered for cyber resiliency and also be secured against cyber-attacks and adversity.

The use of enterprise IT-based platforms in support and maintenance equipment can introduce new security challenges, as these systems can be potential targets for cyber-attacks and may provide attackers with multiple avenues for gaining access to the aircraft platform. To effectively defend against potential cyber-attacks, it is important to implement robust security measures for all air system components, including the aircraft, mission/flight planning systems, maintenance devices, and logistics systems. This includes following industry best practices for IT security.

Securing the complete air system requires a comprehensive and proactive approach to cybersecurity, encompassing all air system components and support systems. By prioritizing cybersecurity and implementing robust security measures, organizations can ensure the safety and security of their aircraft operations and protect against potential threats.

- **Weapons Systems Cybersecurity: DoD Just Beginning to Grapple with Scale of Vulnerabilities, United States Government Accountability Office, Report to the Committee on Armed Services, U.S. Senate, October 2018. (63)**

This report highlights the significant cybersecurity challenges faced by the Department of Defense (DoD) in protecting its weapon systems from increasingly sophisticated cyber threats. The report underscores the critical vulnerabilities in aircraft and interconnected ground support systems due to their extensive reliance on software and network connectivity. Despite the advanced capabilities provided by these interconnected systems, their complexity and the DoD's late start in prioritizing cybersecurity have made them particularly susceptible to cyber-attacks. These vulnerabilities can potentially compromise mission-critical operations, leading to catastrophic failures in both aerial and ground operations. The report emphasizes the necessity for the DoD to develop and implement robust cybersecurity measures to protect these essential systems, as the current efforts are just beginning to address the extensive scale of the challenges. This document is crucial for anyone involved in cyber risk assessment and defense planning, as it provides a detailed analysis of the existing vulnerabilities and the initial steps being taken to mitigate these risks.

Chapter 12

Assessment and Authorization of Aviation Systems

It is critically important that aviation systems are developed to be safe and secure. Consequently, all aircraft are subjected to assessment processes designed to ensure they meet those objectives. Independent evaluation and authorization are an essential part of understanding and minimizing the risk of accidents and cyber incidents. Increasingly, aircraft are subject to assessment and authorization that evaluates the cybersecurity and cyber resiliency posture of the platform.

Assessment and Authorization (A&A) is the current descriptive term for the overall process used to evaluate and approve operational deployment of DoD aircraft systems. While commercial aviation platforms are not subject to DoD standards unless specifically acquired for a military purpose, they are subject to conceptually similar security requirements and authorizations from other organizations, such as the Federal Aviation Administration (FAA).

Regardless of the specific standard or framework used, the fundamental principles of cybersecurity and cyber resiliency remain the same. Rather than myopically focusing on any particular A&A standard, it is recommended that aviation system designers use rigorous systems engineering and cyber resiliency best practices throughout the development lifecycle. That approach ensures that cybersecurity and cyber resiliency are fully considered for the platform. It reduces the risk that the bare minimum of what seems to be required by the relevant authorization standard is all that is implemented. When rigorous systems engineering is performed, along with cyber resiliency best practices, producing evidence of compliance with the A&A standard is relatively straight forward.

At the core, A&A is about ensuring that security controls are in place to protect systems from a variety of threats to mission performance and breaches of confidentiality. It is a systematic method for identifying and assessing risks, selecting and implementing appropriate mitigations, and continuously monitoring the security posture.

This chapter describes how the architects and designers of aviation systems can work effectively with A&A frameworks to set their platform up for success. That includes receiving an Authorization to Operate (ATO) but more importantly, also delivering intrinsically cyber resilient systems to the operational environment.

12.1 The Evolutionary and Interpretative Nature of A&A Standards

The DoD currently follows the National Institute of Standards and Technology (NIST) Risk Management Framework (RMF) for conducting A&A. (23) However, this FSAD Guidebook was intentionally developed to separate the timeless principles that provide the foundation of

cyber resiliency from whatever standard is currently in effect. Regardless of the precise assessment methodology used, the fundamentals remain the same. At the same time, it is important to appreciate how the A&A processes and standards have historically impacted DoD systems as well as how they have evolved and matured over time.

DoD cybersecurity authorization standards started with the inception of the “Rainbow Series” of books in the 1980s (64), which laid the groundwork for modern cybersecurity practices. As the industry understanding matured, so did the standards. In 1997, the DoD Information Technology Security Certification and Accreditation Process (DITSCAP) (65) was released and became the authorization standard for DoD systems. A decade later, DoD Information Assurance Certification and Accreditation Process (DIACAP) (66) came into effect.

The Risk Management Framework (RMF) (67) has been the prevailing standard since 2010. However, there have been a few iterations and revisions to the RMF security controls in the intervening time. While replacement of RMF does not appear to be imminent, it would be naive to fail to anticipate that further evolution in the standards is inevitable.

Therefore, it is crucial to prioritize cyber resiliency in the architecture and design of DoD systems, as the A&A standards will surely continue to be updated. By focusing on the timeless systems engineering principles, aviation system developers can build systems that are intrinsically cyber resilient. That naturally supports whatever is in effect for A&A at the moment.

12.2 The Enterprise IT Foundations of RMF

When applying RMF to embedded aviation systems, it is important to remember that the A&A standard is deeply rooted in an enterprise IT perspective. The cybersecurity risks for large, interconnected networks that were built for public access is fundamentally different from those of the embedded systems that are used on aviation platforms.

Applying RMF to aircraft presents unique challenges due to the distinct nature of those systems. Embedded aircraft operate in highly specialized cyber-physical environments with stringent performance and safety requirements. That includes limitations on computing resources and real-time operating constraints.

Consequently, implementing RMF in this context requires careful adaptation to accommodate the unique operational needs and technical considerations of embedded systems. Implementation of security controls must have minimal impact on system performance and reliability, yet still provide robust cybersecurity protection. A nuanced approach that balances the principles of RMF with the specific requirements of embedded aircraft systems is a necessity.

Domain expertise in the spirit and intent of each RMF control selected and allocated to the aircraft subsystems is essential when designing and developing cyber resilient aviation systems. However, that insight is not useful without a deep knowledge of how the aircraft is designed and implemented as well. Integrating the intrinsic cyber resilience characteristics that satisfy A&A controls requires understanding of the unique hardware, software, and network communications that are unique to these systems.

Many of the RMF controls emerged from enterprise IT networks where security add-on products are feasible and effective. The same paradigm does not work as well for on-board aviation systems where space, weight, and power are tightly constrained. A holistic approach and domain knowledge is required to create embedded aviation systems that are inherently resilient to cyber threats.

Therefore, when applying the RMF standard to aircraft, it is essential to consider these unique characteristics and requirements, and to tailor the security controls and assessment procedures accordingly. That most likely involves adapting the security controls to fit the specific needs of the system, and developing custom assessment procedures that are suited to the unique environment in which the platform operates. Taking a tailored approach is required when applying the RMF standard to aviation systems.

12.3 RMF and the Systems Engineering Lifecycle

NIST 800-37 expresses RMF A&A as seven distinct steps. (67) The cyclic relationship between each of these is typically illustrated as shown in Figure 17.

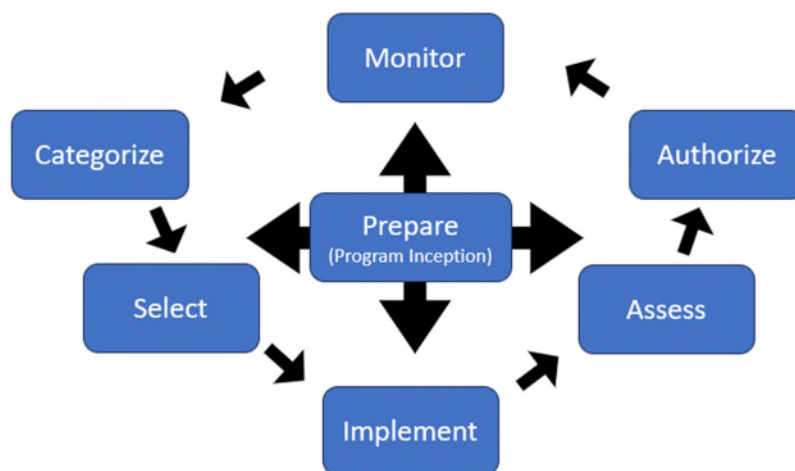


Figure 17. The Risk Management Framework Steps

The definitions for each step of the RMF process are summarized below. This list is the high-level roadmap of the phases of activities to support achieving an Authorization to Operate (ATO) in compliance with the standard.

- **Prepare** to execute RMF from an organization- and a system-level perspective by establishing a context and priorities for managing security and privacy risk.
- **Categorize** the system and the information processed, stored, and transmitted by the system based on an analysis of the impact of loss.
- **Select** an initial set of controls for the system and tailor as needed to reduce risk to an acceptable level based on an assessment.
- **Implement** the controls and describe how they are employed within the system and its environment of operation.

- **Assess** the controls to determine if they are implemented correctly, operating as intended, and producing the desired outcomes with respect to satisfying the security and privacy requirements.
- **Authorize** the system or common controls based on a determination of whether the risk to organizational operations, assets, individuals, and the Nation is acceptable.
- **Monitor** the system and the associated controls on an ongoing basis to include assessing effectiveness, documenting changes to the system and environment of operation, conducting risk assessments and impact analyses, and reporting the security and privacy posture of the system.

While Figure 17 is very similar to the way the steps are represented in NIST SP 800-37, a detailed comparison reveals that this FSAD guidebook depicts the loop in a slightly different rotational orientation. It is represented that way because each step of RMF relates directly to the phases of the Systems Engineering Lifecycle defined in Section 5.5.

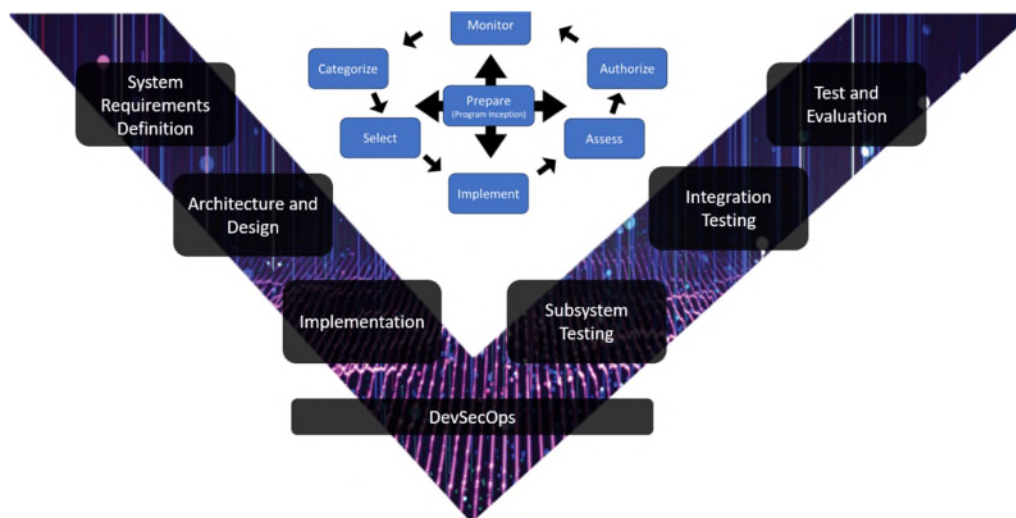


Figure 18. RMF and the Systems Engineering “V” Model

Figure 18 captures how RMF and the Systems Engineering Lifecycle can be thought of as different viewpoints of a common process. In fact, that is exactly what they are. The steps of RMF should be executed in conjunction with the corresponding phases of Systems Engineering.

Taking that mapping a step further, Table 5 matches each step identified in RMF with the corresponding Systems Engineering lifecycle phase along with where that activity is described in this FSAD Guidebook.

If an alternate A&A process is applied using an existing standard or one that emerges in the future, a similar mapping can be performed. Doing so will help program management, architects, and designers ensure that the requirements of that particular A&A methodology is adequately addressed by continuing to exercise the rigorous Systems Engineering process described in Table 5.

RMF Step	SE Lifecycle Phase	FSAD Description / Link
Prepare	Program Inception	The priorities for managing security risk and cyber resiliency objectives are established at program/project initiation. (FSAD Section 6.1)
Categorize	System Requirements Definition	System categorization is performed in accordance with the security and cyber resiliency objectives defined at program inception. (FSAD Section 6.4)
Select		Security controls are selected, allocated, and decomposed based on the system categorization. (FSAD Section 6.3)
Implement	Architecture and Design	Security controls are implemented within the architecture and design. (FSAD Chapter 7)
	Implementation	Security Controls are implemented within the system's hardware and software. (FSAD Chapter 14 and Chapter 13)
Assess	Test and Evaluation	Verification that the controls are implemented correctly, operating as intended, and producing the desired cyber resiliency posture on the platform. (FSAD Chapter 10)
Authorize		ATO issued for the system on the basis of successful security control implementation. (FSAD Chapter 12)
Monitor	Operational Use	Continuously monitor the system for emerging threats, risks, and security control effectiveness. Ongoing cyber risk assessment. (FSAD Section 7.5.2.1 and Chapter 9)

Table 5. RMF, Systems Engineering, and FSAD Mappings

12.4 The Pitfalls of A&A Standards

Program stakeholders must avoid the temptation to regard RMF or any other A&A standard as a minimum set of conditions that must be met to receive an ATO. Engineering for long term cyber resiliency requires analysis and activities above and beyond what is required to receive authorization.

Any system developed to implement only the bare minimum of A&A controls required to receive an ATO is unlikely to be cyber resilient. However, that doesn't mean that authorization standards are optional. An aircraft that does not take into account the RMF controls applicable to the platform is also not likely to be cyber resilient.

To achieve cyber resiliency, the aviation system must address the spirit and intent of RMF controls as well as rigorous security systems engineering. That holistic approach is required to design and build aircraft that are resilient in the face of highly dynamic cyber threats and risks.

12.5 Summary and Additional Resources

This FSAD Guidebook was intentionally decoupled from the prevailing A&A standards required to receive an ATO. That independence and separation is possible because the rigor and completeness of the security systems engineering is more impactful on the cyber resiliency of the aviation system than any enterprise IT oriented checklist could provide. Simply following templates to produce the artifacts required to support security authorization should never be confused with performing the engineering.

At the same time, there is an immutable relationship between the A&A process and the security systems engineering that makes authorization possible. When the engineering is rigorously performed, the information necessary to efficiently and effectively complete security authorization documentation is readily available. When it is not, A&A is more challenging as the documentation tries to retroactively describe the engineering that should have been done.

Producing security authorization documentation for a system is the culmination of the cybersecurity engineering performed during development. That is supported by a deep understanding of the current A&A standard beyond what is included in this FSAD Guidebook. Stakeholders must take steps to become educated and informed on the process and controls within the prevailing authorization process.

- **Risk Management Framework for Information Systems and Organizations: A System Life Cycle Approach for Security and Privacy, National Institute of Standards and Technology (NIST), Special Publication SP 800-37, Rev 2, December 2018.**

This publication describes the Risk Management Framework (RMF) and provides guidelines for applying the RMF to information systems and organizations. It describes the process for managing security and privacy risk that includes information security categorization, control selection, implementation, assessment, authorization, and continuous monitoring.

Part 3 – Cyber Resiliency Specialty Domains



The first F-117 during final assembly at the Lockheed Skunk Works facility in Burbank, California, circa 1980.

Chapter 13

Software Assurance

This chapter was written by Teresa Merklin with significant contributions from Eugene Moore, Trey Jones, Sachin Kamath, and Jeff Langham.

Software Assurance is the discipline of leveraging people, process, and technology to ensure that software products achieve a high standard of reliability and security. For aviation systems, the requisite degree of cyber resilience is significantly higher than in many other industrial domains. Defective or vulnerable software can lead to extensive consequences, imperil mission success, and place human life at risk.

Another definition of software assurance is “the level of confidence that software functions as intended and is free of known vulnerabilities, either intentionally or unintentionally designed or inserted as part of the software, throughout the life cycle.” (29) Developing secure systems requires the identification of potential vulnerabilities or weaknesses throughout the entirety of the system development lifecycle. Consequently, it isn’t enough to apply the best practices outlined in this chapter and expect cyber resiliency. That intrinsic attribute must also be considered as the requirements, architecture, and design are developed. No matter how well software is written, vulnerabilities that are inherent in the architecture cannot be overcome or avoided.

This chapter presents the software security development best practices and mechanisms that support the characteristic of intrinsic cyber resilience that is required for all aviation systems. It also describes the methods, tools, and techniques used during software security testing and evaluation. The ultimate goal is to create inherently secure software that minimizes the need for supplemental security controls or other mitigations. That only happens when the software development process intentionally focuses on cyber resiliency throughout the development and evaluation process.

13.1 Fundamentals of Secure Software Development

It is critically important to establish and follow good fundamental principles of secure software development for aircraft and aviation systems. Successful cyber-attacks or coincidental failures could result in catastrophic risks to safety and mission performance. Practices to ensure that software is resistant to cyber adversity minimizes the likelihood of errors and vulnerabilities that could be catastrophic to aircraft. Consequently, secure development practices are essential to maintaining the integrity, reliability, and trustworthiness of aviation systems.

Secure software development requires achievement of four challenging objectives: (68)

- Ensure that people, processes, and technology are equipped to perform secure software development.
- Protect all aspects of software from tampering and unauthorized access including development environments, distribution, and operational use.
- Develop and release well-secured software with minimal security vulnerabilities.
- Should latent software vulnerabilities be identified, respond appropriately by promptly addressing them and take steps to prevent similar ones from occurring in the future.

The Secure Software Development Framework (SSDF) (68) is a comprehensive guide written to inform software security practices. That resource outlines a robust framework of high-level secure software development practices that can be integrated into any software development life cycle (SDLC). There are many valid approaches that can result in cyber resilient software. Consequently, the SSDF is highly descriptive rather than prescriptive.

Rather than mandating any specific tools, techniques, or methodologies for implementation it focuses on the outcomes of secure development practices. Organizations have the flexibility to adopt and adapt processes and tools to their unique environments. As a result, the SSDF is broadly applicable across many industry sectors including aviation systems.

Software development organizations for aviation systems must have a highly detailed Software Development Plan (SDP). It should specify the structured processes that are rigorous, repeatable, and aligned with the best practices outlined in the SSDF as well as the aircraft specific considerations identified in this FSAD guidebook. That is essential for systematically identifying and mitigating potential risks and vulnerabilities, which is crucial for aviation platforms.

Incorporating these best practices ensures that each phase of software development adheres to established standards for security and quality. Repeatable processes allow for better traceability and accountability, making it easier to demonstrate compliance with cybersecurity requirements and objectives. A comprehensive approach not only enhances the security and functionality of aviation systems but also builds trust and confidence in the platforms.

13.2 People in Software Assurance

People play a vital role in the development of cyber resilient software. While that is true for systems in all industries, it is particularly important for aviation systems where specific domain expertise, experience, and judgement is necessary. Skilled software developers, cybersecurity specialists, and systems engineers who understand the unique processing requirements and constraints of aviation platforms are essential for defending aircraft systems against sophisticated threats.

This section describes the specific ways in which people are the cornerstone of developing cyber resilient software for aviation systems.

13.2.1 Secure coding mindset/culture

Secure coding practices require a mindset and culture that prioritizes cyber resiliency throughout the implementation process. Regardless of how the software is produced, the developers must be acutely aware of how their code can impact the cyber resiliency posture of the aviation platform. This mindset requires adopting a "security-first" approach.

Creating a culture that supports secure coding requires a shift in mindset from reactive to proactive thinking. Instead of waiting for security issues to arise and then fixing them, developers must actively seek out potential threat vectors and vulnerabilities and address them before they are deployed operationally and can be exploited.

The people responsible for writing the software required to make high-performance cyber-enabled aviation systems must internalize the core concepts that were outlined earlier in this FSAD guidebook. That includes understanding that aircraft platforms face the clear and present cyber threats that were introduced in Chapter 2.

Additionally, the people who develop software for aviation systems must embrace the idea that cybersecurity is everybody's responsibility and live out the core behaviors identified in Chapter 3. As a bonus responsibility, software developers also play a key role in creating and maintaining the security of their software development environments in addition to the software products that are released through those systems. The fundamentals of securing the infrastructure were described in Chapter 4 of this FSAD guidebook.

Prioritizing the creation of secure, vulnerability-free code, lays a solid foundation for cyber resilient software. The overarching goal is to produce software to **prevent** cyber incidents and attacks. However, cyber resilient systems will also **mitigate** the effects of cyber incursions and **recover** quickly. Cyber resilient software will also be structured to **adapt** to prevent future attacks. In other words, the people will develop software that fulfills the four pillars of cyber survivability defined in Chapter 5.

People with the mindset who support the culture required for secure coding will be familiar with the foundational knowledge in this FSAD guidebook and leverage those principles throughout the software development process.

13.2.2 Avoidance of Known Software Issues

When developing secure software, it is important to avoid known and obvious weaknesses in the code. One aid to help recognize potential issues is captured in the Common Weakness Enumeration (CWE) list. (55) The CWE list provides a standardized taxonomy for describing types of software security issues. Similarly, another resource that is useful is the Common Vulnerabilities and Exposures (CVE) database which enumerates publicly disclosed cybersecurity vulnerabilities in software products. (56) CVEs are an itemized list of vulnerabilities and risks associated with software products that are used or integrated within the aviation system.

However, it is never enough to only focus on known issues such as those identified in public lists such as the CWE and CVE. Exclusively focusing on vulnerabilities within application code and supporting software products will overlook risks and issues in the full aviation stack. Paradoxically, it is also not safe to assume that any CWEs or CVEs identified within the code are protected by another layer.

The cybersecurity software industry is filled with vendors selling tools that can automatically scan for CWEs and CVEs. In fact, using automated scanning is a best practice integrated throughout this chapter of the FSAD guidebook. While it is important to leverage those tools to the greatest extent possible, it is essential to be aware of limitations and pitfalls of over emphasizing their value.

Vulnerability databases such as the CWE and CVE are heavily influenced by enterprise IT systems. Consequently, the unique operating systems and structure of aviation software is quite likely underrepresented within those resources. The absence of a heuristic pattern makes detection unlikely. It is important to understand that while these scanning tools should be used, they are not as effective for aircraft as for enterprise IT platforms.

Additionally, even software that does not have recognizable CWE and CVEs cannot be assumed to be secure. In fact, it is possible to write malicious code that is free of vulnerabilities. Similarly, automated scanners that check for this type of issue will not detect cybersecurity problems that are inherent to the architecture of the system. Nor will they detect issues elsewhere in the full aviation stack that are not a part of the automated scanning.

Software developers must be familiar with items included in the CWE, CVE, and all other vulnerabilities relevant to their system. That is so they can avoid creating software that replicates those known issues. Additionally, scanning tools should be leveraged to eliminate as many of these obvious errors as possible.

However, it should never be assumed or believed that scanning is sufficient to eliminate all issues with aviation software.

13.2.3 Knowledge, Skills, and Attributes

Having people with the right knowledge, skills, and attributes is essential for software development, especially in high-stakes fields like aviation systems. Human capital is essential for creating secure, reliable, high-performance software. While leveraging technology through digital transformation initiatives and the use of Artificial Intelligence (AI) can greatly enhance the productivity of humans, there is still a need for highly skilled people.

Advanced software development tools can streamline and enhance software development efforts. However, people are still required to use the tools to create high quality architectures, designs, and code. A useful way to conceptualize this is by considering the tools of carpentry. While anyone can purchase the tools required to create intricate woodwork, a high degree of skill and technical competence is required to actually do that. Possessing and using the best tools alone does not guarantee excellence. It is the combination of high-quality tools combined with human expertise, creativity, and attention to detail that results in exceptional craftsmanship.

A high degree of software development expertise as well as strong technical knowledge of the aerospace domain is essential for writing code for high-performance, safe, and cyber resilient aviation systems. First and foremost, software developers must know how to produce high-quality software and solve complex problems effectively. They must also have an in-depth understanding of the programming languages used by the aircraft subsystem they are working on. It is also useful to be well-versed in various software development lifecycle models and methodologies, but particularly the one being used on their projects.

Creating cyber resilient software necessarily includes writing code that can be adapted or maintained as the capabilities of the platform are expanded and in response to evolving threats of cyber-contested airspaces. That requires a strong foundation in data structures, algorithms, design patterns, and best practices for writing clean, maintainable code.

It is also essential that software development personnel possess an acute awareness of cybersecurity principles and can use that knowledge to create cyber resilient software that is free from known vulnerability patterns. They should have specific knowledge of the risks and threats facing aviation systems to achieve real-time performance and resource constraints of the special purpose cyber-physical systems of aircraft.

One of the most essential skills of software developers is excellence in problem-solving by efficiently breaking down complex systems into software solutions. Great developers can write, test, and debug code with high proficiency, manage code changes using configuration management tools, and ensure software quality through rigorous unit and integration testing.

Great software developers pay attention to detail, which drives a meticulous approach to coding and testing. They exhibit the perseverance and determination that helps them tackle challenging problems and develop innovative solutions.

Finally, the best software developers are excellent communicators who embrace a collaborative mindset that fosters teamwork. They freely support and share knowledge with their peers, and have a strong sense of individual accountability. Effective communication skills enable them to collaborate with team members and stakeholders, while adaptability allows them to quickly learn and implement new technologies, tools, and methodologies.

When developing software for aviation systems, there is little room for deficiencies in any of these key knowledge, skill, and attribute areas. The combined qualities make an excellent software developer capable of delivering reliable, efficient, and secure software solutions. That is essential for safe and cyber resilient aircraft.

13.2.4 Training the Software Development Role

People are an indispensable part of developing cyber resilient software. The requisite knowledge, skills, and attributes of highly capable software developers outlined in 13.2.3 establish high standards that are particularly important for the workforce that develops aviation software. Organizations must implement robust training programs and establish a culture of continuous learning to meet their staffing needs.

The software development training curriculum will vary widely based on several factors. That includes the particular knowledge needed for the aviation system or subsystem under development. It must also take into account the programming languages, operating systems, and software development tools used by the project. Additionally, developers must be well versed on the development methodology used within the organization.

This customized approach ensures that the training is relevant and applicable to the developers' day-to-day work, enabling them to understand and effectively apply secure coding practices appropriate for aircraft platforms. Aligning the training curriculum with the organization and project-specific context is essential for writing secure code that contributes to enhanced cyber resilience in the aviation domain.

General programming knowledge suffices for developing software that meets functional requirements, but writing software that is free of vulnerabilities requires experience and knowledge of common pitfalls so they can be avoided.

Training highly capable software developers is an ongoing process of specialized training. As the threat landscape and new vulnerabilities are identified, updates to training curriculum becomes essential. This continuous learning process is vital to ensure that the software development workforce is up to the challenge. Additionally, turnover in the software development workforce means that the organization must be able to quickly bring new developers up to speed.

13.3 Process in Software Assurance

There are many valid methods for creating software. Human developers can write code manually using simple text editors. Code can also be created by leveraging integrated development environments (IDEs) which greatly streamline the production process. Some Model Based Systems Engineering (MBSE) environments may support code creation through human manipulation of visual graphical interfaces.

Regardless of how it is developed, the fundamental way to achieve cyber resilient software is to synthesize code that does not have inherent vulnerabilities or weaknesses in the first place. Unfortunately, that is much easier said than done. This section enumerates the process-based best practices that support that pursuit.

Part 2 of this FSAD guidebook emphasized the importance of integrating cybersecurity and resiliency throughout the entirety of the systems engineering lifecycle. That concept includes the software development lifecycle as well. Organizations committed to achieving cyber resiliency in their software will have processes that integrate those concepts into everything they do, rather than regarding it as an afterthought.

13.3.1 Software Development Plan (SDP)

A Software Development Plan (SDP) is essential for creating cyber-resilient software. The SDP establishes a written and repeatable process that integrates best practices, standards, and methodologies across every phase of the software development lifecycle. Having an SDP is the

hallmark of well-planned and managed software development projects. That foundation is essential to support cybersecurity.

Mature SDPs include the project's objectives, scope, and goals. It includes a process for eliciting, managing, tracing, and verifying both functional and non-functional requirements of the system under development. The SDP also describes how the effort is planned, tracked, and resourced. The project's objectives, requirements, and planning identify the specific ways it will assess and manage cybersecurity.

Risk management is a critical part of the SDP as it describes how risks are identified, along with the way impact and mitigation approaches are determined. The SDP includes recurring assessments of both project risk and the threat landscape. This is crucial as the code is developed and matured as the implementation may impact the aircraft's attack surface and cybersecurity posture.

Software testing encompasses a set of essential processes that ensure the quality, functionality, and security of the developed system. The SDP describes how testing is performed at the unit, integration, and system levels. It will also define the configuration management mechanisms as well as how defects are recorded, tracked, and managed. The SDP will also specify that regression testing is regularly performed to ensure that new features and functionality have not introduced any new security issues and concerns.

A mature SDP will identify the development artifacts that are created and assessed in addition to the code itself. Comprehensive specifications, user manuals, and technical guides will be created and reviewed to ensure that all aspects of the project are well-documented and understood.

Regardless of the specific methodologies used to produce software, an SDP is a necessity. It describes how cybersecurity best practices are integrated throughout the software development lifecycle, ensuring that cyber resiliency is a fundamental component of everything that is done. Regardless of the programming language, development methodology, or specific tools employed, the SDP encompasses the comprehensive approach to ensuring that the software is resilient and secure.

13.3.2 Software and Security Policies

System security policies applicable to the aviation system should be implemented when considering security mechanisms and controls. A software security policy is a set of guidelines, rules, and procedures that governs the development, deployment, and maintenance of software systems to ensure their security and protection against potential threats. It outlines the responsibilities of various stakeholders. In an aviation system that includes the software developers, aircraft operators, maintenance personnel, pilots and support crew.

Software security policies necessarily include guidelines on secure coding practices and the overarching security requirements for both the system and organization. It is a high-level statement of management intent that answers the questions of "what" and "why" without delving into the specifics of "how." Software security policies are strategic and independent of the underlying implementation within the system.

A Security Requirements Guide (SRG) is a set of instructions and standards used to define the security requirements for systems. These guides are developed to ensure that systems are designed, built, and operated in a manner that meets specific security requirements to protect against threats. SRGs provide a structured approach to implementing security controls and measures, ensuring that systems meet minimum cybersecurity requirements.

DoDI 8500.01 tasks the Defense Information Systems Agency (DISA) with developing SRGs that are consistent with DoD cybersecurity policies, standards, architectures, and security controls. An SRG is a compilation of Control Correlation Identifiers (CCIs) which are decomposed NIST controls into single, actionable, measurable statements. DISA is also tasked with creating and maintaining the CCIs. (69)

Chapter 12 cautioned system developers against succumbing to the temptation to regard the Risk Management Framework (RMF) or any other A&A standard as a minimum set of conditions that must be met to receive an Authorization to Operate (ATO). That chapter described how engineering for long term cyber resiliency requires action above and beyond what is required to receive authorization for a system.

For DoD aviation systems, it is likely that aircraft platforms will have a set of CCIs levied against the software as requirements. If those are the only controls implemented, the platform is unlikely to be cyber resilient. However, aircraft software that does not implement those same CCIs on the platform is also not likely to be cyber resilient. Once again, a holistic approach is required to design and build aircraft that are resilient in the face of highly dynamic cyber threats and risks.

13.3.3 Secure Coding Standards and Guides

Secure coding standards and guides provide a structured framework for developers to follow as code is written or generated for a project. Adherence to a rigorous standard is a vital part of the software development of secure systems. Secure coding standards and guides help avoid some types of common vulnerabilities such as buffer overflows and other security flaws that are created because of sloppy programming or errors that are explicitly prohibited by the standard.

Adhering to secure coding practices allows developers to recognize and avoid potential risks early in the development lifecycle. That reduces the likelihood that security vulnerabilities or issues are introduced through the software coding process. Secure coding standards provide the foundation of the overall security posture of the software and fosters a culture of security awareness as code is written and evaluated.

When well-written, secure coding standards and guides serve as a valuable educational resource for developers, especially those who may not have extensive experience in cybersecurity. It provides concrete examples and best practices for writing secure code and helps developers understand the implications of their coding decisions. It is an effective knowledge transfer mechanism crucial for maintaining a consistently high level of security across the codebase.

Compliance with secure coding standards and guidelines reduces style variations that frequently exist from person to person resulting in more consistent and maintainable code. That is

particularly important in larger projects where multiple developers may be independently working on different parts of the program.

Establishing and following secure coding standards and guides is a demonstrated commitment to cyber resiliency. It is fundamental to ensuring that security is built into the entire development lifecycle.

13.3.4 Rigorous Code Reviews

Rigorous code reviews are important because they play a critical role for ensuring the quality, security, and maintainability of software. Thoroughly examining the source code may reveal potential vulnerabilities and bugs that can be resolved as the software is initially developed. This early detection is crucial for eliminating vulnerabilities and errors that can lead to cybersecurity concerns. Code reviews help to enforce coding standards and best practices, which enhances the maintainability, reliability and cybersecurity of the software.

Rigorous code reviews are an effective mechanism for knowledge sharing and skill development within the software development team. It is an effective way for experienced coders to provide valuable feedback and insights to more junior team members which can help them learn and improve their software development skills. When less experienced developers review the code of more senior team members, they can also learn from the process and perhaps even share new innovative techniques.

Paradoxically, there is some evidence that performing rigorous code reviews can reduce the overall cost of system development. Catching coding errors and issues early eliminates the need to detect and correct deficiencies later in the development process. It is a proactive approach that enhances both cybersecurity and overall software quality.

Rigorous code reviews can be supported by both manual and automated methods. In fact, both approaches should be used to ensure comprehensive evaluation of the software's quality and security. In manual code reviews, software developers meticulously examine the code line by line, to identify vulnerabilities, logic errors, deviations from coding standards, and anything else that will negatively impact the performance of the system. Manual reviews are crucial for identifying complex issues that automated tools might miss, such as system logic flaws and nuanced security vulnerabilities.

Automated code reviews complement manual efforts by using specialized tools to quickly and efficiently analyze the source code for common security vulnerabilities and code quality issues. These tools, such as static code analyzers, scan the code to detect patterns that may indicate potential problems, like insecure API usage, deprecated functions, or compliance violations. Automated tools should be run continuously as part of the software development pipeline. The combination of manual and automated reviews creates a robust code review process, leveraging the strengths of both approaches to enhance the overall security, quality, and maintainability of the software.

13.3.5 Software Security Controls and Cybersecurity Models

Software plays a vital role for achieving the properties of Confidentiality, Integrity, and Availability (C-I-A) that were introduced in section 5.2. Consequently, it is beneficial to revisit the fundamental concepts of the C-I-A triad with an expanded look at how software impacts each aspect of this cybersecurity model.

Confidentiality: Sensitive information is not disclosed to unauthorized people or systems.

In aviation systems, confidentiality might be implemented through a combination of robust encryption techniques, strict access controls, and secure communication protocols. It protects sensitive data, such as flight and mission plans, and should not be disclosed to unauthorized individuals or systems either in transit or at rest.

Integrity: Assurance that information has not been improperly altered or destroyed.

Integrity should be implemented to verify that software and data have not been altered during transmission or storage. For aviation systems, substantiated integrity methods are preferred over non-substantiated approaches as they can detect corruption by both accidental events as well as malicious cyber-attack. In aircraft, ensuring that messages and commands come from legitimate and trusted sources are essential for flight critical functions such as navigation and air traffic control. Software should also implement rigorous validation processes, including input validation and error detection algorithms.

Availability: Timely and reliable access to information by authorized people or systems.

Aviation software should implement availability through redundancy, failover mechanisms, load balancing, and continuous monitoring. On an aircraft, redundancy techniques such as providing multiple instances of critical components should be considered. That approach ensures that if one component fails, another can take over with minimal disruption. Highly robust software is necessary to implement those failover mechanisms quickly and seamlessly. High availability is essential for flight and safety critical aviation capabilities.

The C-I-A triad is a foundational model necessary for understanding cyber resiliency and survivability of aviation systems. Additionally, the model is frequently invoked to express cybersecurity core concepts in software across many other domains. However, it should be used with caution for software within aviation systems. That is because it has historically been slanted to focus on risks to confidentiality, integrity, and availability outside the mission context.

Consequently, it is recommended that software assurance engineering efforts also consider the aviation system critical mission framework model that was previously introduced in section 6.5 and repeated here as Figure 19. In this diagram, the highest priority missions are at the foundation of the pyramid, and priorities of the other mission essential functions build from there.

In an emergency, the highest priority for the pilot is to aviate, the next most important thing is to navigate, and the third priority is to communicate. That is a good starting point for prioritization of software functionality for an aircraft system.

For example, software that implements functions for flight control would be the highest priority mission. The property of *availability* is typically more important than integrity, though that aspect is also very important in that context. *Confidentiality* is usually less impactful than either of the other two properties. This thought process is critical for understanding how to consider threats against software and to defend mission execution in a way that makes sense for aircraft systems.

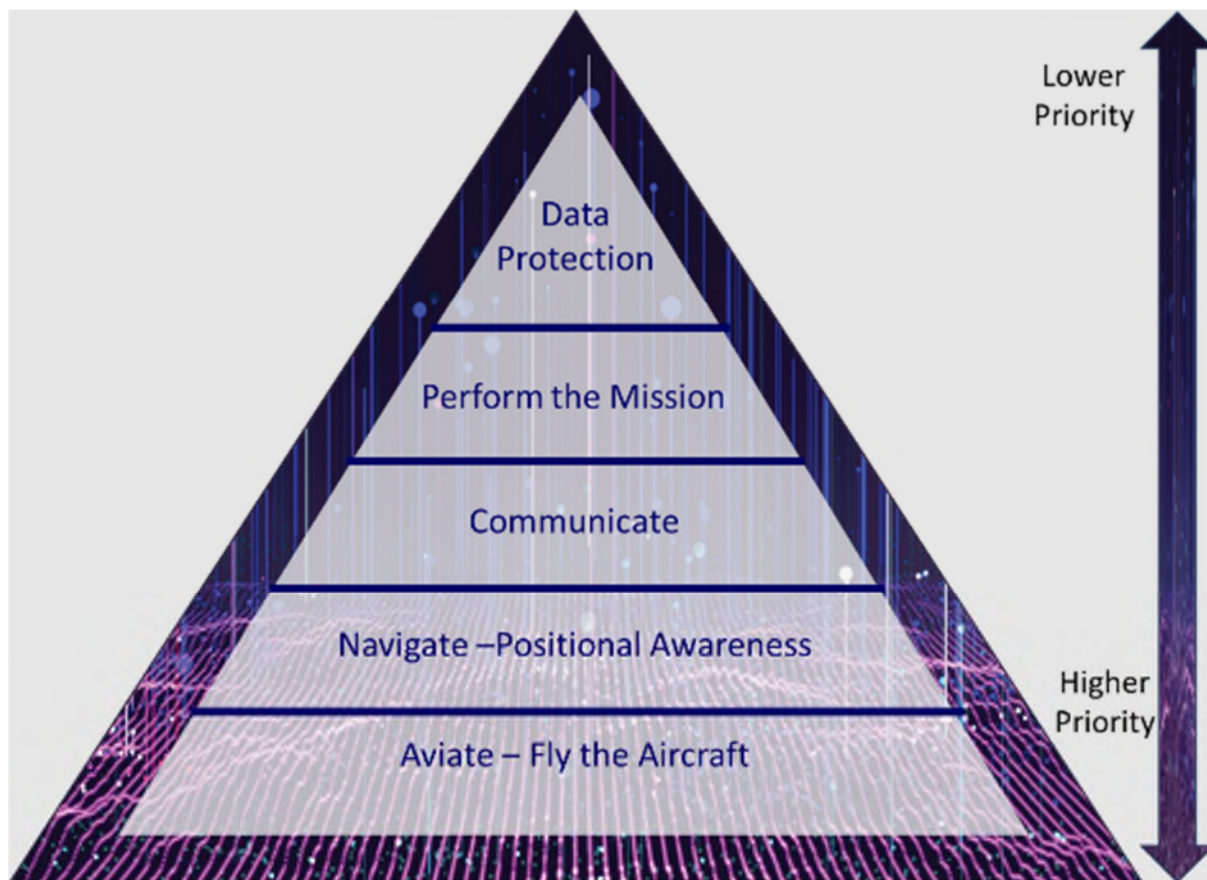


Figure 19. Aviation Platform Critical Mission Framework

A security control is a safeguard used to avoid, detect, counteract, or minimize risks to physical property, information, cyber-enabled systems, or other assets. Ideally, security controls are derived from security policies that provide the overarching principles and guidelines curated for a specific platform. Security controls should flow from clear and concise system requirements.

Software security controls should not be chosen arbitrarily such as blindly selecting and allocating RMF derived CCIs based on a generic checklist. Rather, security controls implemented in code should flow from a system's resiliency requirements and reflect architectural and design decisions.

13.4 Technology in Software Assurance

While the human element and rigorous processes are essential for developing cyber resilient aviation systems, leveraging technology is equally as important. Technology enhances the efficiency and accuracy of assurance processes required to develop secure and reliable software products.

Advanced tools and technologies can automate repetitive and time-consuming tasks such as static code analysis, dynamic testing, and vulnerability scanning. That can speed up the development process and also ensure that tasks are performed consistently and without human error. A project that does not maximize the use of available technology is not performing efficiently.

However, the limitations of technology must be understood. Automated tools are designed to detect known vulnerabilities and patterns of insecure coding practices. However, these tools are less effective against vulnerabilities and classes of software weaknesses that have not previously been identified or catalogued. Technology enhances rather than serves as a substitute for human judgement.

Leveraging technology for software assurance is vital because it enhances efficiency and accuracy for identifying certain types of vulnerabilities and security concerns. It is best used in conjunction with people and process. This technological integration ultimately leads to the development of more secure, reliable, and high-quality software products.

13.4.1 Static Code Analysis / Static Application Security Testing

Static code analysis and static application security testing (SAST) are a software development technique that analyzes source code without execution. This method can be performed using manual methods or automated tools. The rigorous code reviews described in section 13.3.4 are a manual method that uses human labor to examine the source code. Automated tools can also be used to scan source code to identify potential vulnerabilities, coding standard violations, and other quality issues.

When developing software for aviation systems, it is essential to use both manual and automated methods. While automated tools support consistency, efficiency, and accuracy, manual methods are necessary to provide human expertise. Manual methods can help identify unique and complex issues that automated methods may miss, while automated methods can help reduce human error and ensure repeatability.

It is important to emphasize that the majority of the automated tools used for software testing and analysis are based on vulnerabilities and information derived from enterprise IT systems. That can diminish the effectiveness for identifying specific concerns for the embedded cyber-physical systems used in aircraft. The unique characteristics of aviation systems include real-time processing performance and unique hardware and software architectures, which are not typically represented in enterprise IT systems.

As a result, automated tools may not be able to identify all potential vulnerabilities and risks in aviation systems. Manual methods, such as code reviews, threat modeling, and expert analysis, is essential in this domain. Combining both automated and manual methods is the best approach for aircraft software assurance.

One of the most challenging aspects of static code analysis is the inevitability of false positives and false negatives. False positives occur when the analysis tool incorrectly identifies a piece of code as a vulnerability or defect. On the other hand, false negatives arise when the analysis tool fails to identify real issues, leading to potential security risks or system failures. Both false positives and false negatives can have significant consequences for the development process.

It is never sufficient to simply run an automatic static code analysis tool. Any findings detected must be investigated and dispositioned, which takes considerable time and resources. It requires a person with a detailed understanding of the code to understand and address items flagged by the tool. That means that skilled and experienced developers must be involved. Aviation software development organizations should plan for the time and resources required for investigation and dispositioning static code analysis as part of their software development process.

13.4.2 Dynamic Code Analysis / Dynamic Application Security Testing

Dynamic code analysis is a software testing technique that involves executing the code in a controlled environment to identify defects, vulnerabilities, and other issues. Unlike static code analysis, which examines the software without executing it, dynamic code analysis executes the code and observes its behavior to identify issues that may not be apparent from the code alone. Dynamic code analysis can identify concerns that may be missed by static code analysis, such as race conditions, memory leaks, concurrency issues, and performance challenges.

Dynamic application security testing (DAST) is a type of dynamic code analysis that focuses specifically on identifying security vulnerabilities in the application. DAST tools simulate real-world attacks on the software to identify security vulnerabilities. DAST tools can also identify issues related to secure configuration.

One popular approach to dynamic code analysis is fuzz testing. Fuzzing involves sending random, unexpected, or malformed input to an application to identify input validation issues and other security vulnerabilities. Fuzz testing can be used to evaluate the software's robustness and identify weaknesses. It is good at identifying issues such as buffer overflows, heap overflows, and memory leaks, as well as input validation issues.

Fuzz testing can be performed by either automated or manual methods. It is an area of specialty testing that necessarily must be customized to fit the software under evaluation. While fuzzing is a powerful approach for identifying issues that may be missed by other testing techniques, it is extremely labor intensive, and requires skilled and experienced developers with a deep understanding of the software.

It is worth noting that dynamic code analysis and DAST techniques should be used in conjunction with static code analysis and SAST techniques to provide a comprehensive

assessment of the software assurance required for aviation systems. By using both static and dynamic analysis techniques, aviation system developers can identify a broader range of issues and ensure that their software is secure, reliable, and of high quality.

13.4.3 Automated Vulnerability Scanning

Automated vulnerability scanning uses specialized tools to identify security weaknesses in software and systems. These tools automatically scan for known vulnerabilities and configuration issues. A detailed report typically summarizes any potential security problems and sometimes also provides mitigation recommendations.

Scanning software and systems for vulnerabilities are two distinct processes tailored to different aspects of cybersecurity. Scanning software involves analyzing the application's code, either statically or dynamically, to detect security flaws, coding errors, and potential points of exploitation within the software itself.

On the other hand, scanning a system for vulnerabilities focuses on assessing the software within an operationally representative environment. System scanning typically identifies weaknesses such as unpatched software, misconfigurations, weak passwords, and open ports that could be exploited by attackers.

Automated vulnerability scanning is a good idea for aviation systems because it can quickly identify obvious vulnerabilities and issues. However, there are some pitfalls and limitations to consider. Automated tools will sometimes struggle with the complexity of aviation systems and proprietary and legacy software that automated tools aren't designed to handle. In general, while automated scanning is effective at identifying known vulnerabilities, it may not detect new or sophisticated threats that require more advanced analysis.

13.4.4 Security-Focused Software Testing

Security focused testing is a critical component of software assurance in aviation systems which must be cyber resilient against security threats and system failures. Taking a security focused perspective is vital during unit, integration, and system testing. While it is essential to verify and validate that functions and capabilities implemented in systems are working correctly, it is equally important to perform testing that focuses on what can go wrong.

At the unit level, security focused testing can be used to identify potential threats in individual software components such as input validation issues and error handling. During integration, tests should be used to identify potential problems in the interactions between software components. During system testing, security focused evaluations should be performed to identify potential vulnerabilities in the overall aviation platform.

When evaluating aviation systems, it is important to test functional capabilities, performance, and system stability. However, it is equally important for test plans to include negative testing which is a key aspect of cyber-focused testing. It is also important to consider scenarios based on the use and abuse cases described in section 6.6.

Just like so many of the other best practices recommended in this chapter, security focused testing should be performed by skilled and experienced developers who have a deep understanding of the software as well as the aviation domain.

13.4.5 Automated Security Testing

Automated testing is somewhat of an overloaded concept in cybersecurity. At a fundamental level, automated testing encompasses a wide range of tools, techniques, and processes to identify and address security vulnerabilities. Unfortunately, that breadth can lead to potential confusion over what automated testing actually entails.

For example, some sources categorize the static and dynamic code analysis techniques described in sections 13.4.1 and 13.4.2 as testing. In fact, the “T” in the acronyms SAST and DAST stands for exactly that. Alternatively, some restrict the use of the word testing to verification and validation of the functional requirements. Scanning for weaknesses and vulnerabilities is not a substitute for requirements verification through functional testing.

Functional testing can be automated using software tools and techniques to create, execute, and analyze test cases based on the functional requirements for software. It is a type of testing that requires considerably more investment than automated scanning. However, it is ideal for pipeline implementation as this technique automatically evaluates the ability of the system to meet its requirements. That is essential for mission performance.

Additionally, automated functional testing supports continuous regression testing in the development pipeline. That allows requirements that have previously been verified through a test case to be re-evaluated with every new software update to ensure that new functionality didn’t break existing capabilities. Automated regression testing is particularly important in agile development environments, where frequent code changes and releases are the norm.

Every aviation system should have security functional requirements. As described in section 6.3, those can be a part of the security requirement baseline or tracked separately as security controls. Each of those requirements should trace to one or more verification test, many of which are candidates for automation. In fact, it is important to automate security testing to the greatest extent possible to ensure that new capabilities do not introduce new vulnerabilities or break previously verified security functionality.

While automated scanning is an important aspect of cybersecurity evaluation, it is not sufficient on its own to ensure the security and reliability of software for aviation platforms. Testing to verify requirements and capabilities is essential.

13.5 Full Stack Cybersecurity for Aviation Systems

The full stack execution environment for aircraft software requires a complex interplay of platform and supporting infrastructure. That includes hardware, operating systems, middleware, and application layers. Cyber resilient software requires security best practices and mechanisms at all layers of the technology stack.

At the **hardware layer**, aviation system software runs on specialized, robust, and high-performance cyber-enabled systems designed to withstand harsh environmental conditions and cyber adversity. Chapter 14 details more information on developing high-assurance hardware systems for aviation platforms.

The software execution environment also executes on infrastructure traditionally thought of as the Information Technology (IT) stack. It is an enterprise IT concept that can be tailored for aviation systems as summarized in Table 6.

Stack Aspect	Enterprise IT Viewpoint	Aviation Systems Viewpoint
Network Security	Securing devices within the same network including both software and hardware vulnerabilities.	The hardware and software deployed to the aircraft must be intrinsically cyber resilient with no exposed vulnerabilities. Additionally, safeguarding the hardware and software to avoid risks to and from other onboard systems.
End-Point Security	Ensuring the security of devices (laptops, phones, etc.) to prevent unauthorized access.	Security protection from threats and risks from systems that connect to the aircraft including support, maintenance equipment, and pilot aids.
Internet Security	Protecting data and information in transit.	Data transmitted and received from external systems meet availability, integrity, and confidentiality requirements.
Cloud Security	Mitigating software security risks within cloud environments which comprise many of the environments where software is developed and operated.	Software and data received from external systems are subjected to substantiated integrity and validation.

Table 6. Security Aspects of the Aviation System Stack

The software execution environment is typically built on top of an **operating system** that serves as an intermediary between the software and the cyber enabled hardware on the aircraft. The operating system is an integral part of cyber resiliency as it implements security services and controls. The operating system also interfaces with the critical cyber physical systems on the aircraft.

Aviation systems also frequently rely on **middleware** which is software that acts as an intermediary between the operating system and application software. Middleware typically provides common services and capabilities, such as messaging, authentication, and data management. Middleware also frequently features application programming interfaces (APIs), that simplifies the development of applications by handling the underlying complexities of network communication and data exchanges.

In an aircraft, the **application** layer is where the functions and capabilities of the platform are implemented. It provides flight control logic, flight management systems, navigation displays, communication systems, warning systems, and maintenance diagnostic tools. Aircraft applications process and display data from sensors, systems, and other sources to the pilot and support crew. Additionally, the application layer in modern aircraft often includes data exchange with external systems, such as air traffic control, weather information services, and ground-based maintenance systems.

While the application-level software is the first line of defense and is often the major focus when developing and deploying secure software, the aircraft is only as resilient as the aggregate components that make up the full aviation stack. Exploitation of the platform can occur through multiple vectors within all the software. This is why software security must address the entire stack, and cyber risk cannot be accurately determined by partial evaluation of any single layer.

For that reason, curating the selection and provenance of software used within the aviation system's full stack is vital to support cyber resiliency. The importance of managing the supply chain is a recurring theme throughout this FSAD guidebook. It requires a rigorous selection process coupled with robust provenance tracking. The application software is only as secure as the weakest link within the full stack.

This is what makes software security and software assurance such a difficult task. It requires a thorough understanding of the security aspects of the entire software implementation.

13.6 Pipelines: Pulling It All Together

A unified software development effort leverages people, process, and technology in pursuit of achieving cyber resiliency in the system. A high-performance software development team uses the techniques that have been introduced in this chapter. That includes continuously assessing and improving the quality of the code. To the greatest extent possible, automation is used to regularly look for security vulnerabilities, address any identified issues, and to update secure coding practices to prevent future recurrence.

One mechanism for pulling all this together is to create a “pipeline” which is a series of automated process stages used to build, test, and sometimes even deploy software. The pipeline is supported by tools that provide source code management, code compilation/build, static code analysis, testing, and vulnerability scanning. The objective of a pipeline is to establish a consistent and repeatable process for software production that leverages technology and automation.

When developing software for aircraft, it is important to remember that the tools and processes commonly used in pipelines may not be specifically tailored to the unique requirements and constraints of the aviation domain. Nevertheless, these tools should still be used. However, it is crucial to recognize the limitations and supplement with necessary human domain expertise and custom made tools when possible. While enterprise IT systems may use a pipeline that includes automatic deployment, that is currently not prevalent in the aerospace domain due to airworthiness concerns.

Pipelines are a highly recommended best practice for developing secure code due to the automation and standardization they provide throughout the software production process. It is a best practice for maximizing the security and reliability of the software.

13.7 Technology and Tool Evaluation and Selection

Selection of the tools and technology to support the software development lifecycle and pipeline will vary based on a number of factors. This FSAD guidebook does not prescribe any one particular product, because it would be foolhardy to do so. In the aviation domain, there is no such thing as a universally “best” tool, software product, or technology. Instead, the selection process should be tailored to the specific needs and requirements of the aircraft platform.

A new aviation software development project will face a series of decisions on how to implement the deployed operational system. Each of those choices will have a tremendous influence on the tools and technology that best support the effort. The nature of the aviation system under development may dictate performance objectives and constraints that drive the selection of the hardware, operating system, software programming language, and the tools used for compilation and building. Those factors will subsequently drive the selection of the static code and dynamic code analysis tools, vulnerability scanners, and automated testing support.

The prior experience of members of the development team is also a significant factor when selecting tools for new projects. If the developers are comfortable and proficient with a particular product or technology, selecting those tools can reduce or eliminate the learning curve. Additionally, existing licenses, training, and infrastructure may make it impractical to change without a compelling technical reason.

However, it is important to periodically evaluate the current software development tool suites to consider whether there are newer or more suitable tools available that could improve the development process. This evaluation should take into account the specific needs and requirements of the project, as well as the strengths and weaknesses of the current tooling. While changing software development tools can be disruptive and require additional effort, it may ultimately lead to long-term benefits such as increased efficiency, improved quality, and enhanced security.

13.7.1 Selection of the Run-Time Environment

The target processing architecture, hardware, and operating system is a foundational decision made in conjunction with the hardware assurance considerations described in Chapter 14. Additionally, this is a foundational aspect of the “Full Aviation Stack” described in section 13.5.

Unlike enterprise IT desktop systems, the real time embedded nature of aircraft will likely dictate specialized processing architectures and microcontrollers. These might be single-core or multi-core CPUs with tight constraints on power requirements and heat-dissipation. Additionally, the control of real time cyber physical systems used in aircraft likely point toward a variety of specialized devices.

Ruggedized processors and components are often needed to operate effectively in the harsh environments and conditions presented by aviation systems. Ruggedization involves designing and building hardware and software components to withstand extreme temperatures, vibrations, shock, and other environmental factors that can impact performance and reliability.

When making decisions about the run-time environment for aviation system software and hardware, it is also crucial to consider the long-term service life of aircraft platforms. The selection of a processor should be restricted to components that will continue to be available to maintain and sustain aircraft platforms. When hardware parts are no longer being manufactured or become difficult to replace, changing to a new processor will impact the aviation system software as well.

Sometimes the run-time environment may be preordained based on legacy aircraft platforms or other hardware constraints within the system. However, it is important for the software development team to consider the run-time environment as a critical factor in the software architecture and design process as well as when selecting tools and methodologies.

13.7.2 Selecting Programming Languages

One of the most fundamental decisions facing new software development efforts is which programming language to use. Many factors go into making this choice including tool support, target processor, board support, target operating system, and the experience base of the team. The decision is also driven by security features and performance requirements.

The real-time determinate nature of aviation systems typically points toward compiled rather than interpreted languages. Compiled languages, which are translated into machine code before execution, offer several advantages in terms of performance and efficiency compared to interpreted languages, which are executed line-by-line at runtime. Compiled languages typically have faster execution times, lower memory usage, and better optimization capabilities. Those characteristics are critical for aviation systems that require real-time processing and high levels of performance.

Additionally, compiled languages can provide better type safety and memory management, which reduces the risk of errors and improves overall system reliability. While interpreted languages may offer advantages such as ease of use and rapid prototyping, the stringent performance and safety requirements of aviation systems has historically resulted in compiled languages being selected as a more suitable choice.

Memory safe programming languages are gaining advocacy in multiple domains as a panacea for creating intrinsically secure software. This type of programming language is designed to prevent common errors related to memory management, such as buffer overflows, null pointer dereferences, and memory leaks. Those types of errors can lead to security vulnerabilities, system crashes, and other issues. That is a critical consideration in software development, particularly for aircraft.

However, memory safe programming languages can be functionally limited in certain applications, including those that require low-level control over system resources, or the high-

performance processors used by aircraft. This is because memory safe languages often impose additional overhead and constraints on memory management, which can impact performance and limit the ability to optimize code for specific hardware.

Additionally, some memory safe languages may not support certain features or capabilities that are available in other languages, such as direct memory access or low-level system calls. Therefore, when developing software for aviation systems or other high-performance or safety-critical applications, it is important to carefully evaluate the trade-offs between inherent memory safety and functional requirements and select a language that balances those considerations appropriately.

13.7.3 Selection of Software Evaluation Tools

Aviation software development teams are likely to be involved in the selection process for automated software evaluation tools. That includes static code analyzers, dynamic code analyzers, vulnerability scanners, and automated testing environments. Several factors should be considered to ensure that the tools are effective, efficient, and aligned with the specific needs of the software development project.

First and foremost, the tools must be compatible with the selected programming languages and run time operating environment used on the aircraft. Otherwise, they cannot effectively analyze and test the codebase for potential vulnerabilities or issues.

Software evaluation tools should also be selected based on their accuracy, reliability, and false positive/negative rates. This will ensure that the tools can accurately detect vulnerabilities and issues, while minimizing the risk of false positives or negatives that can lead to wasted time and resources. The tools should also be easy to use and support integration within the development pipeline. Scalability, cost, licensing, and ongoing support for updates as new security issues emerge should also be a factor when considering these tools.

The development pipeline will be constrained by the performance of the software evaluation tools it uses. The software development team must make tool selections that are effective, efficient, and aligned with the specific needs of the project. This will ultimately lead to more cyber resilient software.

13.7.4 Configuration and Change Management

Configuration and change management is essential when developing and deploying secure systems. It is absolutely critical for managing multiple software versions and tracking problems or vulnerabilities to closure and resolution. Consequently, configuration management and change management tools are an essential part of the production of cybersecure and resilient software.

When evaluating tools and processes for configuration and change management, several factors should be considered. The configuration management products must be compatible with the specific programming languages, frameworks, and platforms in both the development and target environments. This will ensure that the tools can effectively manage and track versions and updates to the code.

Additionally, the configuration and change management tools should be easy to use and compatible with the development pipeline. Just like any other software product, scalability, cost, licensing, and support options are also a consideration.

Poor configuration management processes and tools can make it difficult to develop and deploy secure software. Without proper configuration management, the software development team may struggle to track and manage various code baselines. That can lead to potential inconsistencies and errors that can impact reliability and security. It is essential to implement effective configuration management processes and tools.

13.8 Specific Techniques for Cyber Resilient Software

A secure software development lifecycle requires not only the right processes and tools but also proper techniques and design patterns for building security into the product baseline. Using these mechanisms correctly reduces the occurrence of weaknesses and defects that can lead to exploitable vulnerabilities.

Security defects are fundamentally different from those resulting from an unmet functional requirement. Software can appear to work perfectly and satisfy capability objectives yet contain vulnerabilities or inherent weaknesses that can make it execute in unintended ways. That can create a multitude of negative consequences including software crashes, or even allowing a cyber adversary to create malicious effects on the system.

This chapter introduces some concepts of the countermeasures and techniques that should be considered for cyber resilient software in aviation systems. However, it is important to note that the breadth and depth of all mechanisms are too numerous to list in this FSAD guidebook. Additional resources are listed in section 13.10 that provide more exhaustive treatment of this vital aspect of secure software development.

13.8.1 Address Space Layout Randomization

The concept of diversity within a cyber-enabled system was introduced in section 7.5.1.8. That technique creates intentional variation of components to enhance system resilience, security, and robustness. One way that design pattern can be implemented in software is through the use of address space layout randomization. That technique creates executable code at unpredictable and ever-changing positions in memory. That negates the viability of certain types of cyber-attacks. Some compilers even directly support randomizing elements of the processing layout.

When address space layout randomization is used, attackers cannot predict exactly where executable software and data will end up on the stack. That makes it difficult to craft exploits that work reliably and consistently.

However, it should be cautioned that address space layout randomization techniques should be used with caution. While it is an effective approach that prevents certain types of cyber-attack, the indeterminate behavior can negatively impact airworthiness certification. When it is used it must be done so with caution and transparency.

13.8.2 Substantiated Integrity

Substantiated integrity is the use of cryptographic checksums to provide assurance that data, systems, and processes, have not been modified. This concept was previously introduced in section 7.5.1.5.

It is important to re-emphasize that while substantiated integrity is similar to other integrity mechanisms such as parity checks and checksums, it is considerably more robust than those non-substantiated methods. An adversary with the sophistication to tamper with software or data on an aircraft will also likely have the ability to easily recalculate parity and checksum values.

Substantiated integrity is typically implemented in software using keyed cryptographic methods to verify the authenticity and integrity of data and software. An attacker without the appropriate cryptographic key cannot generate a digital signature that matches the file. Substantiated integrity provides a much higher level of confidence in the accuracy and authenticity of data, as it is resistant to tampering, modification, and other types of attacks.

Due to the data intensive nature of aviation systems, the substantiated integrity archetype should be leveraged extensively. Candidates for use include all software and configuration data consumed by the aircraft platform. Whenever possible, external communication interfaces should be protected using this technique to prevent against spoofed messaging attempts. Along those same lines, substantiated integrity should also be considered for use on internal messaging on the aircraft. That protects subsystems from attacks should another subsystem be compromised.

Substantiated integrity relies on the concept of a "root of trust," which is a foundational aspect of secure systems. A root of trust ensures integrity and authenticity on a platform by providing a secure, immutable starting point for all trust decisions. It typically involves hardware and software mechanisms that securely manage cryptographic keys and certificates, forming the basis for verifying other software and data on the aircraft.

Since most aircraft do not have persistent connections to cryptographic certificate services, it is challenging to truly establish that immutable starting point. Sometimes the verification keys and certificates are loaded to the platform along with the software and data it is protecting. An attacker with access to both can potentially make modifications at will. Due to operational constraints, establishing a root of trust on an aircraft is a complex and demanding task. This is something that architects, designers, and software developers must take into consideration.

Substantiated integrity provides a high degree of trustworthiness that data, software, or a system hardware component has not been illicitly modified. When those integrity checks pass, the platform can operate with high confidence that it is processing good software and data. When substantiated integrity failures occur, it is a possible indication of cyber adversity.

13.8.3 Checksums and Vertical Parity Checks

A checksum is a mathematical calculation used to verify the integrity of a file or message. It involves calculating a value based on the contents of what is being evaluated using an algorithm, and then comparing the result to a previously calculated value to ensure the contents have not

been altered or corrupted. Checksums are commonly used in data transmission and storage to detect errors or inconsistencies in data.

A vertical parity check is another technique used to detect errors in data transmission. It is implemented by adding an extra bit to each data unit, such as a byte or word, to ensure that the total number of “1” bits in each data unit is even or odd. That helps to detect errors in data transmission, such as when a single bit is flipped or changed.

It is important to emphasize that checksums and vertical parity checks are not an effective security control for protecting against data manipulation from sophisticated cyber-attackers. Both can be easily defeated because it is a simple mathematical calculation. While both techniques are useful for detecting errors and accidents, they are ineffective as a cybersecurity mitigation.

If it isn’t possible to implement substantiated integrity in the aviation system, then checksums and vertical parity checks should be used. However, it is important to understand and document the inherent risks associated with that implementation.

13.8.4 Hashing

Hashing is an integrity verification approach that is relatively inexpensive and easy to implement. As the name suggests, it is a one-way cryptological algorithm that generates a hash of a fixed length. Because it is a one-way process, the data cannot be reconstructed from the hash. If a single bit of the information is altered, the hash produced is completely different. Hashing works by passing the software or data file that was hashed, along with the hash value for verification.

Hashing has similar limitations to substantiated integrity that doesn’t have a well-established root of trust. If the hash is packaged with the data it is protecting, or through a shared communication pathway, the attacker may have the access to change both the data and the hash. Consequently, when hashing is used as a security control, “out-of-band” distribution that separates the data from the hash is essential.

If a system does not have access to the certificate services required for substantiated integrity or access to a viable root of trust mechanism, hashing is a reasonable approach for software and data. However, measures must be taken to create separation between the hash and the information it is protecting. The limitations must be understood when assessing residual risk against the aviation platform.

13.8.5 Input Validation

Input validation is a disciplined process of checking and ensuring that data or input received by the system meets security criteria. While substantiated integrity assesses if data and information comes from an authorized source or has been tampered with, input validation verifies that the data falls within expected parameters and boundaries before it is processed. Input validation is an important security measure to prevent malicious attacks through spoofed messages and malformed data.

In an aviation system, input validation involves examining the data received from various sources such as sensors, user interfaces, and other systems for correctness and integrity. This can include checking for data type, format, length, range, and other constraints.

Effective input validation can help prevent common security threats such as injection attacks, where an attacker attempts to insert malicious code or commands into the system. It can also help prevent errors and anomalies that can cause the system to behave unexpectedly or malfunction.

Input validation should be implemented throughout the aviation system to ensure comprehensive protection.

13.8.6 Bounds Checking

Bounds checking is a software technique used to ensure that a program operates within the bounds or limits of a data structure, such as an array or a string. It involves checking the indices or values used to access or manipulate the data to ensure that parameters are within the valid range for the data structure and type.

Bounds checking is an important best practice to prevent memory corruption, buffer overflows, and other security vulnerabilities that can be exploited by attackers. By ensuring that the program operates within the bounds of its data structures, the technique can prevent memory access or corruption issues. It is also good for preventing unintentional errors or inconsistencies in the data.

Many modern programming languages and frameworks include built-in bounds checking capabilities, while others require the use of external libraries or tools to perform that verification. Regardless of the approach, it is important to implement bounds checking as a standard best practice in programming for all systems.

For aviation platforms, bounds checking is particularly important due to the safety critical nature of flight and the potential consequences of errors. Aircraft failures resulting from exceeding the bounds of data structures can have severe consequences, including loss of mission execution, loss of the platform, and potentially even loss of life.

13.8.7 Malicious Commands and Messages

Aviation software is designed to interact with other systems onboard the aircraft as well as with external platforms using messaging protocols. Correctly processing and responding to commands and messages is a functional requirement of all aircraft. It is necessary to perform the mission essential functions of flight controls, communication, and navigation.

The software on the aircraft must also be able to identify malicious or malformed messages. Categorically, aviation platforms should be designed and built to respond securely and appropriately to unexpected messages. However, sorting out legitimate communication from erroneous or malicious ones is one of the more challenging aspects of software design.

This is why substantiated integrity is recommended for all messaging and communications. However, due to the prevalence of legacy systems and protocols, that simply isn't feasible in many instances. The mechanisms of substantiated integrity must extend across all communications interfaces, and the aviation infrastructure simply isn't ready for that transition.

Consequently, aviation system software designers must implement mechanisms that prevent potentially catastrophic responses to malicious or erroneous messages and inputs. That includes awareness of the state of the aircraft as well as input validation mechanisms to the greatest extent possible. When writing the software that processes messages and commands, the developers must always consider methods to validate what was received along with mechanisms for preventing it from negatively impacting mission performance or safety of flight.

13.8.8 Software Implementation and Zero Trust Architecture

Zero trust architecture is a security model that assumes that all network traffic is untrusted. It typically requires strict identity verification and access controls for all users, devices, and subsystems, regardless of their location. The principles of zero trust were first introduced in section 7.7.6 of this FSAD guidebook.

While zero trust has significant traction for enterprise IT systems, implementation of the concepts for aircraft has proven to be challenging. Verification of source and identity does not always meet real time processing constraints of aviation systems. In addition, the mechanisms for identity and authenticity necessary for end-to-end zero trust implementations are not available across the majority of the industry infrastructure. Consequently, adoption of zero trust architectural concepts is still in the nascent stages for aviation systems at this time.

However, software designers and developers can still embrace the core philosophies underpinning zero trust when implementing aviation systems. That starts with the assumption that all messages and data received are suspect. Development should begin by conceptualizing how communications can induce negative consequences and determine if there is a way to prevent that from happening. Similarly, as software is written, the developers should consider means and methods for potentially identifying inauthentic messages.

While zero trust architecture has significant advocacy for enterprise IT systems, there is still work to be done to figure out how to fully implement the concepts for aircraft. However, software developers should embrace the principles and lean forward into zero trust concepts where possible.

13.8.9 Validate and Control Software Loading

Validating and controlling the software that is loaded on aviation systems is crucial for preventing malicious software injection and intrusion. Ensuring rigorous validation and control of all software and data loaded onto the platform helps protect against unauthorized access, tampering, and the introduction of malicious code.

Substantiated integrity is recommended for all software and data that is loaded to the aviation system. If that is not possible, software developers and the operational community should

collaborate on alternative methods to ensure that software received for the platform is from a reputable source and has not been maliciously tampered with.

Additionally, software loading on the aircraft should include physical interlocks to prevent accidental software updates. That also makes it more challenging for malicious adversaries to remotely add new software to the platform.

13.9 The Software Supply Chain and the Software Bill of Materials (SBOM)

Understanding our collective responsibilities as both a consumer and a producer in the supply chain is a recurring theme throughout this FSAD guidebook. Considering the supply chain is particularly important within the context of software development.

Paying attention to the software supply chain is essential for both development environments and production systems. The software supply chain typically provides the compilers, integrated development environments, and the evaluation tools used to create and deliver software products. Any vulnerability or compromise in the development environment or pipeline can introduce malicious code, defects, or other security risks that jeopardize the entire system.

Diligence and scrutiny are also required when selecting operating systems, third party libraries, and any other software enabled products that are a part of the full aviation stack. Software developers must conduct thorough assessments of purchased software products, and vet suppliers and vendors.

Using trusted and verified sources for software and avoiding unverified open-source components or unreliable software is critical. These practices help developers reduce the risk of introducing vulnerabilities into their systems, resulting in more secure and reliable software for both development and production environments.

A Software Bill of Materials (SBOM) is a crucial artifact for aviation system software security. The SBOM provides a comprehensive inventory of all software components used on a deliverable system, including open-source libraries, third-party modules, and proprietary code. This transparency is essential for managing vulnerabilities, as it allows developers and security teams to understand the software that is actually running on their systems.

An SBOM is vital because it helps developers and security sustainers determine if aircraft software is impacted by emerging issues or vulnerabilities in common modules. The SBOM also enhances supply chain transparency by identifying and managing risks associated with external dependencies. This manifest of software resident on a system is necessary for effective vulnerability management on the platform.

13.10 Summary and Additional Resources

In this chapter we briefly surveyed several key aspects of Software Assurance that establish the foundational concepts necessary to create cyber resilient aviation systems. This overview was not intended to be a comprehensive guide, but rather to highlight specific best practices as they

apply to aircraft. Consequently, this chapter has an extensive list of additional resources for software developers who need more in-depth information.

- **Secure Software Development Framework (SSDF) Version 1.1: Recommendations for Mitigating the Risk of Software Vulnerabilities, National Institute of Standards and Technology (NIST), Special Publication SP 800-218, Version 1.1, February 2022. (68)**

A full description of the Secure Software Development Framework (SSDF) which provides detailed guidance on integrating secure practices throughout the software development life cycle (SDLC). By adopting the SSDF, organizations can significantly reduce the number and impact of vulnerabilities in their software. The framework's focus on early and continuous integration of security measures into the SDLC aligns with the high-security demands of aviation software.

- **DoD Developer's Guidebook for Software Assurance, Dr. William R. Nichols, Jr and Dr. Thomas Scanlon, Carnegie Mellon University, Software Engineering Institute. CMU/SEI-2018-SR-013, December 2018.**

This provides comprehensive guidance on software assurance principles and practices, crucial for ensuring that aviation software functions as intended and is free from vulnerabilities. The guidebook outlines the entire lifecycle of software development, emphasizing secure practices at each stage, from requirements definition to maintenance. This is particularly important for aviation which requires a high degree of software assurance and for platforms that have long service timespans.

- **Program Manager's Guidebook for Software Assurance, Dr. Kenneth E Nidiffer, et al, Carnegie Mellon University, Software Engineering Institute. CMU/SEI-2018-SR-025, December 2018.**

A guide for integrating software assurance into the acquisition lifecycle, which is crucial for DoD aviation systems. The guidebook emphasizes the importance of addressing software vulnerabilities early in the development process to avoid exponential cost growth and mitigate risks effectively. It outlines the roles and responsibilities of various stakeholders in ensuring software assurance. Additionally, the guidebook includes checklists, tools, and techniques for evaluating and improving software assurance, making it a practical resource for developing secure and resilient aviation software.

- **Computer & Internet Security: A Hands-on Approach, Wenliang Du, 3rd edition, May 1, 2022.**

A practical overview of the cybersecurity principles and practices, which are critical for developing secure aviation systems. This book covers essential topics such as network security, cryptography, web security, and vulnerability assessment, all of which are directly applicable to the challenges faced by aviation software developers.

- **Writing Solid Code, Steve Maguire, Greyden Press, 2nd edition, January 1, 2013.**

This book addresses the critical issue of software errors and provides practical guidance on writing bug-free code. It offers insights into common mistakes developers make and emphasizes the importance of detecting and preventing bugs early in the development lifecycle, by learning to ask the right questions about how they can be prevented and detected automatically.

- **Secure By Design, Daniel Deogun, et al, Manning, First Edition, September 10, 2019.**

This book emphasizes the importance of integrating security into the entire software development process. It offers patterns, best practices, and mindsets that are directly applicable to real-world development, helping developers create inherently secure software from the outset. The book covers essential security-promoting constructs such as safe error handling and secure validation, which are critical for aviation software that must adhere to stringent safety and security standards.

- **24 Deadly Sins of Software Security: Programming Flaws and How to Fix Them, Michael Howard, et al, McGraw Hill, 1st edition, September 24, 2009.**

An overview of the most common design and coding errors that lead to security vulnerabilities. The book provides practical solutions for fixing or avoiding these flaws, ensuring that secure coding practices are implemented from the development phase, thereby reducing the risk of critical security issues.

- **Common Weakness Enumeration (CWE), MITRE, cwe.mitre.org**

This is a valuable resource for all software developers because it provides a comprehensive and well-structured catalog of common software vulnerabilities and weaknesses. By utilizing CWE, developers can better understand potential security risks associated with their code, which is crucial in the high-stakes field of aviation where software reliability and security are critical.

- **Common Vulnerabilities and Exposures (CVE), MITRE, cve.org**

This is a standardized publicly available list of known cybersecurity vulnerabilities. This resource allows developers to stay informed about the latest threats and vulnerabilities that could affect their systems. By leveraging CVE, developers can quickly identify and address specific vulnerabilities within their software.

Chapter 14

Hardware Assurance

This chapter was written by Marcus De La Garza and Jeff Langham and edited by Teresa Merklin.

Ensuring the security and reliability of the hardware in aviation systems is critically important for mission performance and safety of flight. Hardware assurance is the art and practice of preventing successful adversarial attacks or disruptive accidental events that impact cyber-enabled electronic assets. This chapter describes comprehensive processes and practices aimed at producing hardware components that are free from vulnerabilities and function as intended.

Modern aviation platforms rely on cyber-enabled systems to create high-performance and advanced capabilities. Gone are the days when aircraft were primarily built using mechanical analog components. As aircraft systems become more sophisticated and reliant on electronic and digital technologies, the potential risks associated with shortfalls of computing hardware parts grows exponentially.

This chapter delves into the critical aspects of hardware assurance in aviation. It explores how attacks occur along with the methodologies, tools, and best practices to mitigate those threats. Robust hardware assurance measures have become a necessity in the aviation industry.

14.1 Introduction to Hardware Assurance

Hardware assurance is a crucial component of cyber resiliency that supports and preserves the reliability and trustworthiness of the cyber-enabled components of a system. It is an important aspect of defending aviation platforms against highly sophisticated cyber-attacks. Robust hardware assurance is essential for maintaining the confidentiality, integrity, and availability of the critical processing elements that are the backbone of advanced modern aircraft. It is vital for establishing and maintaining cyber resiliency by defending the hardware from adversarial cyber-attacks as well as other incidental cyber events.

Defending the physical processing systems on an aircraft is critically important for mission performance and safety of flight. General processors, Application-Specific Integrated Circuits (ASICs), Field-Programmable Gate Arrays (FPGAs), and System on a Chip (SoC) must all be defended. Additionally, various memory devices including Random Access Memory (RAM), Read-Only Memory (ROM), and other non-volatile storage devices must also be considered.

When assessing hardware assurance, it is also important to not overlook the potential impacts of Printed Circuit Boards (PCBs). The PCB provides a structured, organized, and compact

interconnection platform for cyber-enabled electronic components. It provides signal routing, power distribution, and heat dissipation that interconnects devices into a cohesive system. PCBs are vital to hardware assurance, because they present significant attack surface to cyber adversaries.

Additionally, simple integrated circuits (ICs) and passive components such as resistors and capacitors can also have a bearing on the cyber resiliency posture of a system. While simple electronic devices may appear less critical than more complex cyber-enabled hardware parts, they can still contain vulnerabilities that lead to system failures.

The intellectual property (IP) stored inside cyber-enabled hardware parts includes software, firmware, user data, cryptographic keys, and configuration settings. All that information must also be protected as unauthorized access or modification could lead to severe consequences. Intellectual property theft is a serious concern for aviation systems and is sometimes a precursor to more sophisticated attacks against the aircraft.

Historically in the cybersecurity domain, fewer security issues have been found in hardware than in software and networking systems. However, when hardware security vulnerabilities are identified, they are frequently more impactful and disruptive than those which are exclusively confined to software. Additionally, the concept of “trusted hardware” is often used as a basis for a software root of trust. That compounds the risks associated with compromises of hardware assets. Moreover, hardware security issues are extremely difficult to resolve after the physical assets have been manufactured and deployed.

Attacks on hardware systems are broad in nature. The techniques include IP piracy, reverse engineering (RE), counterfeiting, micro-probing of ICs, tampering of traces or components on a PCB, snooping, and accessing privileged information through test and debug interfaces. Understanding the various attack techniques on hardware systems is necessary to effectively design, implement, and maintain robust hardware assurance measures. Knowledge of how hardware components are attacked is essential when designing and developing high assurance systems that support resiliency in cyber contested airspaces.

Cyber resilient aviation systems rely on trustworthy high assurance components. That characteristic is only achieved when cybersecurity is prioritized and intentionally curated throughout the development lifecycle. Trustworthiness is the culmination of secure design principles, stringent fabrication/manufacturing processes, and rigorous validation testing. These best practices are vital for hardware produced or acquired for use on aviation platforms.

14.2 The Engineering Trade-Offs of Hardware Assurance

Designing high assurance hardware for aviation systems is resource intensive. The intellectual capital that provides advanced capabilities and performance often represents a significant investment. Safety of flight considerations necessitates extensive verification and validation testing. For those reasons, the timelines to develop aviation hardware can be significantly longer than what is usually needed for other domains.

Unfortunately, lengthy development timelines make it difficult to quickly update hardware components if problems or cybersecurity issues emerge. Issue resolution is expensive, and it may take years to redesign and update any impacted subsystems. Once hardware is deployed to an aircraft, it is likely to remain in operational use for years.

Aviation systems rely on extremely complex hardware that utilizes a multitude of special purpose components and interconnections. That complexity can make it difficult to ensure that the various electronic parts of the system are functioning correctly and securely. Consequently, creating hardware parts with built in debug and test ports is frequently recommended as a best practice when designing and developing these complex systems. It is essential for testing and debugging efforts.

Additionally, test and debug ports are also frequently regarded as a useful feature for maintenance of systems after the hardware has been deployed. Allowing service personnel to connect specialized tools and equipment to debug and test ports provides visibility into the internal operation. It enables identification and potential resolution of issues that may be difficult or impossible to otherwise diagnose.

Test and debug ports also provide the means to load new firmware or software. That is essential for updating capabilities as well as patching systems should security vulnerabilities emerge. Providing some means to update hardware parts with security patches and bug fixes is a cybersecurity best practice in all domains, including aviation.

Unfortunately, all the benefits of developing and deploying debug and test ports on hardware components can also present a substantial cybersecurity concern. It significantly increases the attack surface for an adversary with physical access. The test and debug ports that are essential for diagnostics and maintenance of the aviation hardware can also be used to gain unauthorized access. Visibility into the internal workings of the hardware could allow a cyber-attacker to extract sensitive information, modify the behavior of the compromised device, or even take control of the system.

Consequently, this is yet another example of the engineering trade-offs between usability and cybersecurity. The decision on the extent to which debug and test ports are included on high assurance hardware involves balancing out the benefits against the drawbacks.

Some of the attacks that are described in the remainder of this chapter require highly specialized tools and assume the attacker has unfettered physical access to the hardware. The appropriate design mitigations and countermeasures will depend on the sensitivity of the data and the importance of the hardware for safe and routine operations. The level of protection needed is the engineering trade-off between the value of the data and the criticality of the hardware.

14.3 Introduction to Hardware Components

This section introduces the cyber-enabled hardware parts that play a vital role in modern aviation systems. These components are critical to delivering capabilities on the platform. They also play an important role in the cyber resiliency posture.

14.3.1 Integrated Circuits (IC)

An integrated circuit (IC) is a small electronic device that contains numerous interconnected electronic circuits on a single piece of semiconductor material. The integration of multiple electronic circuits onto a single chip reduces size, weight, and power consumption which is always a key consideration for aircraft usage. ICs also bring improvements in performance and reliability over using multiple devices to accomplish the same functionality. Integrated circuits are a foundational building block of modern electronic systems.

14.3.1.1 General Purpose ICs

General-purpose integrated circuits are designed to perform a wide range of functions across multiple domains. They provide basic electronic functionality, such as amplifiers, oscillators, and digital logic gates. The widespread commercial availability lends these devices for broad use in things like consumer electronics, industrial automation systems, and automotive applications. General purpose ICs provide a broad set of functions that can be tailored to a specific purpose while keeping costs low.

Cyber-attacks against general purpose ICs are limited to effects based on the functionality of the chip. For example, modifying an amplifier circuit, could cause instability and result in errors. Depending on the role within the aviation platform, an IC may or may not be an attractive target for a cyber adversary. Additionally, given the simplicity of interfaces to the device, it could be difficult to find a reliable trigger to attack it.

While general purpose ICs are not typically regarded as good attack vectors, they cannot be summarily excluded from a comprehensive cyber risk assessment. However, it is generally believed that more attractive targets will be cyber enabled hardware devices that contain, process, or transmit sensitive data rather than these general-purpose ICs.

14.3.1.2 Application-Specific Integrated Circuits

An Application-Specific Integrated Circuit (ASIC) is a microchip designed for a specific application or task. Unlike general-purpose ICs which are designed for use in a wide range of domains, custom ICs are tailored to meet the exact requirements of a particular system.

When working with systems that contain custom ICs, it is important to understand that the design and fabrication process is rarely performed by an individual company. Building and maintaining a semiconductor fabrication foundry is a significant capital investment. Most design organizations cannot afford to build custom chips from start to finish. Typically, custom ICs consist of a mix of proprietary design and purchased third-party IP that are sent to an external foundry for fabrication, testing, and packaging.

This reliance on third-party applications, complex designs, and external fabrication services makes custom ICs significantly more vulnerable to malicious insertions, undocumented functionality, or unintentional vulnerabilities. For example, third-party purchased IP might contain a hardware trojan that creates a vulnerability or backdoor into the chip.

Compromises at any stage of the IP selection, design, or fabrication process could be used to leak sensitive information or trigger in-flight attacks on the aircraft that would obviously be a serious concern. This type of cyber threat is extremely difficult to detect as the search is comparable to the proverbial needle in a haystack. Detection typically hinges on isolating unexpected or triggerable behavior, identifying a malicious circuit at the nanometer level of the physical design, or the use of manufacturing-based tests.

The full design flow for a typical custom IC illuminates the many potential entry points for adversarial cyber insertions into the hardware. Once the need for a custom IC is identified, a designer or integrator identifies the functionality and performance aspects needed. That kicks off a series of build versus buy decision for the IP required to accomplish that task.

IP cores can be purchased at various levels of abstraction including soft IP, which is at the register-transfer level (RTL), firm IP which is at the gate level, and hard IP which is at the layout level. The IC design house then integrates the IP cores it designed or purchased into a hardware-descriptive language (HDL) RTL for the entire IC. Before fabrication, the RTL is put through a series of simulation and functional tests.

Once the RTL design is verified, the designer synthesizes it into a digital logic gate netlist. The designer might also integrate design-for-test (DFT) structures, such as scan-chains or a built-in self-test, to improve the design's testability. Finally, the designer translates the gate-level netlist into a physical layout and checks it for potential manufacturing process issues, voltage, and temperature variations to ensure that timing and power requirements are met.

The final layout is transmitted to a foundry for fabrication. Starting with a blank silicon wafer, the foundry creates the physical design and performs tests on the wafer to identify any manufacturing defects. Once fabrication is complete for multiple wafers, these are sent to an assembly house, where they are cut into individual dies and subsequently packaged into microchips.

In addition to multiple points where errors or intentional injections could occur during the fabrication process, this also presents significant supply chain risk. China dominates the semiconductor global supply chain, and many foundries are either directly in that country's control or direct influence. China has openly declared its intentions to control key aspects of the global supply chain, including leveraging dependencies to threaten and cut off foreign countries during a crisis. China's dominance in semiconductors poses significant cyber supply chain risk to Western aviation system manufacturers. (70)

14.3.1.3 Fixed Architecture Data Processing ICs

ICs that process data are attractive targets for adversarial cyber-attack. These devices are used to input and output significant amounts of data which constitutes both cyber-attack surface as well as command and control channels sometimes used in sophisticated long term adversarial cyber campaigns. Data processing ICs perform a wide range of complex functions, handle sensitive data, provide gateways into systems, and are frequently vulnerable to software attacks. Consequently, protecting data processing ICs requires a multi-layered approach, including

hardware-based security features, secure software development practices, and network security measures.

Data processing ICs are classified into two architectural types. Fixed logic devices have a predetermined logic circuit design that is optimized for specific functions and low power consumption. Programmable logic devices are flexible and customizable, allowing them to be used in a wide range of applications. The choice between fixed and programmable logic devices depends on the specific system and design requirements.

Microprocessors and microcontrollers are fixed logic devices. Each has an instruction set with a defined set of functions. Common architectures include x86, ARM, PowerPC, AVR, PIC, and RISC-V. Developers write firmware for each device in their language of choice, which is compiled and assembled for the specific target processor architecture. The assembly code is then converted to machine code which is the software that executes on the hardware.

Microprocessors require external memory devices and peripherals that communicate over a shared memory bus. However, the fact that most of the data communication is performed external to the IC poses a security issue. An attacker with physical access to the connections on the bus could extract firmware and other sensitive data stored by the microprocessor system.

The risk of that type of data extraction can be mitigated by careful PCB design and IC package selection. One ideal solution is encryption of the data going across the bus and stored in memory. Unfortunately, not all microprocessors support encrypted data. While that can be added by external devices, it is inherently less secure than direct implementation within the microprocessor.

Microcontrollers are very similar to microprocessors as they process instructions from a fixed instruction set one at a time from firmware. However, they are different from microprocessors because they do not require external support components. For example, most microcontrollers have RAM, FLASH memory, and peripherals all built directly into the IC.

From a cybersecurity standpoint, microcontrollers enjoy an advantage. Sensitive data can be encapsulated within the IC and it never has to leave that boundary during normal operation. Additionally, many modern microcontrollers implement security features that prevent read and write operations to internal memory. That provides significant protection.

Unfortunately, the integrated functionality of microcontrollers usually comes at the cost of reduced performance and memory capacity as compared to microprocessors. Those factors are constrained by the available space on the IC. Consequently, despite the security advantages of microcontrollers, aviation systems frequently require a microprocessor implementation to meet performance objectives.

Fixed architecture data processing ICs are subject to reverse engineering attacks. In fact, if a cyber adversary extracts or accesses the firmware, powerful commercial and open source tools can greatly streamline the extraction of IP. A cyber adversary who is able to reverse engineer firmware will gain insight into the internal logic and functionality of these devices. That can be used to identify vulnerabilities, backdoors, or other weaknesses in the device's software or

hardware implementation. This knowledge can be leveraged to exploit the device, modify its behavior, or steal sensitive data.

Reverse engineering can also be the first stage for creating counterfeit or cloned devices that mimic the functionality of the original chip. Sometimes counterfeit or cloned parts are created strictly for economic benefit. However, more nefarious motivations such as malicious modifications or backdoors are also possible. There is currently no known reliable means of detecting if counterfeit or cloned devices are motivated by economic or malicious intent.

The techniques to prevent reverse engineering of firmware includes code obfuscation, encryption, and secure boot mechanisms. Code obfuscation involves modifying the firmware to make it difficult to understand or reverse engineer. Encrypting the firmware can prevent unauthorized access or modification. Secure boot mechanisms use cryptographic techniques to verify the authenticity and integrity of the firmware before it is loaded and executed.

If the architecture of a processor is known and the firmware is not encrypted, it is relatively simple to generate assembly code from the extracted firmware. From that point, understanding the functionality of the system is simply a matter of time.

14.3.1.4 Programmable Logic Devices

Programmable logic devices (PLDs) are a type of integrated circuit that can be programmed or configured after fabrication to perform different logic functions. PLDs are designed to be customizable, which can be an advantage over fixed logic microprocessors and microcontrollers when flexibility is needed.

PLDs can be classified into two main categories. These are Field-Programmable Gate Arrays (FPGAs) and Complex Programmable Logic Devices (CPLDs). FPGAs are ICs that can be programmed to perform various logic functions by configuring the interconnections between logic gates. CPLDs are configurable by modifying interconnections between logic blocks.

PLDs are commonly used in applications where flexibility, customization, and fast time-to-market are critical factors. Typical usage includes prototyping, testing, and low-volume production. They offer a cost-effective solution for applications that require frequent updates, as they can be reprogrammed or reconfigured as needed. PLDs are also frequently used in applications where a customized solution is required to meet system requirements that are unique or cannot be met by off-the-shelf components. Those factors frequently make them ideal for use in aviation systems.

Unlike fixed architecture processors which execute instructions from firmware, the structure and connections of digital logic gates in PLDs are set using a configuration file. That can create significant performance advantages over traditional processors since they can be configured to perform multiple tasks in parallel. That increases the amount of data that can be processed in a single clock cycle.

Rather than the traditional programming languages used for fixed processing architectures, PLDs are programmed using Hardware Description Languages (HDL). These languages are converted

to a configuration file for a specific PLD through a phased process. Synthesis converts the high-level HDL code into a netlist representation of the design that is mapped to the target PLD. A place-and-route phase generates the configuration file that can be loaded to the PLD using a programming device.

FPGAs are the most commonly used PLDs in aviation systems. However, CPLDs are also used. In general, CPLDs are smaller, use less power, and have fewer logic gates than an FPGA. It is also much more common for CPLDs to have internal configuration memory, whereas FPGAs typically have external configuration memory. However, there are exceptions to both those generalities.

While PLDs have many advantages over fixed architecture processors in terms of processing speed, the extremely low-level design process and relatively high cost can make them less desirable for general processing. When needed for performance reasons, PLDs are often paired with a fixed architecture processor which allows the PLD to handle high speed processing and data transmission while the microprocessor performs general tasking.

PLDs are also subject to reverse engineering attacks that are very similar to what was previously outlined for fixed architecture ICs. However, while the configuration file for a PLD can be extracted in the same way as a processor's firmware, it is considerably more difficult to reverse engineer the file. Each PLD manufacturer uses a unique process which can create significant differences in the files between PLD families and even PLDs within the same family. Even if a netlist can be created from the configuration file, turning that into understandable code is a daunting task.

It is important to note that reverse engineering of PLDs has been accomplished against small-scale designs. It is possible that a nation-state level cyber adversary has undisclosed capabilities for larger scale reverse engineering. However, the time and effort required to accomplish such a task should not be trivialized. Unless an attacker knows for certain that valuable data is contained within the PLD configuration file, it might not be worth the investment.

However, the state of the practice is ever increasing and eventually the difficulty of reverse engineering PLD data will surely become more feasible over time.

14.3.1.5 Memory Devices

Memory devices are attractive targets for cyber-attack simply because that is where sensitive and critical data is stored. Additionally, memory is also where a cyber adversary might achieve persistent access to a system. In order to protect against those threats, strong security measures such as encryption should be considered during hardware design to protect memory devices and the data.

Memory can be either volatile or non-volatile. Volatile memory devices can only store data while power is applied. Once power has been removed for some duration of time, all the data stored within is lost. On the other hand, non-volatile memory devices permanently retain whatever data has been written to them, even once power is removed from the device.

Volatile memory devices typically support much faster read and write operations than non-volatile ICs. Consequently, volatile memory is best for storing temporary data used for high-speed processing and data transfer operations. The most common form of volatile memory is Random Access Memory (RAM) which is used in all microprocessor-based devices. Due to the volatility, RAM is not typically used to store firmware, but may process firmware data. Since data busses used to connect microprocessors with RAM are not typically encrypted, it is possible for an adversary to access that sensitive data by observing it in transit.

Non-volatile memory includes devices such as Electrically Erasable Programmable Read-Only Memory (EEPROMs) and Flash memory. Due to the non-volatile nature, these devices are commonly used to store critical and sensitive data such as firmware, log files, and encryption keys.

The nature of the data stored on memory devices makes them a prime target for attackers. As a result, it is important to consider how the devices are connected and what type of data is allowed to be stored on them as hardware subsystems are designed. Encryption techniques are recommended to protect highly sensitive or critical data. However, for performance reasons that is more feasible on non-volatile devices used in less time-sensitive applications.

14.3.2 Printed Circuit Board (PCB)

A printed circuit board (PCB) is a rigid or flexible board that provides a physical platform for the components of an electronic device, such as microprocessors, resistors, and capacitors. PCBs are made by printing a pattern of conductive material onto a non-conductive substrate. This conductive pattern forms the electrical paths that connect the various components on the board, allowing them to function together as a cohesive system.

Depending on the intent of the design, circuits made of active and passive components are mounted on the top and bottom of the PCB. The electrical connections between the components are established using traces and vias, which are generally made of copper. Modern PCBs vary in complexity, ranging from a single layer of traces up to 20 to 30 layers of conductive material interconnecting with hidden vias between layers. Many times, the PCB form factor, such as an enclosure and any mounting methods needed to secure it, is a defining feature of the design necessary to accommodate the environment.

Numerous steps go into designing and building PCBs. As is the case with ICs, many entities are involved in design and manufacturing. The process starts with parts research and parts selection. The design engineer then performs schematic entry, which uses symbols to represent all the electrical connections of each component. The netlist is created by a design engineer who visually draws the interconnections between all the components. Simulation and verification are then performed to ensure that the design functions as intended and meets all requirements.

At that point, the netlist is typically handed off from the design engineer to a person or team that performs the layout. Using the netlist, the design house uses symbols that represent the true mechanical dimensions of the components used in the design. Generally, the initial arrangement starts with the larger components as well as the power and ground necessary for signal integrity.

This is followed by the placement of all the other smaller components as well as any debug headers or test points.

Once the initial placement is done for the components, the nets are routed on the board. Critical nets are usually routed manually while less critical nets may be auto-routed using rules provided by the designer. If the number of routes is congested making it impossible to accommodate all the nets, the layout engineer will add layers to the design. This increases the space required to make all the connections and meet the layout rules and guidelines for an effective design.

The design is verified through post-processing steps, which include automated checks on spacing rules, netlist opens and shorts, and a signal integrity analysis. These verifications ensure that the propagation of digital signals is clean and will not cause errors during operation. Finally, a set of descriptive design files (Gerber) and drill files are generated with the associated assembly diagrams for the fabrication house that creates the physical boards.

Once the bare boards have been fabricated, the PCBs go through additional steps. All the physical components are stenciled and placed using automated pick-and-place machines. Then, they proceed to a reflow oven, where the components are electrically connected to the fabricated boards. After the boards have been built, they are inspected to ensure that they were built correctly. Automated optical inspection and x-ray equipment is commonly used for quick visual inspections during this manufacturing step.

Once the physical integrity of the design has been verified, a powered test of the board is performed through various automated methods. Common test approaches include in-circuit testing, functional testing, and JTAG-based boundary scan testing. Each PCB is tested to verify and validate as much interconnectivity and functionality of the design's operation as possible. Only then is the PCB considered a success for either a prototype or production run.

PCB designers follow a general set of rules and best practices to ensure the devices are compact and reliable. Unfortunately, that historically has not involved considering hardware security. For example, it is generally considered a good design practice to label important connections on the silkscreen layer so they can be easily identified during troubleshooting. This also allows attackers to quickly locate important connections which can be exploited.

PCBs are the mechanism which connects all the cyber-enabled components together. By following the traces of a PCB, it is possible to reverse engineer the design without access to the schematic. This process is not trivial, and the complexity depends heavily on the number of layers used on the PCB. However, X-Ray imaging can be used to reveal the inner layers of a PCB without any physical modification of the board.

Reverse engineering is a real and legitimate threat against PCBs.

14.4 Common Hardware Attacks

This section focuses on the most common and straightforward types of attacks that could be carried out against hardware components within a system. These attacks can be executed using readily available electronic test equipment and can typically be completed in a relatively short

amount of time. Given their simplicity and the ease with which they can be carried out, it is particularly important to implement protective measures against these types of attacks.

14.4.1 Exploiting Debug Interfaces

Most modern programmable ICs include debug interfaces as that is considered a best practice in design. Debug interfaces allow users to program the device, monitor memory, and to perform stepwise execution through the firmware. Those features are incredibly useful for troubleshooting. However, debug ports can also be used for data extraction and exploitation. They present a significant security risk to the system.

14.4.1.1 JTAG

JTAG (Joint Test Action Group) is a standard for testing and debugging electronic circuits and systems. It specifies a set of test access ports (TAPs) and a boundary-scan architecture that can be used to test and debug digital circuits. JTAG allows engineers to test and debug electronic systems at the board or chip level, without the need for physical access to the individual components.

ICs that support JTAG can have their pins monitored and controlled externally through the debug interface. That enables debug software to check connections for faults. JTAG enabled devices also allow for monitoring and control of memory spaces within processors.

The JTAG interface consists of four mandatory signals (TCLK, TMS, TDI, TDO) and one optional signal (TRST). (71) While this is a somewhat large amount of pins required, one of the important features of JTAG is that multiple JTAG devices can be chained together through the TDI & TDO pins so that multiple devices can be accessed through the same debug port.

By default, JTAG has access to the memory spaces within processors. Attackers can utilize this debug interface to extract firmware, monitor sensitive data within RAM, and upload malicious software. While the most effective solution to prevent these attacks would be to completely remove JTAG connections from these devices, that is often not feasible. Most aviation systems are expected to be repairable and upgradable, and removing JTAG would limit those abilities.

14.4.1.2 Manufacturer/Architecture Specific Debug Ports

While JTAG is the most common standardized debug interface, alternative debug mechanisms are not uncommon. For example, Serial Wire Debug (SWD) is utilized by ARM based microcontrollers. SWD only requires two pins for operation, but still provides most of the functionality of JTAG including accessing memory regions of the device.

Most IC manufacturers have developed proprietary or custom debug protocols for their products. Every one of those interfaces can be used by an attacker to extract and manipulate data within a device. Since these debug ports are less common than JTAG, the information and tools required to use these interfaces are not as readily available. However, that obscurity does not equate to security. Any functional debug port can be exploited, no matter how poorly documented or supported.

14.4.1.3 Minimizing Debug Port Risk

Debug ports are a potential vulnerability in electronic devices, as they can provide a means of gaining unauthorized access or extracting sensitive data. In order to mitigate the risks associated with debug ports, it is important for designers to implement appropriate security measures. This section describes a range of potential mitigations to reduce debug port risk. They are listed in order from least effective to most effective.

- **Disguising the debug port header.** Manufacturers can place the debug connections in unexpected spots to hide them. For example, instead of placing a header for JTAG, a manufacturer could route the connections to unpopulated resistors. However, some automated tools can discover undocumented JTAG ports, so this method will most likely delay rather than prevent attack.
- **Enabling a password for the debug interface.** Implementing password protection is not a default feature of most debug interfaces. However, it can be implemented by the manufacturer. This is another delaying tactic, as password bypass is possible in many instances. This is still a viable solution for non-critical components.
- **Partially disabling debug features.** Some ICs allow the user to disable certain features of the debug interface, such as access to specific memory regions. The effectiveness of this solution is dependent on the implementation from the manufacturer. It is another security mechanism that can be bypassed in many instances.
- **Fully disable debug features.** The most effective method to prevent these attacks is to completely disable the debug port. For this to be effective, the device must utilize security fuses or bits that cannot be accessed or altered after programming. This is generally very secure. However, it prevents debugging and reprogramming of the device which may be desired during normal operation and maintenance. Additionally, depending on the manufacturer's implementation, bypass may still be possible.

14.4.2 Monitoring/Manipulating Exposed Connections

In addition to the debug ports mentioned previously, many PCB designers also add additional pads and connectors on a board, which allows for easy measuring of signals such as power lines and data buses. This can be very useful when verifying functionality. However, if left in the final design of the system, these extra connections can provide easy access for an attacker. Additionally, if the ICs on a board have exposed pins, an attacker could connect directly to those to extract data.

In addition to data extraction, these exposed connections can also be used to add additional ICs to a PCB. For example, an attacker could add a microcontroller that will begin sending data across a bus once some trigger has been met.

Very little can be done to prevent these kinds of attacks after the component selection, placement, and routing of a PCB has been finalized. Therefore, it is imperative that careful consideration is given to the cyber risks, the mitigation options at hand, and the trade-offs between functionality and security during the design phases of a hardware asset.

14.4.2.1 Minimizing Exposed Connection Risk

Exposed connections on a device can be a potential vulnerability, as they can provide a means of gaining unauthorized access to the device or extraction of sensitive data. In order to mitigate the risks associated with exposed connections, it is important for designers to implement appropriate security measures. This section describes a range of potential mitigations for exposed connection risk, listed in order from the least to most effective.

- **Don't label connections.** If an attacker doesn't have access to a schematic for a board, then it can be challenging to determine which pads are used for what purpose. Leaving a PCB unlabeled forces the attacker to first reverse engineer aspects of the design, which can be costly and time consuming.
- **Use ICs without exposed pins.** If a component is available in a package without exposed pins, it may be more difficult for an attacker to monitor the signals from the chip without removing it.
- **Remove test points from final production boards.** While test points on a PCB are useful in the design stage, they can result in security issues. If possible, remove them from the final production PCBs.
- **Route security-critical PCB traces on internal layers.** An attacker only has easy access to the external layers of a PCB, so routing signals within the inner layers will limit the exposed attack surface.
- **Cover exposed traces and ICs in a non-removable coating.** Non-removable coatings on the PCB, such as a hard potting compound, prevent most signal monitoring attempts. While it is possible to remove the coating, doing so is time consuming and may cause damage to the board.

14.4.3 Data Extraction Flowchart

Attack flowcharts are a useful tool that helps hardware designers conceptualize, identify, evaluate, and understand risks and threats associated with the design of devices and systems. Visualizing the attack process can guide decision-making as hardware systems are constructed.

Figure 20 is a flow chart that illustrates the thought process a cyber adversary might use when attempting data extraction from a device. This flowchart only covers the common attacks that do not require high skill or specialized equipment that are described in this section.

The flowchart is useful for identifying the likely steps that an attacker may take when attempting data extraction. That can assist with the decision-making process on how to prioritize defending the data extraction attack surface.

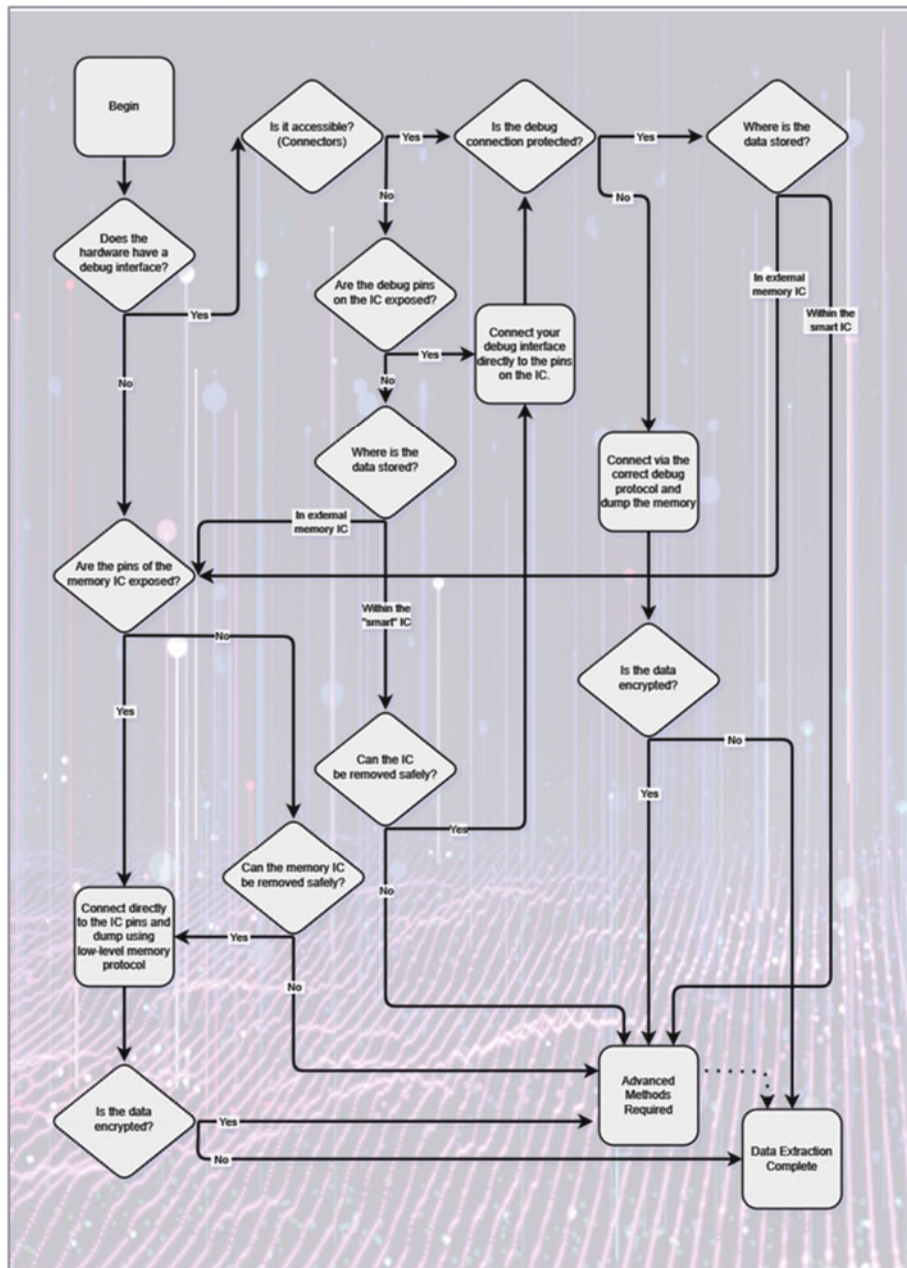


Figure 20. Data Extraction Methodology Flowchart

14.5 Advanced Hardware Attacks

This section delves into more advanced techniques that often demand a substantial time investment and carry a higher risk of causing irreversible hardware damage. This risk can be a significant deterrent, as it may render the hardware unusable for further exploit attempts, unless additional components are available. Although these advanced attacks may not be more complex to execute than the common hardware attacks covered in the previous section, they typically require a greater investment of time and resources. However, the potential payoff can be substantial, making the extra effort worthwhile for attackers who succeed.

14.5.1 Side-Channel Attacks

The attacks described in section 14.4 focused on extracting data directly at the source, usually through a direct physical connection. However, an attacker may not always have direct access to the physical connections which transmit data.

In the side-channel scenario, an attacker does not have direct physical access and is forced to leverage other methods of extracting data through indirect means. Some of the most common types of side-channel attacks are described in this section. However, it is important to note that other methods are known or suspected to exist. Additionally, new attacks could be discovered at any point. Side-channel attacks can be difficult to prevent since they utilize inherent properties of electricity and electronics. Therefore, mitigations that anticipate and prevent side-channel attacks must be implemented from the very beginning of the design.

14.5.1.1 Power Analysis

A side channel power analysis attack is carried out by analyzing the power consumption of a device to extract sensitive data or to infer information about the device's internal operations. This method works based on the principle that the device's power consumption varies depending on the data being processed or the operations it is performing.

By analyzing the power consumption of a device, an attacker can infer information about its data or operations, and potentially extract sensitive information such as cryptographic keys. Power analysis side-channel attacks require a variety of tools typically including oscilloscopes, logic analyzers, and power consumption monitors.

Physical modification of a PCB may also be required to gain access to the power lines of a device or to remove support components such as capacitors which may make detection of small current changes difficult. The success of this attack is dependent on the precision of the measurement equipment and the complexity of the device and data that is involved.

14.5.1.2 Electromagnetic Emissions

An intrinsic property of electricity is that a current running through a conductor generates an electromagnetic field at a right angle to the flow of the current. This field can be detected by an antenna and analyzed similar to the power analysis side-channel attack described in section 14.5.1.1.

Since this technique relies on electromagnetic fields rather than a direct physical connection, it has some advantages over power analysis side-channel attacks. Physical modifications to a device may not be a viable option in some instances. For example, that can lead to detection that an attack was attempted or damage the device preventing extraction of the data. In some cases, electromagnetic emissions side-channel attacks may be a better option to the attacker.

However, since there is no direct connection to the device, the placement and quality of the antenna require precision. The small electromagnetic fields generated by ICs are extremely

susceptible to noise. Care must be taken to accurately decipher the signals captured by the antenna. This is not an easy task.

14.5.1.3 Minimizing Side-Channel Attack Risk

Two potential mitigations against side-channel attacks are as follows:

- **Restrict physical access.** Since side-channel attacks utilize intrinsic and unalterable properties of electronic systems, one of the best defenses is to restrict physical access to the device. For example, encapsulating the PCBs inside a hard potting compound prevents direct access to IC pins. Alternatively, the PCB could be housed within an enclosure that is difficult to disassemble.
- **Use electromagnetic shielding.** Circuits which handle sensitive data should be encapsulated within electromagnetic shielding to prevent leaking data through electronic emanations. For the shielding to be effective against this attack, it must be attached in a permanent manner that would prevent an attacker from simply removing it. Conscious use of PCB ground layers could also help to prevent this attack.

14.5.2 Fault Injection Attacks

Fault injection is another broad category of attack that takes advantage of the intrinsic properties of electronic devices. While the side-channel techniques described in section 14.5.1 passively observe normal operation of the device, fault injection intentionally alters normal operation to achieve data disclosure.

By injecting extra signals or selectively removing ones that are already present, fault injection attacks can induce devices to bypass security features that may be protecting its sensitive contents. While these attacks can be very effective, success is the culmination of a lot of trial and error. Consequently, for an attacker to successfully pull off a fault injection attack, unfettered physical access to the device is required. Additionally, a high degree of patience and time are also necessary.

14.5.2.1 Power Glitching

A power glitching attack temporarily and selectively removes or interrupts power to a device causing it to behave erratically. That may induce it to skip security processing. This attack is typically carried out by interposing a transistor between the power pins of an IC and its normal power supply. The transistor is controllable via an external triggering device which enables precise timing of the power fluctuations. Some systems may also need to have other modifications, such as removing decoupling capacitors, to create the desired effect.

Quickly removing and restoring power is outside of the normal operating conditions for any IC. In theory, when power is removed from the chip, some internal sections of the device will lose power before others. That will likely cause the internal logic of the IC to behave unexpectedly. One potential effect is potentially skipping processor instructions. If timed when security features are being executed, that could produce significant impacts.

14.5.2.2 Clock Glitching

Inserting erroneous pulses into the normal clock cycle of an IC is another way to potentially induce erratic processor execution. This attack is conceptually similar to the power glitching techniques described in 14.5.2.1. The extraneous clock pulses can cause a processor to skip instructions or perform incorrect operations. That could potentially bypass internal security protections.

This method requires physical access to the pins of the ICs in order to insert the new clock pulses. However, this vector is less invasive than the modifications required to carry out power glitching. If the pins of the chip are exposed, it is possible to perform this attack without obvious signs of tampering.

14.5.2.3 Electromagnetic Glitching

Electromagnetic glitching is a technique that uses an antenna connected to a voltage pulse generator. When this apparatus is close to an IC, it can be used to create unexpected voltage levels within the device. That can result in various effects including errors and disabling built-in security features of the IC.

One of the advantages of this method is it doesn't require a direct physical connection to the IC to cause any effects. However, the electromagnetic pulses are difficult to manage and there is significant risk of damaging the IC's internal components. It is extremely difficult to regulate or monitor the voltage induced within any device without a direct connection.

14.5.2.4 Minimizing Fault Injection Risk

The potential mitigations for the risks associated with fault injection attacks on a device are as follows:

- **Restrict physical access.** The full range of potential fault injection vulnerabilities for specific devices are not widely known. That means that these attacks are not trivial, and success requires unfettered physical access and a lot of time. Iterating through multiple test cases could conceivably take days of constant testing to uncover a susceptibility. Additionally, there is no guarantee that any fault injection vulnerability even exists. Consequently, limiting close physical access to devices is an effective security control.
- **Use electromagnetic shielding.** Shielding can help prevent electromagnetic pulses from entering the device. Any IC which may be susceptible to an electromagnetic attack should be encapsulated in shielding to minimize the effectiveness of such an attack.
- **Select ICs that have fault injection detection.** Some ICs are implemented with technology that can perform fault detection. These systems can raise an alert if a fault injection attempt is detected. However, these ICs are not very common. In some instances, devices that meet the design requirements may not be available.

14.5.3 IC Deconstruction Attacks

Another class of advanced hardware attacks is IC deconstruction. If a device has security features that prevent reading of data using the other non-destructive attacks described throughout this chapter, one last ditch option is to deconstruct the IC.

Removing the IC packaging can provide direct access to the semiconductor die, which is the actual circuit of the IC. This die is typically considerably smaller than the package which contains it. The casing can be removed using either mechanical or chemical means.

This technique also requires access to powerful magnification devices as well as knowledge of IC manufacturing. However, with those two things, it is possible to reverse engineer the IC circuit to determine the precise location of memory as well as the security bit settings that might be preventing reads from the device. That might be enough to allow access either through modification of the security bit settings or potentially reading memory directly by micro-probing the exposed memory circuits.

IC deconstruction is usually the last resort for an attacker. It is likely to only be attempted when all other methods of data extraction have failed. The process of exposing the die and manipulating the state of the security bits is incredibly risky and has a high likelihood of causing permanent damage to both the device and data.

Unfortunately, beyond the deterrence of potential damage to the data, the options for protecting this attack are extremely limited. The following suggestions might help mitigate the risks associated with IC deconstruction attacks:

- **Encapsulate the IC in a coating.** Coating the IC in something such as a hard potting compound will make removing the casing much more difficult. This will also increase the risk of damaging the IC during the removal process.
- **Utilize ICs with more robust package materials.** Most ICs currently on the market are in plastic packages which provide little resistance against mechanical and chemical removal. Utilizing ICs which are made of more robust materials increases the difficulty of exposing the die. That also increases the risk of damaging the data.

14.5.4 Hardware Implant Attacks

So far in this chapter, the majority of the attack methods described culminate with some form of unauthorized data access. The purpose of such attacks is either to directly steal the IP or to perform firmware modification.

This section introduces the concept of hardware implant attacks. That opens the aperture of cyber physical effects on a system. For aviation platforms, hardware implants could result in devastating physical events during aircraft operation. With insight into the design of the platform, it is possible for a malicious adversary to create hardware implants that produce triggerable cyber effects.

Those attacks could either be activated through environmental conditions or potentially even intentional remote stimulation. The type of damage that can occur varies depending on the target systems and the sophistication of the implant. However, such attacks could imperil the mission, aircraft, and even human life.

While the potential effects of hardware implant attacks are quite severe, it is logistically complicated. Success requires extensive knowledge of the specific hardware target, as well as the internal operation and messaging of the aircraft. The attacker would also require physical access to the aircraft or the supply chain to inject the hardware implant. Additionally, the implant itself would have to be discrete to avoid detection as the parts are maintained and inspected. Carrying out this attack represents a significant investment in intellectual capital.

The following suggestions can be used to mitigate the risks associated with hardware implant attacks:

- **Minimize exposed connections on PCBs.** Exposed connections on a PCB are a prime location for placement of hardware implants. For example, if there is direct access to data buses, an implant could be designed to spoof messages and data on the bus to cause cyber effects. Removing any exposed connections, including exposed pins of an IC, will greatly reduce the feasibility of successful hardware implant installation.
- **Limit physical access to devices.** Limiting and monitoring access to devices starting from manufacturing and continuing throughout the operational use on an aircraft limits the opportunities of an attacker to install a hardware implant.
- **Implement secure firmware and software.** Implementing software and firmware that is intrinsically cyber resilient can reduce the opportunities for the insertion of malicious hardware implants. That is because high assurance software limits the places in the system where triggerable effects can be created. Chapter 13 described specific techniques that can be used to develop cyber resilient software.
- **Randomized inspection of devices.** In the event that an implant is installed on a hardware subsystem, the signs of tampering may be evident to trained personnel. Continuous randomized inspection of devices during and after production can reveal the presence of implants and limit the duration of exposure. Additionally, maintenance procedures and repair depots should implement routine inspections looking for hardware implants.

14.5.5 Hardware Trojan Attacks

A hardware trojan is a malicious design modification that can be implemented inside an IC or PCB during the design and fabrication process. While this is conceptually similar to hardware implants, the trojan is built into the part itself. Hardware subjected to a trojan attack could be maliciously altered to modify functionality, lower performance, leak sensitive information, or exhibit reduced reliability.

Hardware trojans are difficult to detect because there is no obvious physical indication of their presence. A skilled trojan designer will make it look like a part of the intended system or

otherwise keep it well-hidden. Since these attacks are executed at the hardware level, software countermeasures are most likely ineffective in mitigating the threat. Hardware trojans lie dormant in the system until activated by some triggering condition. That is likely the first indication of its presence.

Detailed understanding of the design is required to create an effective hardware trojan. The most likely insertion points are in the design process itself, untrusted foundries, IP vendors, computer-aided design tools, and design facilities. A trojan can potentially be inserted at any time in the design and fabrication process. This is one of the reasons why the supply chain management practices recommended in Chapter 15 are critically important.

While hardware trojans are relatively easy to insert by someone legitimately involved with the design flow, it is also possible for an adversary who lacks that access. For example, an adversary can identify the circuit functionality via reverse engineering and leverage that insight to fabricate a malicious counterfeit part with the addition of a trojan. The counterfeit part is then injected into the supply chain.

At the IC and the PCB levels, the source at the register-transfer level (RTL) and the schematic are the simplest methods for reverse engineering that enables the design and insertion of a trojan into a real or counterfeit part. Consequently, it is very important to protect those design artifacts to make reverse engineering significantly more difficult and time consuming.

The following suggestions can be used to mitigate the risks associated with hardware trojan attacks:

- **Obfuscation.** Inserting locking mechanisms into the original design increases the difficulty of accessing it. That makes it more difficult for adversaries to characterize and understand the design. Alternatively, the design's intent can be concealed by adding dummy circuits, hiding traces, or designing elements to camouflage the purpose of all parts of the IC. This type of technique can hinder attackers from extracting a gate-level netlist in an IC. It also can prevent direct reverse-engineering from a schematic. This requires relatively minor effort during the design phase.
- **Limit access to design documentation.** Only trusted vendors and design facilities should be used when outsourcing the various steps of the design flow. Access to documentation should only be extended to those who require that information to perform their tasks. Stringent access control mechanisms decrease the probability of inadvertent or intentional disclosure to an adversary.

14.6 Additional Methods of Hardware Assurance

Hardware security is a critical aspect of aviation systems that requires a holistic and comprehensive approach to defend against potential attacks. It takes more than the mitigation strategies in response to the adversarial techniques previously described in this chapter. This section introduces additional best practices and consolidates those recommendations as part of a multi-layered approach to hardware assurance.

14.6.1 Hardware Assurance and Supply Chain Risks

Hardware assurance within aviation systems is inextricably linked to supply chain management. As electronic components are designed and selected, it is critically important to consider cyber threats against them and to implement security controls and best practices to mitigate those risks.

That necessarily includes understanding the potential dangers lurking within the full design flow of PCBs, and ICs. There is ample opportunity from a host of individuals and organizations to intentionally or inadvertently create security vulnerabilities or weaknesses in electronic parts. Consequently, it is important to select reputable sources and perform verification testing to identify unexpected or malicious behavior.

These concerns are amplified by the continued globalization and modularity of modern supply chains. Trust issues in the entire supply chain includes risks of counterfeit and potentially malicious parts. Hardware assurance for small electronic components is an incredibly challenging task due to their intricate designs, complexity, and the sheer scale of integration in modern aviation systems. Despite these challenges, neglecting hardware assurance is not an option.

14.6.1.1 PCB Supply Chain Security

Hardware assurance is essential for PCBs. That requires particular attention to how the boards are manufactured, acquired, and managed within the supply chain.

As previously described in this chapter, protection of IP is crucial throughout the design process. The IEEE Standard 1735-2023 (72) provides comprehensive guidance on technical protection measures for electronic design IP and addresses the risks associated with unauthorized access. One of the key recommendations in that standard is the use of encryption. It also defines a common markup language format for encrypted IP.

Complete protection against reverse engineering of PCBs is challenging. However, using anti-reverse engineering techniques can sometimes make the process prohibitively expensive and time-consuming for an attacker. The following strategies are specifically recommended in support of that objective:

- **Unmarked ICs.** Unmarked ICs support protection of IP on PCBs by making it more difficult for attackers to identify and target specific components for reverse engineering or tampering. Unmarked ICs do not provide any visible indication of their functionality. That makes it harder for attackers to determine which components are important.
- **Unmarked silkscreens.** Information that is typically printed on PCB silkscreens, such as part numbers, reference designators, and other identifiers, should be suppressed to protect IP. That makes it more difficult for attackers to identify and target specific components for reverse engineering or tampering. It also makes it more challenging for an attacker to determine the interconnections of the components.
- **Eliminate debug ports.** Debug ports on PCBs can expose IP to potential threats. Elimination of those access points reduces the risk of unauthorized access to the data.

- **Ball Grid Array (BGA) components.** Using BGA components protects IP on a PCB by making it difficult to probe signals, as they do not have exposed pins. This makes the design more challenging for attackers to extract or modify the IP.
- **Tamper-proof fittings.** Custom screw shapes and tamper-proof screws can protect IP on a PCB by making it more difficult for attackers to physically access the board and modify or extract sensitive information.
- **Adhesively bonded enclosures.** Adhesively bonded enclosures can protect IP on a PCB by providing a physical barrier that prevents attackers from accessing the board and modifying or extracting sensitive information.
- **Full potting PCB casings.** Full potting is a process of filling a complete electronic assembly with a solid compound. The technique is frequently used to exclude water and to increase resistance to shocks and vibrations, which can be beneficial in aviation platforms. Full potting is also used to prevent reverse engineering as it protects the PCB and ICs from external access.
- **Hidden signal traces.** The signal traces on a PCB should be routed only on inner layers of the PCB using blind and buried vias. That prevents physical access and can also hide the traces in the design from visual inspection.

Hardware assurance is a critical aspect of PCB design and manufacturing. That is particularly true for aviation systems where security and safety are of the utmost importance. It is essential to embrace the strategies that make reverse engineering more difficult, expensive, and time-consuming. To protect IP, it is necessary to implement these best practices and mitigations throughout the design process and as the supply chain is managed.

14.6.1.2 Electronic Component Supply Chain Security

Protecting the electronic components in aviation systems is critically important due to their essential role in aircraft safety, reliability, and performance. Those tiny parts control vital functions including navigation, communication, and flight control. That makes them attractive targets for malicious attacks, unauthorized access, and tampering. Compromised hardware devices can lead to catastrophic failures, jeopardizing mission execution, the aircraft, and even human life.

When developing ICs and PLDs for aircraft, it is an imperative to exercise supply chain best practices when using third-party IP cores. It is essential that trustworthy and reputable sources are selected. Additionally, those designs must be inspected for potential hardware trojans or backdoors.

Functional testing and code coverage analysis should be performed to assess the thoroughness of validation testing for hardware. Additionally, formal verification techniques which can be prohibitively costly for very complex software systems are potentially more feasible for validation of the security properties in register-transfer level (RTL) designs.

Bitstream encryption using techniques recommended by IEEE Standard 1735-2023 (72) should be considered as technical protection measures of electronic design IP. For use in aviation

systems, hardware designers should select ICs and PLDs that implement security features and properties.

It is critical to use these best practices throughout the design process and the supply chain to prevent the leakage of IC and PLD bitstreams. Prioritizing cybersecurity and protection of the IP in aviation system hardware is essential to prevent potential threats and ensure the safe and reliable operation of these systems.

14.6.2 Mutual Authentication

Mutual authentication is a critical security control performed between hardware components of a system. This process ensures that both are authentic before data or information exchanges occur. In an aircraft platform, mutual authentication can be implemented with a challenge-response mechanism using a shared secret or asymmetric keying material.

Mutual authentication can help prevent unauthorized access, tampering, and other forms of hardware-based attacks that can compromise the safety and security of aviation systems. Implementing mutual authentication is a technique that should be considered to enhance the hardware security posture, protect sensitive information, and ensure the safe and reliable operation of the aircraft.

14.6.3 Hardware Root of Trust

A Hardware Root of Trust is a fundamental concept in system security that provides a foundation for building a chain of trust. It is a secure hardware-based mechanism that provides system resistance against certain types of software attacks. A hardware root of trust is a relatively strong method used to validate the authenticity and integrity of the system's boot processes and other critical security functions.

Establishing a hardware root of trust prevents attackers from inserting themselves into the trust chain as cyber-enabled devices are powered on. It is a protection mechanism that makes it very difficult for adversaries to execute unauthorized malicious code. Without a hardware root of trust an attacker can hijack the device as it is booted and take control of the software execution and bypass software-based security controls. That can have catastrophic consequences for aviation systems.

However, implementing hardware root of trust on an aircraft platform can be a complex challenge. The linchpin is a hardware-based security module that provides secure storage for cryptographic keys, digital signatures, and other security-related data. The security module is integrated with the hardware boot process to ensure that only authenticated and authorized code is executed. The chain of trust is established through verification of each stage of the boot process using digital signatures or cryptographic keys stored in the security module.

Hardware root of trust faces the same fundamental challenges associated with the substantiated integrity software-based approach described in section 7.5.1.5. Implementation of secure key management including generation and storage are complicated within a system that does not have a persistent network connection to public key infrastructure services.

Ultimately, the digital signatures and cryptographic keying material must be securely stored within the hardware security module. Those are the keys to the kingdom that is an attractive target for the IP motivated attacks described throughout this chapter. Adding mechanisms to update those materials in the event of suspected compromise only adds more vectors for potential attacks. The challenge of implementing hardware root of trust in aviation systems is immense.

A hardware root of trust implementation requires more initial investment than software-based approaches. The specialized hardware security module increases design costs and integration complexity. Additionally, under some circumstances the hardware may need to be replaced if believed to be compromised. That is expensive.

Since software-based solutions don't require specialized hardware, the initial investment is not as significant. On the surface, the flexibility to update keying material is an advantage for software methods. However, that comes at the cost of adding vectors to the attack surface that are easier to defeat than hardware approaches.

A hardware root of trust can also be used to detect and respond to attempted cyber-attacks against the hardware. By monitoring the boot process and critical security functions, it can detect anomalies which are indicative of unauthorized access attempts. That can be used to alert aircraft operators and maintainers of an attempted incursion.

Establishing a secure root of trust in aviation systems is an advanced best practice that can be used to protect the confidentiality and integrity of data on sensitive platforms. Whether to include this security control at all is a complex decision influenced heavily by the intended purpose and sensitivity of the aircraft.

14.7 Summary and Additional Resources

Hardware is the foundation of the design of an aviation system. Consequently, ensuring that hardware components are secure is crucial to the overall resiliency of any aircraft subjected to malicious cyber activity. Many of the attacks presented in this chapter do not have mitigations that completely close the potential for compromise. Despite that, it is critical to understand the risks and the level of trust necessary to make informed decisions during the hardware design process.

As aviation system designers select and create the hardware implementations on their platforms, the following key considerations should be used to guide and inform their decisions:

- Adding mitigations for hardware attacks after the design has already been completed is costly and impractical. Therefore, hardware security must be kept in mind, starting from inception through all stages of design.
- Since hardware attacks rely on some amount of physical access to the device, it is of paramount importance to limit who has access to these devices to only those who need it. The same applies for design documentation and details since developing a hardware attack without an understanding of the design can be incredibly time consuming and costly.

- Adversaries with enough time, knowledge, and skill will eventually be able to bypass most commonly used security features. Implementing security mechanisms may also limit the ease of use, speed, or maintainability for a device throughout its lifecycle. It is important to find a balance that meets the functionality needs of the customer while remaining secure. Therefore, it is up to the design team to determine the acceptable amount of risk associated with a device and implement mitigations to meet those requirements.

The following additional resources are recommended for more in-depth information about the topics and concepts presented in this chapter.

- **DoD Anti-Tamper Executive Agent Website. (73)**

The Anti-Tamper (AT) Executive Agent publishes an information page that describes the systems engineering activities intended to prevent and/or delay exploitation of Critical Program Information (CPI) in US weapons systems. These activities involve the entire life cycle of systems acquisition, including research, design, development, implementation, and testing of AT measures. Properly employed, AT adds significant longevity to CPI by deterring efforts to reverse-engineer, exploit, or develop countermeasures against a system or system component.

- **IEEE Recommended Practice for Encryption and Management of Electronic Design Intellectual Property (IP), Institute of Electrical and Electronics Engineers (IEEE), Std 1735™-2023, 29 June, 2023. (72)**

The IEEE Std 1735-2023 provides comprehensive guidance on technical protection measures for electronic design intellectual property (IP). This standard addresses the risks associated with the distribution of IP, which can lead to unauthorized use and the erosion of the investment in its creation. Key protective measures include encryption, the specification and management of usage rights granted by IP producers, and methods for integrating license verification. By implementing these measures, the standard aims to safeguard the integrity and value of electronic design IP against malicious exploitation.

Chapter 15

Cyber Supply Chain Management

This chapter was written by Jeff Chang and Tambre Paster and edited by Teresa Merklin.

Aviation system cyber resiliency is inextricably linked to the supply chain. In fact, the initiative that ultimately culminated with this FSAD Guidebook was driven by a need to help all participants in the vast supply chain for aircraft fully understand their roles and responsibilities essential for the development, production, and sustainment of cyber resilient platforms.

In previous chapters, the fact that this book was intentionally structured around the fundamental principles of cybersecurity systems engineering rather than specific details of any assessment and authorization methodology was repeatedly emphasized. That is because the foundational engineering must be performed regardless of what any standard may specify. Cyber resiliency isn't achieved by only doing the bare minimum.

That same philosophy applies to the cyber supply chain risk management best practices discussed in this chapter. Supply chain security is an evolving competency, just like the threats it seeks to address. Organizations in the aviation supply chain that master and lean forward into the fundamental principles will outperform all others who regard the standards as checklists of minimum requirements. The development of cyber resilient aircraft requires that all stakeholders go above and beyond in every aspect of cyber supply chain risk management.

Securing the supply chain is a critical aspect of aircraft system cyber resiliency. While this chapter is primarily targeted at individuals and organizations who are officially assigned to supply chain management roles, the reality is that these responsibilities apply more broadly. Just as cybersecurity is everybody's job, so is cyber secure supply chain risk management.

15.1 Introduction

Every cyber enabled component integrated on an aviation platform has a foundational role in the resiliency posture of the aircraft. The individual parts don't simply materialize out of thin air. Every organization that designs, produces, and supplies aviation hardware and software to the aircraft is a critical part of the aviation supply chain.

The conceptual model of the aviation mission stack introduced in Chapter 1 is reprinted here as Figure 21. That conceptual model encapsulates the complex relationships between various aspects of the supply chain and the ultimate mission of the aircraft. Each layer in the stack simultaneously represents an opportunity to support cyber resiliency as well as a potential threat vector of vulnerability or attack.

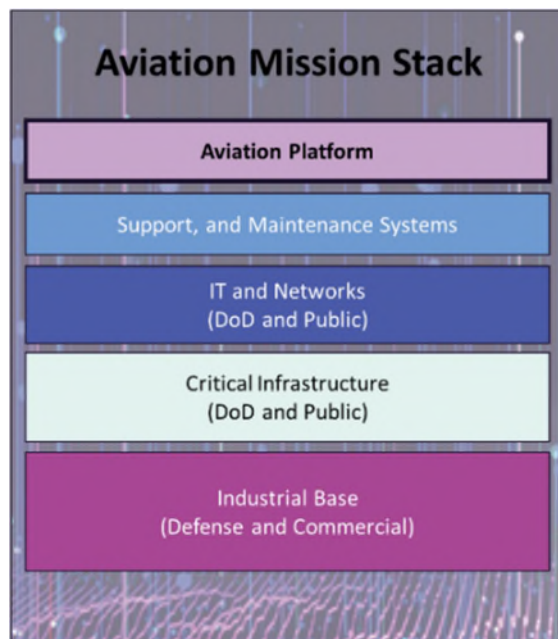


Figure 21. Aviation Mission Stack

The aircraft is positioned at the top of the aviation mission stack. However, cyber resiliency for the flight platform cannot be achieved in isolation. The aviation mission stack is essential for understanding the full scope of what is required to achieve cyber resiliency objectives throughout the supply chain. Each layer in the stack simultaneously represents an opportunity to support cyber resiliency as well as a potential vector for hostile cyber-attack.

The aviation mission stack includes support and maintenance systems, Information Technology (IT) networks, critical infrastructure, and the aviation industrial base. The aviation mission stack is the perfect lens through which to frame the enormous scope of the supply chain that produces and supports flight platforms.

It is imperative for all members of the aviation industrial base to implement best practices that secure the parts, software, and services supplied to aircraft. At the same time, those manufacturers must be aware of the impact their own suppliers also have on the security of deliverable products.

Cybersecurity of every aviation system relies on locating and eliminating potential points of vulnerability. The vast supply chain for complex modern aircraft contains many potential weak points. It is the duty and responsibility of the entire industrial base to ensure that products installed in aviation systems are not the source of cyber resiliency failures.

The goal of this chapter is to provide a high-level summary of the broad topic of supply chain security. It highlights the *what, why, and how* the discipline is performed. While much of the material in this chapter is specific to the DoD, many of these concepts are equally applicable to commercial aviation systems.

This chapter is intended to help suppliers connect the dots between key concepts, published articles, standards, and best practices for managing and securing the supply chain. Due to the inordinate amount of changing supply chain security requirements, best practices, and special publications, it can be overwhelming for many suppliers to stay up to date and feel confident about their responsibilities relating to supply chain cybersecurity.

The National Institute of Standards and Technology (NIST) has published guidance in a special publication titled "Cybersecurity Supply Chain Risk Management Practices for Systems and Organizations." NIST Special Publication (SP) 800-161 (74) provides essential guidelines for managing cybersecurity risks across the supply chain. The fundamental concepts of cybersecurity supply chain risk management (C-SCRM) were originally introduced in that document.

To illustrate the dynamic and evolutionary nature of the C-SCRM guidance, NIST SP 800-161 (74) was updated in May of 2022 with enhanced C-SCRM requirements and recommendations. Key updates in that revision broadened the scope to include operational technology (OT) and Internet of Things (IoT) in addition to traditional information and communications systems. This revision also updated the risk management practices and noted that executive leaders are responsible for managing cybersecurity risks in the supply chain.

It is important to understand that C-SCRM is an evolving competency, just like the threats it seeks to address. All stakeholders in the aviation supply chain must stay abreast of their interdependencies and recognize that multiple vendors and service providers all have an impact on the cyber resiliency posture of every delivered product.

15.2 The Full Significance of the Aviation Supply Chain

It is important to recognize that the supply chain risks faced by aviation systems include the potential for harm or compromise that may arise from suppliers, their supply chains, their products, and their services. These types of risks are the results of threats that exploit vulnerabilities or exposures that traverse or are within the supply chain itself.

Securing the supply chain depends on identifying, assessing, and mitigating all types of risk associated with an organization's procurement and production processes. That includes 3rd party suppliers, manufacturers, integrators, logistics, and transportation services. Supply chain security is more than just protecting the confidentiality of sensitive information associated with hardware, software, or firmware. It is equally important to ensure the integrity and trustworthiness of each item. Cyber resiliency depends on components that are authentic, have not been maliciously altered, and function as designed.

Supply Chain Risk Management (SCRM) applies best practices for mitigating a broad range of supply chain risks. Cybersecurity Supply Chain Risk Management (C-SCRM) is an

augmentation of SCRM that focuses specifically on mitigating cyber-based supply chain risks. C-SCRM is most effective when built on a solid foundation of SCRM principles and practices.

The phrase “flow downs” refers to the process of extending SCRM requirements and expectations to lower-tier suppliers in the supply chain. This is an important aspect of SCRM because it helps ensure that all suppliers, regardless of their tier, are aware of and adhere to the same security and risk management standards.

SCRM flow downs typically involve establishing clear expectations and requirements for suppliers, communicating those requirements effectively, and verifying the necessary controls and processes to manage supply chain risks have been implemented. This may include conducting risk assessments, establishing supplier ratings or scores, and implementing monitoring and reporting mechanisms to track and mitigate supply chain risks.

Building more secure and resilient supply chains requires effectively managing SCRM flow downs. It reduces the risk of supply chain disruptions, and better protects critical assets and mission-critical functions.

Figure 22 is an example set of SCRM related flow downs under the overall DoD SCRM definition. That diagram illustrates a requirement for a “Multi-tier SCRM” plan for mitigating “Other Risk Categories.”

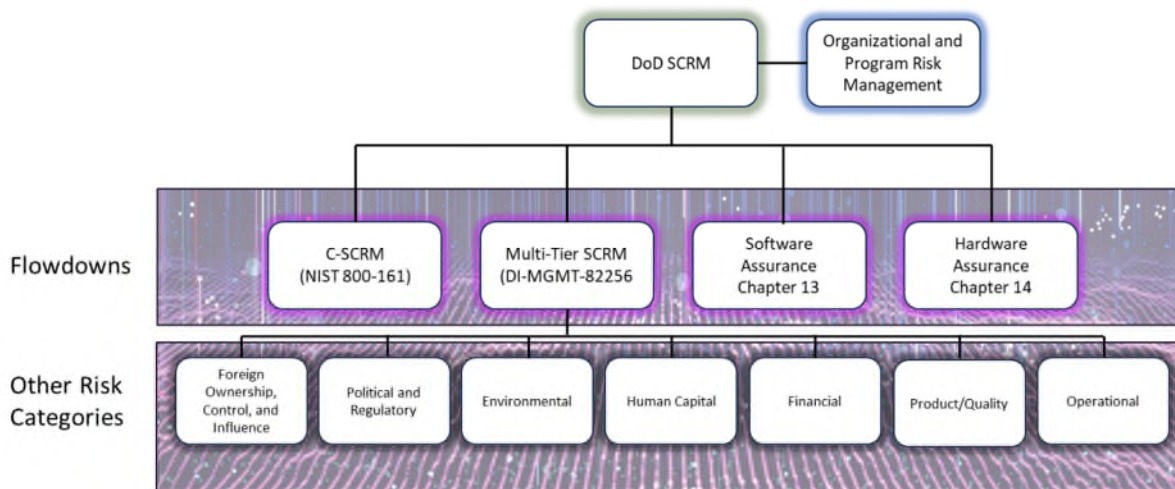


Figure 22. Supply Chain Risk Management Flow Downs

A notional contract paragraph for that requirement might read as follows:

“The Contractor shall develop, implement, and maintain a Supply Chain Risk Management (SCRM) plan for monitoring, reporting, and controlling risks and opportunities associated with executing the requirements. Additionally, the Contractor shall address key supply chain activity processes including but not limited to supplier selection, supplier evaluation / audits, supplier rating system, receiving test and inspection, conditional source approval, source inspections, procurement, and metrics in

the SCRM plan. The Contractor shall deliver an SCRM Plan in accordance with DI-MGMT-82256.”

This new requirement for an SCRM plan directs prime contractors to proactively map out their multi-tier supply chain network and to perform an assessment of the risks associated with every illuminated tier 1 and sub-tier supplier. That also includes a daily monitoring requirement. This is a new fundamental supply chain risk management practice that not only calls for a specific SCRM plan, but also some method of monitoring and alert notifications.

To understand the significance of supply chain security, it is important to identify which (or *what*) critical assets require protection and *why* it is important. Mission-critical functions are those functions of the system that would likely lead to mission failure or degradation if corrupted or disabled. (75)

There are many valid methods for identifying the mission critical functions and components of a system. Section 6.5 in Chapter 6 described the use of Mission Impact Assessment for performing that identification. It should be noted that there are many additional published methods and standards as well. Another public method for finding mission critical functions is performance of a criticality analysis as defined by NIST IR 8179. (76) That defines a methodology for identifying and prioritizing information systems and components to increase understanding of an organization’s assets. (77)

As NIST SP 800-161 sums it up, “The level of exposure to cybersecurity risks throughout the supply chain depends largely on the relationship between the products and services provided and the criticality of the missions, business processes, and systems that they support.” (74)

Critical assets can also include sensitive information associated with programs, systems, sub-systems, components, and subcontracted services. For example, within the scope of DoD contracts, Controlled Technical Information (CTI) is a category of Controlled Unclassified Information (CUI). That includes research and engineering data, engineering drawings, 3D models, specifications, standards, manuals, technical reports, studies, computer software executable code, and source code. Moreover, per National Archives and Records Administration (NARA), CTI related to military or space application is subject to controls on the access, use, reproduction, modification, performance, display, release, disclosure, or dissemination. (78)

Collectively, these critical assets are endangered by threats and vulnerabilities such as malicious intent of foreign governments, poor manufacturing, lax development practices, counterfeit products, product tampering, theft, insider threats, malicious software, and other threats at all levels of the Aviation Mission Stack as identified in Figure 21.

When these threats and vulnerabilities are not identified and properly managed or mitigated, foreign and domestic bad actors can exploit them to cause significant harm and disruption to U.S. and allied nations’ national security and critical infrastructure. The resulting exposure and risk must be assessed and mitigated by implementing appropriate safeguards and countermeasures. The gravity of the ever-present danger of supply chain threats was previously discussed in detail in Chapter 2.

In 2021, a study published by the U.S. House of Representatives' Defense Critical Supply Chain Task Force (79) highlighted the vulnerabilities in the security of the U.S. supply chain made obvious by the COVID-19 pandemic. The report called out two specific supply-chain concerns which were illuminated by that global crisis:

- 1) The DoD has a critical lack of visibility into the defense supply chain which impedes the understanding of its vulnerabilities and surge capacity to handle the next crisis.
- 2) The United States has an over-reliance on foreign made critical microelectronics, particularly from China, which could be used as leverage against the United States in the future.

The task force was chartered to review, identify, and analyze supply chain threats and vulnerabilities. It ultimately recommended that the DoD treat supply chain security as a defense strategic priority. To do that, visibility into the defense supply chain is required to understand its vulnerabilities and to develop risk mitigation strategies. Additionally, the DoD must reduce its reliance on adversaries for resources and manufacturing. Those findings are particularly pertinent to the aviation supply chain for both military and commercial platforms.

A similar call to action was issued by NIST in Revision 2 of SP 800-37, the Risk Management Framework for Information Systems and Organizations. (80) That regulation states that “The significant increase in the complexity of the hardware, software, firmware, and systems within the public and private sectors, including the U.S. critical infrastructure, represents a significant increase in attack surface that can be exploited by adversaries. Moreover, adversaries are using the supply chain as an attack vector and effective means of penetrating our systems, compromising the integrity of system elements, and gaining access to critical assets.”

Supply chain security directly impacts both military and non-military assets. Consequently, the DoD and commercial aviation service providers are demanding more visibility into the supply chain. Contracting decisions are increasingly influenced by which suppliers the acquiring organizations trust, and which ones it does not. As a result, future success of any supplier in the aviation industry hinges on being trustworthy in the eyes of the purchaser.

15.3 Specific Concerns for DoD Aviation

The DoD has become increasingly aware of the challenges associated with increasingly obscured visibility and understanding of how aircraft platforms are acquired, developed, integrated, and deployed. The processes, procedures, and practices used to assure the integrity, security, resilience and quality of aviation products and services have been obscured by the complexities and lack of transparency throughout the downstream supply chain tiers. The same challenge is also apparent in commercial aviation systems.

Capable and motivated adversaries will most certainly seek to exploit vulnerabilities in the DoD aviation platform supply chain. That is in anticipation and recognition that cyber warfare will be a significant factor in future armed conflicts. Since the security of any supply chain is only as strong as its weakest link, it is imperative to gain greater visibility into the full multi-tiered DoD

supply chain to identify, assess, and address the risks. This is particularly true of cybersecurity risks associated with the suppliers that have insufficient cybersecurity posture and hygiene.

Figure 23 illustrates the reduction of visibility that occurs through multiple layers of the supply chain. That diagram was derived from the NIST Cybersecurity Supply Chain Risk Management Practices for Systems and Organizations. (74) Additionally, insight is frequently hampered by numerous legal, contractual, and logistical data sharing challenges. Full traceability of any multi-tiered supply chain is extremely challenging for the acquiring enterprise.

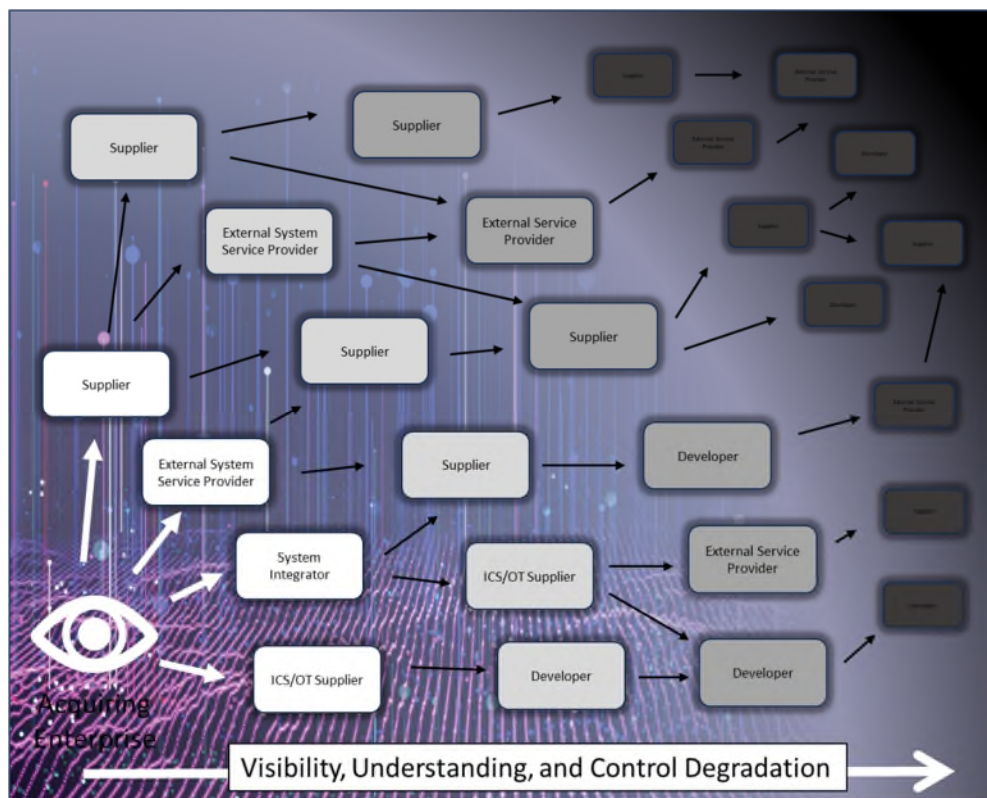


Figure 23. Enterprise Visibility, Understanding, and Control of a Supply Chain. (74)

As a practical example, ensuring that a mission-critical military aircraft is delivered uncompromised requires identification of every subsystem, component, software, and service on the platform. The provenance, including the place of origin and history of ownership, must be assessed. Additionally, verification or certification of the security, safety, integrity, quality, reliability, authenticity, and trustworthiness of every part must be determined. This is required at every tier of the supply chain.

A Bill of Materials (BOM) is a comprehensive list of all the parts, assemblies, and other items included in a product. The BOM is a critical document for supply chain management as it provides detailed information about components including part numbers, descriptions, and supplier information. For cyber-enabled parts, it is particularly important that the BOM include provenance information.

Additionally, a comprehensive BOM is needed for suppliers at every tier of the supply chain. Sharing that detailed information can be challenging due to the shifting lines between partners and competitors from program to program. BOM data may be competition-sensitive and possibly reveal proprietary information. Clear guidelines and protocols are essential to balance the necessity of data sharing with the need to protect intellectual property and to preserve competitive advantage.

15.4 Supply Chain Risk

Supply chain risk is defined by Federal Acquisition Regulation (FAR) 252.239-7018 as “the risk that an adversary may sabotage, maliciously introduce unwanted function, or otherwise subvert the design, integrity, manufacturing, production, distribution, installation, operation, or maintenance of a system so as to surveil, deny, disrupt, or otherwise degrade the function, use, or operation of such system.” (81)

The U.S. National Counterintelligence and Security Center (NCSC) characterizes supply chain risk as a function of threat, vulnerability, and consequence. (82) A supply chain threat is specific and credible information indicating that adversaries might target a component, system, or service. It is important to emphasize that this concern is separate from intelligence or indications that any adversary is actively doing that. Proactive engineering analysis identifying the components, systems, or services where targeting would be beneficial is an excellent way to identify threats in the supply chain.

A supply chain vulnerability is a weakness that is either inherent in a component, system, service, or has been introduced by an outside agent. That includes the potential for counterfeit or maliciously altered components, which could lead to system failures or accidents. Vulnerabilities could also arise from inadvertent weaknesses in systems that an adversary has discovered through reverse engineering or other data sources.

Supply chain risk occurs when the capability and intention of an adversary align with the opportunity to exploit a vulnerability. Once again it is important to emphasize that supply chain risk exists whether an adversary is actively pursuing that vector or not.

One consequence of a successful exploit or attack might result in adversarial exfiltration of intellectual property or sensitive government data. Even worse, it might allow an adversary to surveil, deny, disrupt, or otherwise degrade a component, system, or service. Understanding these risk elements and the relationships among them are key to managing supply chain risks. Specifically, mitigation of risks should be driven by specific threats and vulnerabilities. (82)

15.5 Supply Chain Risk in the Program Protection Plan

A Program Protection Plan (PPP) is a comprehensive outline of the measures and procedures needed to defend the program against various threats and risks. That necessarily includes the ones associated with the supply chain. In addition to detailed analysis of potential risks to the program, the PPP identifies the specific measures to address them. It is a critical artifact that describes the management and mitigation of all risks throughout the program's lifecycle.

While PPPs are a part of the broader framework for managing defense acquisition programs, it raises the question of exactly what requires protection. One source is DoD Instruction (DoDI) 5000.02 (83) which asserts the necessity to protect certain “critical technologies, components, and sensitive information” throughout the program lifecycle. That instruction also highlights the need to identify, assess and mitigate supply chain risks for DoD programs, including aircraft.

Figure 24 organizes some of the key policies and guidance into a framework of technology, components, and information. It is a great model to understand what to protect, which drives identification of the protection measures and goals. These high-level objectives should be written into requirements statements and flowed throughout the supply chain.

For DoD systems, critical technologies include critical program information (CPI) which is a capability that contributes to a decisive military advantage. CPI is any data or knowledge that is vital to the success, security, or mission of a specific program or project. It is often sensitive in nature and may be subject to data classification, and other security measures to protect it from unauthorized access, use, or disclosure. The goal is to prevent the loss of intellectual property (IP) and CPI.

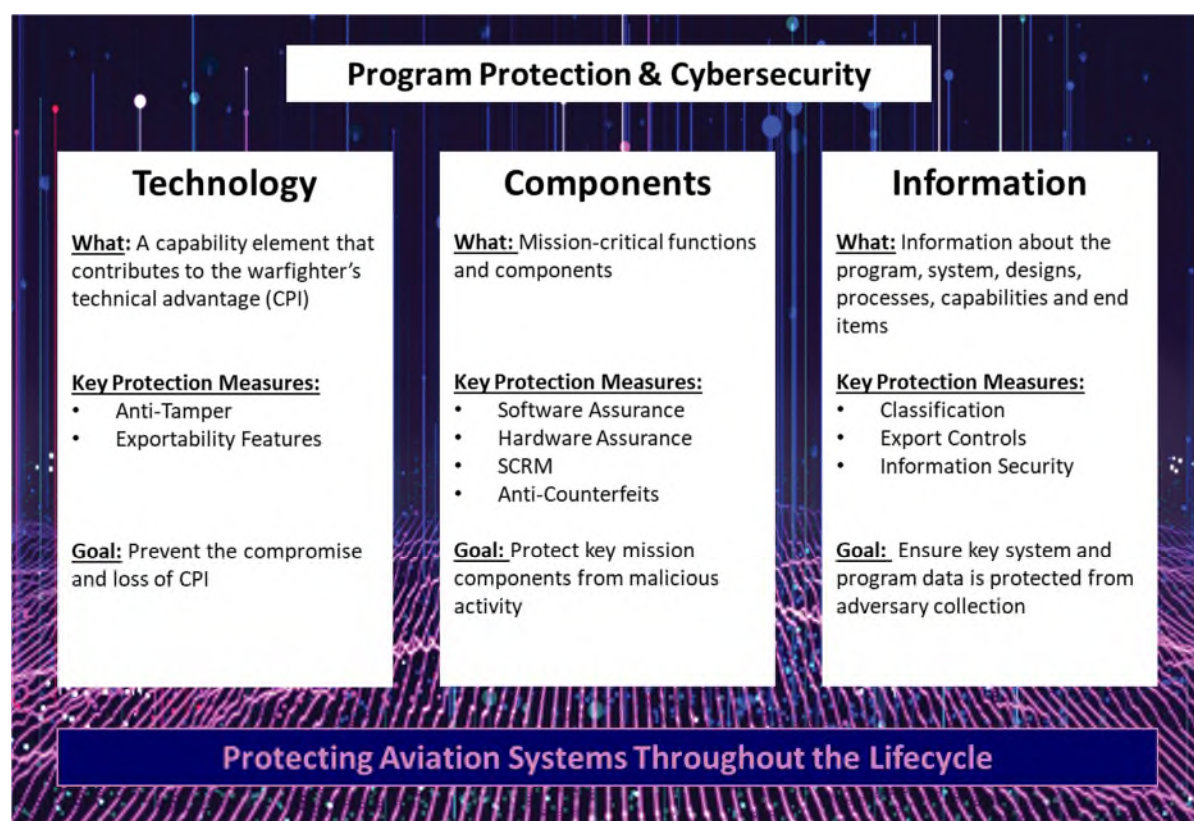


Figure 24. What Are We Protecting? (84)

Critical components are those which fulfill a mission-essential purpose within the system. That can include software and hardware devices that require protection throughout the supply chain. The goal is to protect key mission components from cyber adversity.

Sensitive information includes program data, designs, processes, and capabilities. It can also sometimes include proprietary data, Covered Defense Information (CDI) or Controlled Unclassified Information (CUI) associated with the critical technologies and components. The goal is to ensure that key system and program data is protected from disclosure.

Additionally, within the context of program and sensitive information protection, the Federal Acquisition Supply Chain Security Action (FASCA) of 2018 (Title II of SECURE Technology Act) requires federal agencies to perform Supply Chain Risk Assessment (SCRA) and use NIST standards and guidance when assessing and developing mitigation strategies to address supply chain risks. (85) These mitigation strategies should also be flowed as contractual requirements to the prime and sub-contractors/suppliers to address supply chain risks.

15.6 Potential Sources of Supply Chain Risk

There are many sources of supply chain risk, and the list continues to grow as new ways to disrupt, attack, and corrupt the supply chain are discovered and devised. This section presents a high-level overview of the types of risks that should be considered when protecting cyber-enabled systems within the supply chain.

15.6.1 Data and Cyber Risk

Data supply chain risks can be broadly categorized into four types: exposure and theft of sensitive data, insider threats, third-party cyber risk, and software/hardware risk. It is vital to understand these risks to identify the appropriate protection measures.

- **Exposure and Theft of Sensitive Data.** Sensitive data is at risk of being exposed, stolen, or corrupted. It is important to understand the sensitive data associated with the system, and to take the appropriate measures to protect it. NIST's framework "FIPS 199, Standards for Security Categorization of Federal Information and Information Systems", can be used to assess data sensitivity and security categorization in terms of confidentiality, integrity and availability, commonly referred to as the C-I-A triad. (86) (87) The C-I-A triad was previously described in Section 5.2 of this FSAD Guidebook.
- **Insider Threats.** People are frequently a critical factor in cyber-attacks because targeting *human* weaknesses instead of *digital* ones can be very effective. While the risk of human compromise is frequently more closely associated with enterprise IT systems, it cannot be discounted for aviation platforms. Section 2.4.2 provided more information on conceptualizing the insider threat.
- **Third-Party Cyber Risk.** The companies that aviation suppliers do business with can create significant risk. A supplier with a strong cybersecurity posture can sometimes be compromised through its less well-defended vendors and subcontractors. Consequently, companies must anticipate potential threats and risks through their relationships to partners, suppliers, and service providers. (88)
- **Software/Hardware Risk.** Running untrustworthy software or hardware can create significant threat vectors. The U.S. government has banned some untrusted products from use in the DIB. For any aviation system, vigilance and skepticism when procuring

hardware and software should be exercised. As part of their situational awareness, companies must ensure that their procurement personnel are aware of this risk.

15.6.2 Hardware Manipulation/Product Tampering Risk

Supply chain risks associated with hardware manipulation and product tampering can have severe consequences for aviation platforms. Such tampering can be difficult to detect and may remain undiscovered until the component is in use. That increases the potential of catastrophic consequences.

Hardware and product tampering can also be carried out within the supply chain through interdiction of parts as they are transported to production and sustainment. Additionally, risks associated with counterfeit components, including hardware and software, could also lead to failure and catastrophic consequences. Product tampering could also occur by an adversary who gains access to the development environment and injects vulnerabilities or malicious software at the source.

To mitigate these risks, it is essential to establish robust security measures and protocols throughout the entire supply chain. That includes stringent vendor selection, ongoing monitoring, and regular audits. Additionally, implementing secure manufacturing processes, such as using tamper-evident packaging and hardware root of trust, can help detect and prevent unauthorized modifications. Chapter 14 describes specific techniques for addressing hardware manipulation or product tampering risks.

15.6.3 Geopolitical Risk

Geopolitical risk poses significant challenges to the aviation supply chain, as it can disrupt the flow of goods and services and cause delays. These risks can arise from political instability, trade disputes, sanctions, and military conflicts. For example, natural disasters, terrorist attacks, or civil unrest in a region where a critical supplier is located can disrupt the availability of raw materials or components. Similarly, trade disputes or sanctions imposed by governments can restrict the flow of goods and services across borders.

To mitigate geopolitical risks, organizations should conduct regular risk assessments of their supply chains, identify critical suppliers and components, and develop contingency plans to address potential disruptions. Additionally, diversifying the supplier base and investing in supply chain visibility and monitoring can help organizations anticipate and respond to geopolitical risks.

Per the Global Risk Institute, “Global affairs are in a state of disruptive transition. Trends like the growth of emerging economies, the rise of populism, and evolving security threats are increasing multipolarity in international relations, upsetting trade and investment flows, steering markets, and shaping regulation. Multilateral organizations, national and subnational governments, and civil society groups often hold competing interests and objectives. Of greatest consequence, the relationships among these constituencies can shape the business environment in which financial institutions operate.” (89)

15.6.4 Procurement Risk

Procurement is a critical function in managing supply chain risks, as the components and services sourced from suppliers can significantly impact quality. Selecting the lowest cost items may seem like an attractive option to reduce costs, but it can also introduce significant risks to the supply chain. Low-cost parts may be of inferior quality, leading to reduced reliability, increased maintenance costs, and potential safety issues. Additionally, the most competitively priced commodity items are more likely to be counterfeits.

To mitigate supply chain risks associated with procurement, organizations should adopt a risk-based approach to sourcing decisions that specifically consider criteria beyond just cost. These factors include the supplier's trustworthiness and stability.

15.6.5 Physical Risks in Transportation and Distribution

The entire aviation supply chain must have a robust verification process to ensure the supplier deliveries of hardware and software components are authentic and have not been tampered with during transit. Supply chain interdiction can occur at any point in the transportation and distribution process. That includes items that are stored in warehouses.

To mitigate these risks, it is essential to integrate cybersecurity programs with physical security measures. For software, substantiated integrity techniques are an effective way to detect and prevent illicit modification. Tamper-evident seals, holograms, or other physical security measures can be used to ensure that hardware has not been tampered with during transit.

Visibility technology can potentially be used to monitor the location and status of critical components throughout the supply chain. This can help detect potential interdiction attempts. Another mitigation strategy to consider is establishing a secure chain of custody for critical components, including the use of tamper-evident bags, containers, or other physical measures to support integrity during transit. Companies can also require suppliers to provide documentation, such as a certificate of authenticity, to verify the origin and integrity of the components.

15.6.6 Software Development Risk

The participants in the aviation supply chain who receive software from external suppliers, vendors, or use software downloaded from open-source platforms must take steps to ensure that it is authentic and has not been illicitly modified. Additionally, it is equally crucial to protect in-house developed software from malicious modifications during the design and delivery phases.

Externally acquired software comes with an inherent risk of adversarial injections into upstream software components or libraries. That same injection threat exists should an adversary gain access to the in-house development or integration environments. Software vulnerabilities are a significant part of supply chain risk. It presents a vast and powerful attack surface for cyber adversaries.

15.6.7 Sustainment and Maintenance Risk

Sustainment and maintenance risk refers to the potential dangers and uncertainties that can arise during the post-sale or post-shipment phase of the aircraft. Even after goods are sold or shipped, supply chain risk to cyber-enabled parts does not disappear. Instead, it transforms and presents itself in new ways throughout the sustainment and maintenance lifecycles.

One key risk during sustainment and maintenance is incomplete or unsecured digital records, which can make it challenging to track and manage ongoing aircraft support. Additionally, the introduction of replacement or spare parts introduces new hardware risks including malicious modification or injection of parts throughout the supply chain. Replacement parts may also be counterfeits.

The maintenance and sustainment phase also comes with software risks. Mechanisms for updating and patching software can sometimes be leveraged by malicious actors as a vector of attack to inject unauthorized or malicious software. Additionally, software patches can introduce compatibility issues, software bugs, or security vulnerabilities that weren't present in prior versions.

Diminishing manufacturing sources (DMS) and material shortages (MS) are significant challenges to the sustainment and maintenance of aviation systems. The gradual reduction or unavailability of critical components and materials will impact manufacturing and repair of aviation platforms. This can occur due to a variety of factors that includes the discontinuation of production by original equipment manufacturers and changes in supply chain dynamics. Additionally, when cyber-enabled components go end-of-life, it can be extremely impactful to aircraft platforms due to the relatively long service timespans.

The addition of new suppliers to the supply chain can also change the risk profile through expanded attack surface and potential uncertainty of the cybersecurity posture of the new vendor. Additionally, aviation maintenance and sustainment faces risk from aging components, end-of-life support, and technology obsolescence. To manage these risks effectively, a comprehensive risk management approach that considers the entire lifecycle of the system that includes maintenance and sustainment is essential.

15.6.8 Commercial Off the Shelf (COTS) Product Risk

The DoD has embraced a strategy of using Commercial-Off-The-Shelf (COTS) components to the greatest extent possible to reduce overall costs. The products that are readily available on the commercial market will usually be less expensive than custom solutions developed for military aviation platforms. In fact, many commercial components can be used for military purposes without modification.

However, the use of COTS components in DoD aviation platforms introduces new supply chain risks due to the reliance on commercial suppliers. The defense supply chain is global and modular, comprising commercial and DIB suppliers worldwide. This reliance on commercial vendors, who may not have the same cybersecurity standards as DIB suppliers, increases supply chain opacity and decreases security.

To combat these risks, companies must adopt tools and processes that ensure defense-level security while being commercially viable, low-cost, and friction-free. By empowering commercial actors to enhance their products' security, the overall security of defense platforms can be improved.

A successful risk mitigation strategy must evolve to address potential disruptions to standard operating procedures. Any organizational disruption should be evaluated using a probability and risk model that assesses the likelihood of occurrence and potential fallout. This approach enables organizations to proactively identify and manage supply chain risks associated with COTS components and ensure the security and reliability of aviation platforms.

15.6.9 Product Traceability Risk

Product traceability risk is a critical concern in ensuring the safety and reliability of complex aviation platforms. While much attention is often given to the traceability of logic-bearing devices, mechanical parts can have an equally significant impact on platform resilience. In high-value systems like a hundred-million-dollar aircraft, a two-dollar part can have far-reaching consequences if it is not genuine or fails to operate as intended. Therefore, understanding the provenance of all parts is essential to managing product traceability risk to ensure the overall safety and reliability of aviation platforms.

15.6.10 Natural Risk

In the context of supply chains, natural risk refers to the potential dangers and uncertainties that arise from unpredictable events such as natural disasters, accidents, or equipment failures. These events can disrupt the supply chain by causing delays, increasing costs, and potentially compromising the safety and reliability of aviation systems.

When assessing supply chain risk, some natural events will have extremely low probabilities, yet still deliver the most devastating consequences if they occur. That risk is particularly impactful for the defense aviation supply chain as many crucial parts and components come from a single supplier. Disruptions of production can lead to vulnerabilities and severe impact. This can have unintended consequences upstream and downstream within the supply chain.

15.6.11 Mapping Sources of Supply Chain Risk to the Aviation Mission Stack

Each of the potential sources of supply chain risk introduced in this section can be mapped to the Aviation Mission Stack that has been referenced throughout this FSAD Guidebook. Completing that exercise provides a comprehensive view of how each risk asserts itself across various tiers. By mapping supply chain risks to this framework, organizations can better understand how various risks can impact different aspects of their systems and develop targeted risk management strategies to address them.

Table 7 presents a mapping that reveals that some potential risks apply to all tiers of the Aviation Mission Stack, while others are more isolated. This example is provided as a tool that suppliers can use to perform a mapping unique to their own tier and context.

Risk Categories	Aviation Platform	Support and Maintenance Systems	IT and Networks (DOD and Public)	Critical Infrastructure (DOD and Public)	Industrial Base (Defense and Commercial)
Exposure and Theft of Sensitive Data	X	X	X	X	X
Insider Threats	X	X	X	X	X
Third-Party Cyber Risk	X	X	X	X	X
Software/Hardware Risk	X	X	X	X	X
Hardware Manipulation/Product Tampering Risk	X	X		X	X
Geopolitical Risk				X	X
Procurement Risk				X	X
Physical Risk in Transportation and Distribution	X	X		X	X
Software Development Risk	X	X			
Sustainment and Maintenance Risk	X	X	X		X
Commercial Off the Shelf (COTS) Product Risk	X	X	X		
Product Traceability Risk				X	X
Natural Risk	X	X	X	X	X

Table 7. Relationship Between the Aviation Mission Stack and Common Supply Chain Risks

15.7 Applied Cyber Supply Chain Risk Management (C-SCRM)

Supply Chain Risk Management (SCRM) is a comprehensive approach to identifying, assessing, and mitigating risks that can disrupt the supply chain. It is focused on ensuring the continuous flow of goods, services, and information from suppliers to customers. That requires managing a broad range of risks associated with the supply chain including financial, operational, reputational, and logistical factors. SCRM focuses on ensuring that all elements of the supply chain, from suppliers to customers, operate smoothly and without interruption.

Cyber Supply Chain Risk Management (C-SCRM) is a subdiscipline of SCRM that focuses on cybersecurity risks. This vital aspect arose from recognition that potential cyber threats and vulnerabilities can compromise systems, products, and services. C-SCRM specifically includes risks from software vulnerabilities, hardware tampering, malicious code insertion, and other cybersecurity threats that can exploit weaknesses in the supply chain. C-SCRM is an integral part of managing the aviation supply chain to minimize exposure to cyber-attacks using supply chain vectors.

“Cybersecurity Supply Chain Risk Management Practices for Systems and Organizations,” NIST SP 800-161r1, (74) is the current authoritative standard for performing C-SCRM. As the definitive yardstick by which C-SCRM efforts are evaluated, it is important for the suppliers of

cyber-enabled parts within the aviation supply chain to embrace and be compliant with this standard.

However, it is also important to prioritize the fundamental principles and best practices of C-SCRM rather than only the specific controls it specifies. Suppliers must be willing to implement new and emerging best practices in this area to maximize protection of the supply chain. It is important to remain adaptable as both the best practices and standards will certainly continue to evolve and change over time. Robust risk management practices that can withstand and adapt to these inevitable updates are essential.

Effective C-SCRM also requires a commitment from all members of the aviation supply chain, from the lowliest individual contributor at the smallest supplier through the most senior officials at the acquisition organizations. Effective execution and management of SCRM requires the collaboration of all stakeholders.

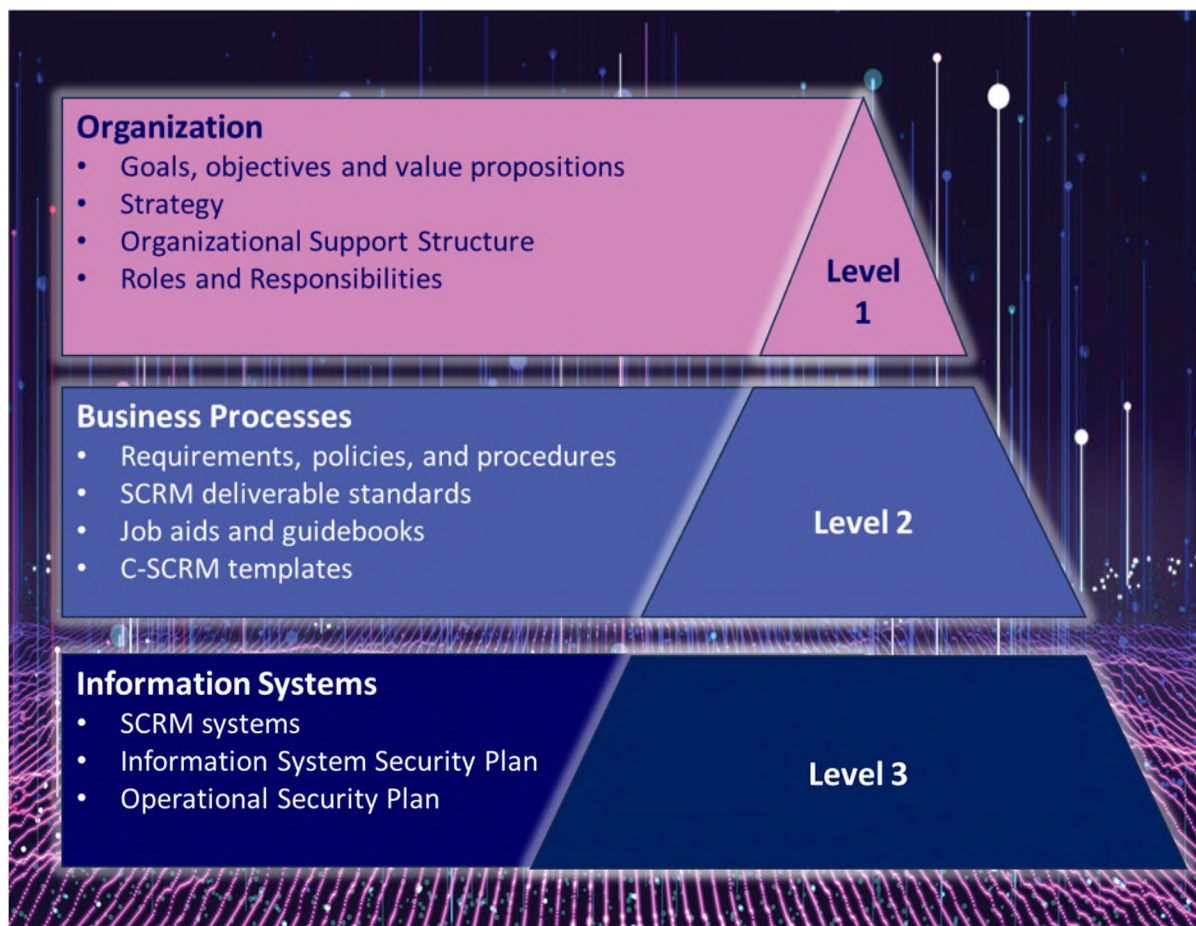


Figure 25. NIST SCRM Framework Model

Figure 25 illustrates a multi-tiered, cross-functional Framework model for SCRM. (90) It highlights that Organization, Business Processes, and Information Systems are all necessary aspects of a holistic SCRM program. (91) The pyramid structure reflects a top-down approach, starting with clearly defined SCRM program goals and objectives, securing the necessary

executive support, and establishing the supporting organizational structure, processes, policies, procedures, SCRM templates, job aids, and system level deliverables.

Figure 26, illustrates the four aspects of a robust SCRM strategy. It is a multifaceted approach that supports **integrity** (accurate and trustworthy data), **resilience** (withstanding disruptions), **quality** (meeting requirements and standards), and **security** (protecting people, products, and information). These aspects overlap, ensuring a comprehensive approach to mitigating risks.



Figure 26. Four Aspects of Supply Chain Risk Management (74)

15.8 Cyber Supply Chain Risk Management (C-SCRM) Key Practices

C-SCRM is built on a foundation of standardized practices and evolving capabilities. To effectively manage C-SCRM, enterprises should focus on achieving a baseline level of maturity that provides a solid foundation for their C-SCRM programs. These key practices should be tailored to the organization's unique context. Essential activities include integrating C-SCRM across the enterprise, establishing a formal program, and closely collaborating with critical suppliers.

Each of the following subsections identifies and enumerates key practice areas to guide organizations in developing a robust C-SCRM process to mitigate their own supply chain risk. Each of these areas were derived from NIST Cybersecurity Supply Chain Risk Management Practices for Systems and Organizations. (74)

15.8.1 Foundational Practices

Foundational practices establish the groundwork for integrating C-SCRM into an organization's operations by developing policies, procedures, and governance structures. This includes defining roles and responsibilities, aligning supply chain risk management with enterprise-wide risk strategies, and promoting a culture of risk awareness across the organization.

To establish an effective C-SCRM program, the following essential elements should be in place:

- **Dedicated Program Management Office:** Establish a core, dedicated, multidisciplinary C-SCRM team with senior leadership support.
- **Risk Management Framework:** Implement a risk management hierarchy and process, including an enterprise-wide risk assessment process.
- **Governance Structure:** Establish an enterprise governance structure that integrates C-SCRM requirements into enterprise policies.
- **Supplier Risk Assessment:** Develop a process for identifying and assessing the criticality of suppliers, products, and services.
- **Awareness and Training:** Raise awareness and foster understanding of C-SCRM, and develop competencies among staff.
- **Acquisition and Procurement:** Integrate C-SCRM into acquisition/procurement policies and procedures.
- **Supplier Management:** Establish a supplier management program, including guidelines for purchasing from qualified original equipment manufacturers.
- **Incident Management:** Implement a robust incident management program to identify, respond to, and mitigate security incidents.
- **Vulnerability Disclosure:** Establish internal processes to validate that suppliers and service providers disclose vulnerabilities in their products.
- **Governance and Monitoring:** Establish a governance capability for managing and monitoring components of embedded software to manage risk across the enterprise.

These essential elements provide a foundation for a robust C-SCRM program, enabling organizations to effectively manage the risks associated with their supply chain.

15.8.2 Sustaining Practices

As aviation platforms rely on complex and interconnected cyber-enabled parts and components, the importance of an effective C-SCRM program is essential. The aviation supply chain presents a unique set of challenges, with its vast and complex web of suppliers, manufacturers, and service providers.

To ensure the safety and security of aircraft platforms, it is essential that aviation organizations implement robust C-SCRM practices that go beyond foundational measures. Sustaining practices, which build upon foundational elements, are critical to advancing C-SCRM capabilities and staying ahead of emerging threats.

To further enhance C-SCRM capabilities, the following sustaining practices should be implemented:

- **Threat-Informed Security Program:** Establish and collaborate with a threat-informed security program to stay ahead of emerging threats.
- **Supplier Security Assessment:** Use confidence-building mechanisms to assess critical supplier security capabilities and practices.
- **Continuous Monitoring:** Establish formal processes for continuous monitoring and reassessment of suppliers and the supply chain.

- **Risk Appetite and Tolerances:** Quantify risk appetite and risk tolerances to empower leaders to make informed C-SCRM decisions.
- **Information Sharing:** Engage with industry groups, government agencies, and other organizations to enhance supply chain cybersecurity threat and risk insights.
- **Training and Awareness:** Embed C-SCRM-specific training into the curriculums of applicable roles across the enterprise.
- **Integration into System and Production Lifecycles:** Integrate C-SCRM considerations into every aspect of the system and production lifecycles.
- **Contractual Language:** Integrate C-SCRM requirements into contractual language with suppliers and service providers.
- **Supplier Engagement:** Include critical suppliers in contingency planning, incident response, and disaster recovery planning and testing.
- **Metrics and Reporting:** Define, collect, and report C-SCRM metrics to ensure risk-aware leadership and drive the efficacy of supply chain management processes and practices.

These sustaining practices provide a foundation for advanced C-SCRM. That enables organizations to further improve their cybersecurity posture and mitigate supply chain risks.

15.8.3 Enhancing Practices

To stay ahead of emerging threats and protect the cybersecurity and resiliency of aircraft platforms, organizations must transition from sustaining to enhancing C-SCRM practices. This next generation of C-SCRM capabilities focuses on adaptive and predictive frameworks and processes that enable organizations to anticipate and mitigate risks before they materialize. By embracing an enhancing culture, aviation organizations can further strengthen their cybersecurity posture, improve their ability to respond to emerging threats, and ensure the continued integrity of the aviation supply chain.

To advance towards adaptive and predictive C-SCRM capabilities, the following essential elements should be in place:

- **Automation and Efficiency:** Automate C-SCRM processes to drive consistency, efficiency, and free up critical resources.
- **Quantitative Risk Analysis:** Adopt probabilistic approaches to reduce uncertainty and optimize resource allocation for risk response and measuring return on investment.
- **Predictive Metrics and Insights:** Apply leading C-SCRM metrics to shift from reactive to predictive strategies, adapting to risk profile changes before they occur.
- **Community of Practice and Continuous Improvement:** Establish or participate in a community of practice to enhance and improve C-SCRM initiatives. Keep up to date with emerging best practices and threats.

These essential elements provide a foundation for optimizing C-SCRM enhancements, enabling organizations to optimize resource allocation, improve predictive capabilities, and stay ahead of emerging threats.

15.9 Integrating C-SCRM with Risk Management Framework (RMF) Security Controls

DoD and commercial aviation platforms are subject to assessment and authorization processes. Military aircraft are subject to comprehensive assessment of cybersecurity and resiliency controls under the Risk Management Framework as described in Chapter 12.

C-SCRM is an integral aspect of the Risk Management Framework (RMF) and a part of a well-structured approach to managing risk in the supply chain. The integration of C-SCRM within the context of RMF security controls is critical for managing overall cyber risk to aircraft platforms.

For the complex and global supply chains in the aviation industry, effective risk management is crucial to ensuring the security and integrity of aircraft systems and components. To mitigate the risks associated with the supply chain, aviation companies must adopt a proactive and collaborative approach to risk management. This involves working closely with customers, suppliers, and partners to identify and prioritize risks. It also necessitates developing a comprehensive risk management strategy that includes compliance, metrics, checklists, and identification of root causes.

Furthermore, aviation companies must invest in education and training programs to promote workforce awareness of their responsibilities in ensuring supply chain security. By taking a collaborative and proactive approach to risk management, aviation organizations can minimize the risks associated with the supply chain, protect against cyber threats, and ensure the safety and security of products and services.

15.10 Summary and Additional Resources

The aviation supply chain is complex and global, with multiple tiers of suppliers, manufacturers, and service providers. That makes C-SCRM a vital aspect for ensuring the security and integrity of aircraft. A proactive and collaborative approach to risk management necessarily involves close working relationships with customers, suppliers, and partners across the supply chain.

Suppliers in the aviation industrial base must account for the critical role of the cyber supply chain on their development lifecycle and production processes. This involves implementing a risk management hierarchy and process, establishing a governance structure, and developing a supplier risk assessment process.

To advance towards adaptive and predictive C-SCRM capabilities, aviation suppliers and manufacturers must implement foundational, sustaining, and enhanced C-SCRM practices. The aviation industrial base must continuously strengthen our cybersecurity posture, improve our ability to respond to emerging threats, and ensure the continued integrity of the aviation supply chain.

The following additional resources are recommended for more in-depth information about the topics and concepts presented in this chapter.

- **Workshop Brief on Cyber Supply Chain Best Practices. (92)**

This workshop provided crucial insights into managing cyber risks within the supply chain, particularly emphasizing the importance of integrating cybersecurity measures throughout all stages of supply chain management. The detailed best practices, risk assessment questions, and examples of successful strategies offers invaluable guidance to enhance the security, reliability, and integrity of supply chains.

- **Baker's Dozen: 13 Elements of an Effective SCRM Program. (93)**

This flyer outlines 13 critical elements for an effective Supply Chain Risk Management (SCRM) program, emphasizing the importance of securing the supply chain at all levels of an organization. It highlights essential practices such as obtaining executive commitment, communicating across organizational stakeholders, prioritizing critical assets and suppliers, and implementing integrated risk reduction measures. This provides a structured approach to mitigating risks that can compromise the integrity and security of aviation components and systems.

- **Deliver Uncompromised: A Strategy for Supply Chain Security and Resilience in Response to the Changing Character of War. (94)**

This MITRE document provides a strategic framework for enhancing supply chain security and resilience, particularly focusing on the Department of Defense (DoD) amid evolving threats. It outlines various courses of action (COAs) aimed at elevating security as a primary metric in acquisition, forming a National Supply Chain Intelligence Center, and implementing continuous monitoring of software integrity. This document is highly relevant to professionals in Cyber Supply Chain Risk Management (C-SCRM) for aviation systems because it addresses the critical need for securing complex supply chains that aviation systems rely on.

- **NIST Cybersecurity Supply Chain Risk Management (C-SCRM) Program Website. (95)**

This NIST C-SCRM webpage includes the NIST Cybersecurity Supply Chain Risk Management Practices for Systems and Organizations (NIST SP 800-161r1), which offers comprehensive guidelines, and the C-SCRM Fact Sheet, providing a concise overview of key concepts. Additionally, the page features important resources and activities essential for implementing C-SCRM and details on risk management processes for identifying, assessing, and mitigating supply chain risks.

- **NIST Case Studies in Cyber Supply Chain Risk Management**

<https://csrc.nist.gov/Projects/cyber-supply-chain-risk-management/key-practices>

The National Institute of Standards and Technology (NIST) has researched Cyber Supply Chain Risk Management (C-SCRM) practices through two rounds of industry engagement. These case studies provide insights into how organizations approach C-SCRM, including tools, techniques, and processes. These resources offer a detailed exploration of C-SCRM practices and their evolution over time.

- **NIST C-SCRM Software and Supply Chain Assurance Forum**

<https://csrc.nist.gov/Projects/cyber-supply-chain-risk-management/SSCA>

The Software and Supply Chain Assurance Forum (SSCA) is a collaborative effort between government, industry, and academia to share knowledge and expertise on software and supply chain risks, effective practices, and mitigation strategies. This forum provides a platform for diverse stakeholders to share information, collaborate, and receive feedback on current and future work. Meetings are held 2-3 times a year and are open to all interested parties, with presentations and discussions encouraged to facilitate open interaction. The forum's scope includes supply chain assurance and offers a valuable opportunity for stakeholders to engage with each other and advance the field of software and supply chain risk management.

- **NIST C-SCRM Federal Cyber Supply Chain Risk Management Forum**

<https://csrc.nist.gov/projects/cyber-supply-chain-risk-management/federal-c-scrm>

The Federal C-SCRM Forum is a collaborative platform for organizations to share information and best practices on cybersecurity supply chain risk management. The forum brings together federal employees and contractors with dedicated C-SCRM responsibilities to discuss issues of interest, share knowledge, and inform policy makers. By facilitating the exchange of information and expertise, this initiative aims to strengthen relationships, promote awareness, and build understanding of C-SCRM while increasing transparency and capabilities across the federal ecosystem.

Part 4 – Program Management and Executive Leadership



Block 30 F-16C from Cannon AFB, New Mexico, in flight.

Chapter 16

The Critical Role of Executive Leaders in Cyber Resilient Aviation

This Chapter was written by Lt Gen (retired) J. Kevin McLaughlin, President, McLaughlin Global Associates LLC. It was lightly edited by Teresa Merklin.

United States national cyber strategy and policy all confirm that cybersecurity threats are among the highest risks that exist for small, medium and large organizations in both the private and public sectors. As a result, organizations that develop and acquire systems critical to national security have a growing need to adopt an advanced and integrated approach to identifying and mitigating cybersecurity risks across their entire enterprise. That includes the industrial base that designs and develops aviation systems. Unfortunately, in many organizations, the approaches taken are flawed in fundamental ways.

Of greatest importance, the senior leadership teams (SLT) in many organizations have been inattentive. Specifically, they have not accepted direct accountability and responsibility for understanding, quantifying, and mitigating the cyber risks most critical to their organizations. In every organization, senior leaders can understand and are accountable for many complex areas of risk. Cyber-related risks should be no exception.

As a first priority, SLTs must assess the quantifiable strategic risks tied to brand, reputation, mission, market standing, public trust, product safety and product effectiveness. Then, they must directly connect the sensitivity of those risks to cybersecurity. This allows them to effectively allocate cybersecurity resources to retire risks commensurate with those investments. In many cases, cybersecurity spending is growing, but senior leaders are unsure as to whether it is sized appropriately or targets the right areas.

Compounding that issue is the failure of many SLTs to recognize that there can be internal conflicts of interest concerning how cybersecurity risks are identified, mitigated and reported inside the organization. Often, the same people who are responsible for executing cybersecurity-related tasks are those who report on the organization's cyber posture and risk stance. Senior leaders must ensure the appropriate degree of independence between those who execute day-to-day cybersecurity tasks and those who assess their effectiveness and risk exposure.

Also of note is the fact that many organizations approach cyber risk as something that exists only in information technology (IT) networks. With that flawed perspective, cyber risks are solely regarded as the responsibility of the Chief Information Officers (CIOs) and Chief Information Security Officers (CISOs). However, the reality is that cyber risks and threats can manifest in any part of an organization's ecosystem including developed and delivered products.

Accordingly, SLTs must ensure that cybersecurity risks are managed appropriately across the scope of the entire organization.

This chapter provides information to help SLTs correct the above deficiencies. Organizations of all sizes can adopt these principles, even if their implementation is managed and scaled to the uniqueness of each organization. This guidebook also has specific applicability to government organizations whose missions are affected by cybersecurity risks.

16.1 Senior Leadership as a Recipe for Success

Almost all cyber enabled systems are interconnected and networked in some way. This trend results in new products and capabilities that transform our lives for the better. However, it also creates an avenue for constant and relentless cyber-attacks that are causing strategic harm to the nation. The magnitude of these threats cause governments and societies to demand steps to be taken to protect the country and its citizens from harm. Government and private sector leaders are increasingly being held accountable for cyber-related incidents. Moreover, future competitiveness is increasingly being tied to an organization's ability to mitigate the cyber risks most directly tied to their mission. This requires new behaviors and approaches for SLTs.

16.1.1 Accepting Responsibility and Accountability

SLTs must first make a conscious decision to take ownership of their organization's cyber related risks, making that decision clear and visible to internal and external stakeholders. This does not mean that the SLTs should become responsible for the day-to-day execution of their organizations' cybersecurity programs at the detailed level. Rather, it means that there are key responsibilities held only by SLTs – ones that should not be delegated.

16.1.2 Tying Strategic Risk Exposure to Targeted Cybersecurity Risk Mitigation

As SLTs take on responsibility and accountability for cyberspace risks, their first step is one not directly related to cyberspace. It entails identifying and quantifying the areas of strategic business risk to the organization. In many cases, the SLT already understands risks that could negatively impact the organization's brand, reputation, mission, market standing, public trust, product safety and/or product effectiveness. As a result, these SLTs can quantify their risks in terms of dollars or the ability to succeed in carrying out an assigned mission.

Only by taking this step, can the SLT effectively conduct a true business case analysis of the investments aimed at retiring any area of risk exposure. Although SLTs know how to do this, and many do it every day, most do not use the same approach for identifying and mitigating the associated cyber risks that also directly impact the business. As a result, as cybersecurity investment grows, SLTs cannot be sure as to whether the investments are sized appropriately or target the right areas. Moreover, they cannot tie cybersecurity investments to retiring the corporate risks they worry about every day.

16.1.2.1 Become Informed About Cybersecurity Threats and Risks

SLTs should educate themselves on the nature of the cybersecurity threats that exist in their domain. There are countless sources of threat data in government reports, reports from think tanks and research centers, and information from sector-specific information sharing and analysis centers. For SLTs with security clearances, there are venues where classified threat data can be shared as well.

Additionally, SLTs should work to understand cybersecurity-related policy stemming from the White House, Department of Defense (DoD), Department of Homeland Security (DHS) and others throughout the U.S. government (USG) and Intelligence Community. By obtaining this guidance and direction, SLTs can appropriately allocate funding in accordance with policy and regulation.

USG budgets will also have an effect. Government cybersecurity operations often impact acquisition policies and authority, and may shift when there are changes to roles, functions and responsibilities for cybersecurity development and acquisition programs. Even companies with no public sector customers would likely benefit from having these governmental organizations' insight.

16.1.2.2 Determining Which Strategic Business Risks Have a Cybersecurity Component

Once an SLT has identified and quantified strategic organizational risks and has become familiar with the threat and overall cybersecurity environment, it must determine whether any of the identified strategic organizational risks have a cybersecurity component. This step requires taking a multidisciplinary approach that spans the entire organization. That includes its supply chain, corporate administration, production systems, and mission environments. It also includes deliverable products such as aviation systems when deployed to the operational environment.

Each SLT must ensure that cybersecurity risks are managed across the entire ecosystem, and that cross-functional processes exist where appropriate. In addition, the SLT should determine whether there is a significant human element to identified cybersecurity risk areas. To the degree possible, mitigations of these human risks should be understood and quantified. These mitigations could include new internal cybersecurity policies, procedures, training, and exercises for the company's entire workforce.

16.2 Conducting Both Internal and Independent Assessments

Once the above steps have been taken, the SLT should conduct internal and external assessments of the cybersecurity elements that directly contribute to strategic corporate risk exposure. Much of this can be done initially by internal personnel, but key independent assessments are also warranted.

As a first step, the SLT must identify and select cybersecurity organizations capable of conducting independent cybersecurity risk assessments. The SLT must also ensure sufficient oversight of what will most likely be a series of discrete evaluations. Although these assessments

should be executed under the overall direction of the SLT, they should be performed in very close coordination with the organization's cybersecurity personnel.

The assessments should include both material and non-material approaches to cybersecurity, such as assessing all organizational written directives and procedures. That applies regardless of whether those procedures are implemented, and whether they comply with any applicable government direction and guidance.

Next, the SLT should also ensure that independent cybersecurity assessments are performed against both material and non-material cybersecurity measures. Each evaluation should also include an examination of the organizational business/administrative and production operations infrastructure, as well as deliverable products and platforms.

To achieve independence, these assessments must include adversarial network-penetration testing, third-party blue team evaluations, and internal threat hunting campaigns to search for evidence of compromise. Whoever conducts the assessments should be capable of using the advanced offensive skills required to penetrate the systems under evaluation. Cybersecurity tabletop exercises and other forms of Cyber Risk Assessment (CRA) should be routinely performed.

Cybersecurity awareness training must include both academics and live exercises such as spear phishing tests. It should be delivered by a qualified counterespionage team and cover all facets of sophisticated threats against all employees and systems.

The SLT should establish clear chains of cybersecurity accountability and responsibility to be observed within all elements in the organization. That includes ensuring that the results of all independent assessments are provided directly to the SLT. The results should identify the levels of risks and vulnerabilities identified with each exercise.

Finally, the SLT should consolidate the results from internal and independent assessments to create a comprehensive and cohesive view of the cyber related risks that impact their organization. In doing so, the team should document the cost and timelines required to mitigate the identified cybersecurity risks and should map each to previously identified and quantified strategic organizational risks.

16.3 Prioritizing the Mitigation and Tracking of Cybersecurity Risks

The steps identified in the previous section provides the SLT with the data needed to determine which cybersecurity mitigations offer the greatest return on investment for reducing corporate risks. This data can be used to build a list of prioritized investments and implementation plans to mitigate the highest priority risks.

Similarly, this data can also be used to create an organization-wide incident response plan to be executed in the event of a significant cyberattack. To prepare for such an event, the plan should be implemented routinely in drills (including no-notice events) using realistic tabletop exercises covering myriad scenarios.

16.4 Placing Special Emphasis on the Microelectronics Supply Chain

Microelectronics supply chain risks pose a threat to all modern integrated platforms and systems. These risks pertain to whether microelectronics:

- Can be trusted.
- Will always be available.
- Can be sustained throughout the life of a product or key system.
- Are built to specification and without any backdoors or built-in deficiencies.
- Meet the minimum cybersecurity standards for prime contractors and subcontractors.

Understanding and managing these risks is increasingly important to the USG. Accordingly, it is reasonable to have a dedicated microelectronics supply chain risk mitigation team that can take ownership of this issue as an enterprise risk area. The SLT should ensure that this team analyzes the most basic elements of the supply chain in order to understand where risks and issues regarding trust, availability and/or sustainability exist.

Issues concerning trust pertain to whether a component can be relied on to work as intended, with no functionality added by an adversary. Availability issues consider whether a component can be procured during an adversarial attack. Sustainability issues focus on whether the component will continue to be manufactured once it becomes commercially obsolete. This is especially important in many continuously operating national security and national critical infrastructure systems that rely on microelectronics that might be many generations behind what is used by new and emerging systems.

The SLT should expect the results of this analysis to yield recommendations, in the context of how each mitigates a strategic organizational risk. Risk mitigation tactics might include finding other sources for parts that have better trust and availability; moving processes from foreign locations to U.S. locations, or from higher-risk areas of the world to lower-risk areas; or stockpiling certain parts to mitigate manufacturing stoppages or other causes for delays in part availability. Also, the SLT should pay very close attention to USG policies, laws and investments in domestic sources of critical microelectronics.

16.5 Conclusion

This chapter focused on principles tied to senior leader accountability and responsibility, as well as to tying cyber risk mitigation to strategic corporate business/mission risks. The principles discussed apply to organizations and SLTs of all sizes.

Finally, it is essential for SLTs to take an active role in empowering their organizations to implement the best practices identified throughout this FSAD Guidebook. This can be achieved through leading by example, providing the necessary resources and training, and fostering a culture of continuous improvement. Senior leaders should also ensure that there is clear communication and alignment around the goals and objectives of the FSAD Guidebook, and that there are mechanisms in place to monitor progress and make adjustments as needed.

It is important to remember that implementing the best practices in the FSAD Guidebook is not a one-time event, but an ongoing process that requires commitment, collaboration, and a willingness to adapt. By taking these steps, SLTs can help their organizations develop and deploy cyber resilient aviation systems.

Part 5 – Appendices



First International F-35 Rolls Out of the Factory.

Appendix A Acronyms and Abbreviations

Acronym	Definition
APT	Advanced Persistent Threats
ATC	Air Traffic Control
BYOD	Bring Your Own Device
CA	Certificate Authority
C-I-A Triad	The Confidentiality-Integrity-Availability Triad
CIO	Chief Information Officer
CISO	Chief Information Security Officer
CM	Configuration Management
CMMC	Cybersecurity Maturity Model Certification
CNSS	Committee on National Security Systems
CNSSI	Committee on National Security Systems Instruction
COTS	Commercial Off the Shelf
CPS	Cyber Physical System
CRA	Cyber Risk Assessment
CRSE	Cyber Resiliency Systems Engineering
C-SCRM	Cyber/Cybersecurity Supply Chain Risk Management
CTI	Cyber Threat Intelligence
CTT	Cyber Table Top
CUI	Controlled Unclassified Information
CVE	Common Vulnerabilities and Exposures
CWE	Common Weakness Enumeration
DFARS	Defense Federal Acquisition Regulation Supplement
DHS	Department of Homeland Security
DIB	Defense Industrial Base
DoD	Department of Defense
EW	Electronic Warfare
FAA	Federal Aviation Administration
FSAD	Fundamentals of Secure Aviation Design
GAO	Government Accountability Office
GPS	Global Positioning System
IC	Integrated Circuit
ICAO	International Civil Aviation Organization
ICS	Industrial Control System
IDS	Intrusion Detection System

IP	Intellectual Property
IPS	Information Protection System
IT	Information Technology
JCIDS	Joint Capabilities Integration and Development System
JROC	Joint Requirements Oversight Council
JROCM	Joint Requirements Oversight Council Memorandum(s)
KEW	Kinetic Energy Weapon
LRU	Line Replaceable Unit
MBSE	Model Based Systems Engineering
MOSA	Modular Open Systems Architecture
NIST	National Institute of Standards and Technology
NOTAM	Notice to Air Missions
NSAS	National Strategy for Aviation Security
OT	Operations Technology
PCB	Printed Circuit Board
PIRA	Public Information Release Authorization
PKI	Public Key Infrastructure
RMF	Risk Management Framework
SAST	Static Application Security Testing
SBOM	Software Bill of Materials
SCADA	Supervisory Control and Data Acquisition
SCRM	Supply Chain Risk Management
SE	Systems Engineering
SLT	Senior Leadership Team
SP	Special Publication
SPA	Software Product Assessment
SSE	Security Systems Engineering
STPA	System Theoretic Process Assessment
TTP	Tactics, Techniques, and Procedures
USCYBERCOM	U.S. Cyber Command
USG	U.S. Government
VPN	Virtual Private Network

Appendix B Bibliography

1. **Ron Ross (NIST), Victoria Pillitteri (NIST), Richard Graubart (MITRE), Deborah Bodeau (MITRE), Rosalie McQuaid (MITRE).** NIST SP 800-160 Vol. 2. *NIST*. [Online] November 2019. Withdrawn on December 09, 2021. Superseded by SP 800-160 Vol. 2 Rev. 1. <https://doi.org/10.6028/nist.sp.800-160v2>.
2. **Mr. John Garstka, SES, Director, Cyber, Office of the Chief Information Security Officer.** Cyber Risk to Mission: NDIA Systems & Mission Engineering Division Planning. [Online] November 9, 2020. https://www.ndia.org/-/media/sites/ndia/divisions/systems-engineering/se-monthly-meetings/se---december-2020-meeting/garstka_presentation.pdf?download=1.
3. **Cybersecurity Infrastructure Security Agency (CISA).** Transportation Systems Sector. [Online] <https://www.cisa.gov/topics/critical-infrastructure-security-and-resilience/critical-infrastructure-sectors/transportation-systems-sector>.
4. **U.S. Government Accountability Office.** Contested Information Environment: Actions Needed to Strengthen Education and Training for DOD Leaders. [Online] January 2023, 2023. <https://www.gao.gov/products/gao-23-105608>. GAO-23-105608.
5. **William D. Bryant, Joint Force Quarterly 88.** Surfing the Chaos: Warfighting in a Contested Cyberspace Environment. *National Defense University Press*. [Online] January 9, 2018. <https://ndupress.ndu.edu/Publications/Article/1411713/surfing-the-chaos-warfighting-in-a-contested-cyberspace-environment/>.
6. **President Donald J. Trump, The White House.** National Strategy for Aviation Security of the United States of America. [Online] 2018. <https://www.hsdl.org/c/view?docid=821736#:~:text=The%20new%20NSAS%20directs%20a,engagement%20with%20government%20and%20private%2D.821736>.
7. **Joint Task Force Transformation Initiative.** Managing Information Security Risk: Organization, Mission, and Information System View. *National Institute of Standards and Technology (NIST)*. [Online] March 2011. <https://csrc.nist.gov/pubs/sp/800/39/final>. NIST SP 800-39.
8. **Joint Task Force.** Security and Privacy Controls for Information Systems and Organizations. *National Institute of Standards and Technology (NIST)*. [Online] September 2020. <https://csrc.nist.gov/pubs/sp/800/53/r5/upd1/final>. NIST SP 800-53 Rev. 5.
9. **Langewiesche, William.** What Really Brought Down the Boeing 737 Max? *The New York Times Magazine*. [Online] September 18, 2019. <https://www.nytimes.com/2019/09/18/magazine/boeing-737-max-crashes.html>.

10. **Joint Task Force Transformation Initiative.** Guide for Conducting Risk Assessments. [Online] September 2012. <https://csrc.nist.gov/pubs/sp/800/30/r1/final>. NIST SP 800-30 Rev. 1.
11. **United States Government Accountability Office.** Weapon Systems Cybersecurity: DoD Just Beginning to Grapple with Scale of Vulnerabilities. [Online] October 9, 2018. <https://www.gao.gov/products/gao-19-128>. GAO-19-128.
12. **Zetter, Kim.** The Untold Story of the Boldest Supply-Chain Hack Ever. [Online] May 2, 2023. <https://www.wired.com/story/the-untold-story-of-solarwinds-the-boldest-supply-chain-hack-ever/>.
13. **Zetter, Kim.** *Countdown to Zero Day: STUXNET and the Launch of the World's First Digital Weapon*. s.l. : Crown, 2023. ISBN 9780770436193.
14. **Greenberg, Andy.** The Colonial Pipeline Hack Is a New Extreme for Ransomware. [Online] May 8, 2021. <https://www.wired.com/story/colonial-pipeline-ransomware-attack/>.
15. **Greenburg, Andy.** Hackers Remotely Kill a Jeep on the Highway - With Me in It. [Online] July 21, 2015. <https://www.wired.com/2015/07/hackers-remotely-kill-jeep-highway/>.
16. **Greenberg, Andy.** *Sandworm: A New Era of Cyberwar and the Hunt for the Kremlin's Most Dangerous Hackers*. s.l. : Anchor Books of Penguin Random House, LLC, 2019.
17. **Perlroth, Nicole.** This Is How They Tell Me the World Ends: The Cyberweapons Arms Race. s.l. : Bloombury Publishing, February 9, 2021. ISBN 1635576059.
18. **The Drucker Institute, Claremont Graduate University.** Did Peter Drucker Say That? [Online] <https://drucker.institute/did-peter-drucker-say-that/>.
19. **International Civil Aviation Organization.** Security Culture. [Online] [Cited: August 1, 2023.] <https://www.icao.int/Security/Security-Culture/Pages/default.aspx>.
20. **Paul Cichonski (NIST), Thomas Millar (DHS), Tim Grance (NIST), Karen Scarfone (Scarfone Cybersecurity).** Computer Security Incident Handling Guide. [Online] August 2012. <https://csrc.nist.gov/pubs/sp/800/61/r2/final>. NIST SP 800-61 Rev. 2.
21. **Reid, Thomas.** Essays on the Intellectual Powers of Man. [Online] 1785. <https://philpapers.org/rec/REIEOT-19>. B1533.I4 2002.
22. **Barnes, David E. Sanger and Julian E.** U.S. Hunts Chinese Malware That Could Disrupt American Military Operations. [Online] July 29, 2023. <https://www.nytimes.com/2023/07/29/us/politics/china-malware-us-military-bases-taiwan.html#:~:text=The%20Biden%20administration%20is%20hunting,military%2C%20intelligence%20and%20national%20security>.
23. **National Institute of Standards and Technology (NIST).** Framework for Improving Critical Infrastructure Cybersecurity. [Online] April 16, 2018. <https://nvlpubs.nist.gov/nistpubs/cswp/nist.cswp.04162018.pdf>.

24. **Keith Stouffer (NIST), Suzanne Lightman (NIST), Victoria Pillitteri (NIST), Marshall Abrams (MITRE), Adam Hahn (WSU).** s.l. : National Institute of Standards and Technology, May 2015. Superseded by SP 800-82 Rev. 3. NIST SP 800-82 Rev. 2.
25. **L. Johnson (NIST), Kelley Dempsey (NIST), Ron Ross (NIST), Sarbari Gupta (Electrosoft Services), Dennis Bailey (Electrosoft Services).** Guide for Security-Focused Configuration Management of Information Systems. [Online] August 2011.
<https://csrc.nist.gov/pubs/sp/800/128/upd1/final>. NIST SP 800-128.
26. **Ron Ross (NIST), Victoria Pillitteri (NIST), Kelley Dempsey (NIST), Mark Riddle (NARA), Gary Guissanie (IDA).** Protecting Controlled Unclassified Information in Nonfederal Systems and Organizations. [Online] February 2020. Superseded by SP 800-171 Rev. 3.
<https://csrc.nist.gov/pubs/sp/800/171/r2/upd1/final>. NIST SP 800-171 Rev. 2.
27. **Chief Information Officer, U.S. Department of Defense.** Cybersecurity Maturity Model Certification. [Online] <https://dodcio.defense.gov/CMMC/>. CMMC 2.0.
28. **Joint Requirements Oversight Council (JROC).** Manual for the Operation of the Joint Capabilities Integration and Development System. [Online] October 30, 2021.
<https://www.dau.edu/sites/default/files/2024-01/Manual%20-%20JCIDS%20Oct%202021.pdf>. JROCM 079-21.
29. **Committee on National Security Systems, U.S. Government.** *National Information Assurance (IA) Glossary*. 2010. CNSS Instruction No. 4009.
30. **Ron Ross, Mark Winstead, Michael McEvilley.** Engineering Trustworthy Secure Systems. [Online] November 2022. <https://csrc.nist.gov/pubs/sp/800/160/v1/r1/final>. NIST SP 800-160 Vol. 1 Rev. 1.
31. **U.S. Federal Aviation Administration.** Aviation Cyber Initiative (ACI) Informational Page. [Online] September 1, 2023. https://www.faa.gov/air_traffic/technology/cas/aci.
32. **U.S. Department of Defense.** Airworthiness Certification Criteria, Department of Defense Handbook. December 12, 2014. MIL-HDBK-516C.
33. **Paul M. Ragard, CYBERSAFE Program Director.** MIL-HDBK-516C Airworthiness Certification Criteria Cybersecurity Supplement. Version 1.0 Patuxent River, MD : Airworthiness and CYBERSAFE Office (ACO), May 24, 2023. MIL-HDBK-516C Cybersecurity Supplement.
34. **National Institute of Standards and Technology (NIST) Joint Task Force.** Control Baselines for Information Systems and Organizations. [Online] October 2020.
<https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-53B.pdf>. NIST SP 800-53B.
35. **(INCOSE), International Council of Systems Engineering.** *INCOSE Systems Engineering Handbook*. 5th Edition. s.l. : International Council of Systems Engineering (INCOSE), 2023. ISBN: 978-1-119-81429-0.

36. **Saydjari, O. Sami.** Engineering Trustworthy Systems: Get Cybersecurity Design Right the First Time. 1st Edition s.l. : McGraw-Hill Education, July 5, 2018. ISBN 1260118177.
37. **United States Government Accountability Office.** Weapon Systems Cybersecurity: Guidance Would Help DOD Programs Better Communicate Requirements to Contractors. [Online] March 4, 2021. <https://www.gao.gov/products/gao-21-179>. GAO-21-179.
38. **Joint Staff J6, Deputy Director for Information Warfare, Requirements Division.** Cyber Survivability Endorsement (CSE) Implementation Guide, Version 3.0. [Online] July 2022. https://events.afcea.org/afceacyber23/Custom/Handout/Speaker0_Session10259_1.pdf.
39. **Graubart, Deb Bodeau and Richard.** Cyber Resiliency and NIST Special Publication 800-53 Rev.4 Controls. [Online] September 2013. <https://www.mitre.org/sites/default/files/publications/13-4047.pdf>. SP 800-53.
40. **Lockheed Martin.** *Mission Impact Assessment Guidebook, Version 4.3.* 2020. PIRA #AER202007005.
41. **INCOSE.** Guide to the Systems Engineering Body of Knowledge (SEBoK). [Online] May 6, 2024. [https://sebokwiki.org/wiki/Guide_to_the_Systems_Engineering_Body_of_Knowledge_\(SEBoK\)](https://sebokwiki.org/wiki/Guide_to_the_Systems_Engineering_Body_of_Knowledge_(SEBoK)).
42. **National Institute of Standards and Technology (NIST).** The NIST Cybersecurity Framework (CSF) Version 2.0. [Online] February 26, 2024. <https://nvlpubs.nist.gov/nistpubs/CSWP/NIST.CSWP.29.pdf>.
43. **Ron Ross (NIST), Victoria Pillitteri (NIST), Richard Graubart (MITRE), Deborah Bodeau (MITRE), Rosalie McQuaid (MITRE).** Developing Cyber-Resilient Systems: A Systems Security Engineering Approach. [Online] National Institute of Standards and Technology (NIST). <https://csrc.nist.gov/pubs/sp/800/160/v2/r1/final>. NIST SP 800-160 Vol. 2 Rev. 1.
44. **Franklin, Benjamin.** On Protection of Towns from Fire. [Online] February 4, 1975. <https://founders.archives.gov/documents/Franklin/01-02-02-0002>.
45. **U.S. Defense Standardization Program.** Modular Open Systems Approach (MOSA). [Online] [Cited: May 13, 2024.] <https://www.dsp.dla.mil/Programs/MOSA/>.
46. **United States Code.** Title 10 U.S.C. 4401(b). [Online] <https://uscode.house.gov/view.xhtml?req=granuleid:USC-prelim-title10-section4401&num=0&edition=prelim>.
47. **Eric M. Hutchins, Michael J. Cloppert, Rohan M. Amin, Ph.D.** Intelligence-Driven Computer Network Defense Informed by Analysis of Adversary Campaigns and Intrusion Kill Chains. s.l. : Lockheed Martin Corporation Whitepaper, 2011.

48. **Scott Rose (NIST), Oliver Borchert (NIST), Stu Mitchell (Stu2Labs), Sean Connelly (DHS).** Zero Trust Architecture. [Online] August 2020.
<https://csrc.nist.gov/pubs/sp/800/207/final>. NIST SP 800-207.
49. **Karen Uttecht et al., MIT Lincoln Laboratory.** Application of Zero Trust (ZT) Concepts to Air Force Weapon Systems (WS). November 17, 2020.
50. **Shostack, Adam.** *Threat Modeling: Designing for Security*. s.l. : Wiley, 2014. ISBN: 978-1-118-81005-7.
51. **M. Muckin and S.C. Fitch.** A Threat-Driven Approach to Cyber Security. s.l. : Lockheed Martin Corporation, 2019.
52. **Kohnfelder, Loren and Garg, Praerit.** *The threats to our products*. 1999.
53. **MITRE.** Common Attack Pattern Enumerations and Classifications (CAPEC). [Online] July 6, 2023. <https://capec.mitre.org/>.
54. —. ATT&CK Framework. [Online] MITRE. <https://attack.mitre.org/>.
55. —. Common Weakness Enumeration. [Online] MITRE. <https://cwe.mitre.org/>.
56. —. *CVE Program Mission*. [Online] MITRE. <https://www.cve.org/>.
57. **OWASP.** OWASP Top Ten. [Online] OWASP. <https://owasp.org/www-project-top-ten/>.
58. **NIST.** National Vulnerability Database (NVD). [Online] NIST. <https://nvd.nist.gov/>.
59. **Deputy Assistant Secretary of Defense for Systems Engineering and Department of Defense Chief Information Officer.** Trusted Systems and Networks (TSN) Analysis. [Online] June 2014. <https://rt.cto.mil/wp-content/uploads/2019/06/Trusted-Systems-and-Networks-TSN-Analysis.pdf>.
60. **NIST.** Guide for Conducting Risk Assessments. s.l. : National Institute of Standards and Technologies (NIST), September 2012. Supersedes: SP 800-30 (07/01/2002). NIST SP 800-30 Rev. 1.
61. **Director of Defense Research and Engineering (Advanced Capabilities), et al.** The Department of Defense Cyber Table Top Guide. [Online] September 16, 2021.
<https://www.cto.mil/wp-content/uploads/2023/06/DoD-Cyber-Table-Top-Guide-v2-2021.pdf>.
62. **United States Department of Defense.** Cybersecurity Test and Evaluation Guidebook. [Online] February 10, 2020.
<https://www.dau.edu/sites/default/files/Migrated/CopDocuments/Cybersecurity-Test-and-Evaluation-Guidebook-Version2-change-1.pdf>. CASE # 20-S-0618.
63. **United States Government Accountability Office.** Weapon Systems Cybersecurity: DOD Just Beginning to Grapple with Scale of Vulnerabilities. [Online] October 9, 2018.
<https://www.gao.gov/products/gao-19-128>. GAO-19-128.

64. **FAS.** Rainbow Series and Related Documents. *Federation of American Scientists Intelligence Resource Program*. [Online] <https://irp.fas.org/nsa/rainbow.htm>.
65. **Valletta, Anthony M.** DoD Information Technology Security Certification and Accreditation Process. [Online] [https://www.fismacenter.com/DITSCAP%20\(2\).pdf](https://www.fismacenter.com/DITSCAP%20(2).pdf).
66. **Department of Defense.** DoD Information Assurance Certification and Accreditation Process (DIACAP). [Online] <https://web.archive.org/web/20090825030249/http://www.dtic.mil/whs/directives/corres/pdf/851001p.pdf>.
67. **NIST Joint Task Force.** Risk Management Framework for Information Systems and Organizations: A System Life Cycle Approach for Security and Privacy. [Online] December 2018. Supersedes: SP 800-37 Rev. 1 (06/05/2014); CSWP 3 (06/03/2014). <https://csrc.nist.gov/pubs/sp/800/37/r2/final>. NIST SP 800-37 Rev. 2.
68. **Murugiah Souppaya (NIST), Karen Scarfone (Scarfone Cybersecurity), Donna Dodson.** Secure Software Development Framework (SSDF) Version 1.1. [Online] February 2022. <https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-218.pdf>. NIST SP 800-218.
69. **Department of Defense.** DoDI 8500.01. *Department of Defense Instruction: Cybersecurity*. [Online] March 14, 2014. https://www.esd.whs.mil/Portals/54/Documents/DD/issuances/dodi/850001_2014.pdf.
70. **Office of the Director of National Intelligence.** Annual Threat Assessment of the U.S. Intelligence Community. [Online] February 6, 2023. <https://www.dni.gov/files/ODNI/documents/assessments/ATA-2023-Unclassified-Report.pdf>.
71. **XJTAG.** Technical Guide to JTAG. [Online] 2024. <https://www.xjtag.com/about-jtag/jtag-a-technical-overview/>.
72. **IEEE.** Recommended Practice for Encryption and Management of Electronic Design Intellectual Property (IP). [Online] November 2023. <https://ieeexplore.ieee.org/document/10328536>. IEEE 1735-2023.
73. **Agent, DoD Anti-Tamper Executive.** What Is Anti-Tamper. *What Is Anti-Tamper*. [Online] Department of Defense. [Cited: 07 21, 2024.] <https://at.dod.mil/What-Is-Anti-Tamper/>.
74. **Jon Boyens (NIST), Angela Smith (NIST), Nadya Bartol (Boston Consulting Group), Kris Winkler (Boston Consulting Group), Alex Holbrook (Boston Consulting Group), Matthew Fallon (Boston Consulting Group).** Cybersecurity Supply Chain Risk Management Practices for Systems and Organizations. [Online] May 2022. Supersedes: SP 800-161 (04/08/2015). <https://csrc.nist.gov/pubs/sp/800/161/r1/final>. NIST SP 800-161 Rev. 1.
75. **Defense Acquisition University.** Critical Function/Component Risk Assessment. [Online] https://content1.dau.edu/DAUMIG_se-brainbook_189/content/Management%20Processes/Critical-Function-Component-Risk-

Assessment.html#:~:text=Criticality%20Analysis%20is%20the%20process,to%20mission%20failure%20or%20degradation..

76. **Celia Paulsen, et al.** Criticality Analysis Process Model. [Online] <https://nvlpubs.nist.gov/nistpubs/ir/2018/NIST.IR.8179.pdf>. NISTIR 8179.

77. **Celia Paulsen (NIST)**. Identifying Critical Assets for Risk Management. [Online] May 16, 2018. <https://csrc.nist.gov/CSRC/media/Presentations/NISTIR-8170-Criticality-Analysis-Process-Model-C/images-media/Criticality%20Analysis%20051618%20-%20Celia%20Paulsen.pdf>.

78. **National Archives**. CUI Category: Controlled Technical Information. [Online] April 10, 2024. <https://www.archives.gov/cui/registry/category-detail/controlled-technical-info.html>.

79. **U.S. House Armed Services Committee**. Report of the Defense Critical Supply Chain Task Force. [Online] July 22, 2021. https://democrats-armedservices.house.gov/_cache/files/e/5/e5b9a98f-9923-47f6-a5b5-ccf77ebbb441/7E26814EA08F7F701B16D4C5FA37F043.defense-critical-supply-chain-task-force-report.pdf.

80. **Joint Task Force (NIST)**. Risk Management Framework for Information Systems and Organizations: A System Life Cycle Approach for Security and Privacy. s.l. : National Institute of Standards and Technology, December 2018. NIST SP 800-37 Rev. 2.

81. **Defense Federal Acquisition Regulation Supplement (DFARS)**. 252.239-7018 Supply Chain Risk. 252.239-7018.

82. **U.S. National Counterintelligence and Security Center**. Supply Chain Risk Management: Reducing Threats to Key U.S. Supply Chains. [Online] September 25, 2020. <https://www.dni.gov/files/NCSC/documents/supplychain/20200925-NCSC-Supply-Chain-Risk-Management-tri-fold.pdf>.

83. **U.S. Department of Defense**. Operation of the Defense Acquisition System. [Online] January 7, 2015. <https://www.acq.osd.mil/fo/docs/500002p.pdf>. DoD Instruction Number 5000.02.

84. **Kristen J. Baldwin, Acting Deputy Assistant Secretary of Defense for Systems Engineering (DASD(SE))**. Cybersecurity in Program Acquisition. [Online] August 18, 2016. <https://www.ndia.org/-/media/sites/ndia/policy/documents/cyber/baldwin-kristen.ashx?la=en>. DOPSR Case # 16-S-1757.

85. **U.S. 115th Congress (2017-2018)**. S.3085 - Federal Acquisition Supply Chain Security Act of 2018. [Online] <https://www.congress.gov/bill/115th-congress/senate-bill/3085/text>. S. 3085.

86. **LM Space to use Cyber SDK to monitor/detect fundamental health attributes and RF characteristics**. Standards for Security Categorization of Federal Information and Information Systems. [Online] February 2004. <https://nvlpubs.nist.gov/nistpubs/FIPS/NIST.FIPS.199.pdf>. FIPS 199.

87. **Jennifer Cawthra, National Cybersecurity Center of Excellence, NIST.** Data Integrity: Detecting and Responding to Ransomware and Other Destructive Events. [Online] December 2020. <https://www.nccoe.nist.gov/projects/building-blocks/data-integrity/detect-respond>. NIST SP 1800-26A.
88. **Kassner, Michael.** Anatomy of the Target data breach: Missed opportunities and lessons learned. [Online] February 2, 2015. <https://www.zdnet.com/article/anatomy-of-the-target-data-breach-missed-opportunities-and-lessons-learned/>.
89. **Global Risk Institute (GRI).** Geopolitical Risks, Responding to global affairs in a state of disruptive transition. [Online] 2024. <https://globalriskinstitute.org/category/geopolitical-risks/>.
90. **National Institute of Standards and Technologies (NIST).** Cybersecurity Supply Chain Risk Management C-SCRM. [Online] May 06, 2024. Created May 24, 2016, Updated May 06, 2024. <https://csrc.nist.gov/Projects/cyber-supply-chain-risk-management#:~:text=The%20NIST%20Cyber%20Supply%20Chain%20Risk%20Management%20%28C-SCRM%29,cyber%20supply%20chain%20compromise%2C%20whether%20intentional%20or%20unintentional>.
91. **Joint Task Force Transformation Initiative.** Guide for Conducting Risk Assessments. s.l. : National Institute of Standards and Technology, September 2012. NIST SP 800-30 Rev. 1.
92. **NIST.** Best Practices in Cyber Supply Chain Risk Management. [Online] <https://csrc.nist.gov/CSRC/media/Projects/Supply-Chain-Risk-Management/documents/briefings/Workshop-Brief-on-Cyber-Supply-Chain-Best-Practices.pdf>.
93. **Director of National Intelligence.** Baker's Dozen: 13 Elements of an Effective SCRM Program. [Online] <https://www.dni.gov/files/NCSC/documents/supplychain/20190326-Baker-Dozen.pdf>.
94. **NIST.** Deliver Uncompromised: A Strategy for Supply Chain Security and Resilience in Response to the Changing Character of War. [Online] <https://www.dni.gov/files/NCSC/documents/supplychain/20190327-Deliver-uncompromised.pdf>.
95. —. Cybersecurity Supply Chain Risk Management Resource Page. [Online] <https://csrc.nist.gov/projects/cyber-supply-chain-risk-management>.
96. **Officer, Office of the DoD Chief Information.** DoD Manual 8140.03. *Cyberspace Workforce Qualification and Management Program*. [Online] 02 15, 2023. <https://dodcio.defense.gov/Portals/0/Documents/Library/DoDM-8140-03.pdf>.
97. **Ponemon Institute.** 2020 Cost of Insider Threats Global Report. s.l. : Observe IT, IBM Security, 2020.
98. **National Archives.** About Controlled Unclassified Information (CUI). [Online] August 1, 2019. <https://www.archives.gov/cui/about>.

99. **U.S. Cybersecurity & Infrastructure Security Agency (CISA).** Defense Industrial Base Sector. [Online] <https://www.cisa.gov/topics/critical-infrastructure-security-and-resilience/critical-infrastructure-sectors/defense-industrial-base-sector>.
100. **The White House.** Executive Order on Improving the Nation's Cybersecurity. [Online] May 12, 2021. <https://www.whitehouse.gov/briefing-room/presidential-actions/2021/05/12/executive-order-on-improving-the-nations-cybersecurity/>.
101. **National Institute of Standards and Technology.** Glossary: Supply Chain Risk Management (SCRM). [Online] https://csrc.nist.gov/glossary/term/supply_chain_risk_management.
102. **Robert D. Buzzell.** Is Vertical Integration Profitable? [Online] January 1983. <https://hbr.org/1983/01/is-vertical-integration-profitable>.
103. **Jon Boyens (NIST), Celia Paulsen (NIST), Nadya Bartol (Boston Consulting Group), Kris Winkler (Boston Consulting Group), James Gimbi (Boston Consulting Group).** Key Practices in Cyber Supply Chain Risk Management: Observations from Industry. [Online] February 2021. <https://csrc.nist.gov/pubs/ir/8276/final>. NIST IR 8276.
104. **Dawn Bailey (NIST).** Lumberjacks and Supply Chain Cybersecurity: Take Time to Prepare. [Online] August 2, 2018. <https://www.nist.gov/blogs/blogrife/lumberjacks-and-supply-chain-cybersecurity-take-time-prepare>.
105. **National Institute of Standards and Technology (NIST).** NIST Shares Key Practices in Cyber Supply Chain Risk Management Based on Observations from Industry. [Online] February 22, 2021. <https://www.nist.gov/news-events/news/2021/02/nist-shares-key-practices-cyber-supply-chain-risk-management-based>.
106. **Carol Woody, Robert J. Ellison.** Supply-Chain Risk Management: Incorporating Security into Software Development. [Online] March 2010. https://insights.sei.cmu.edu/documents/435/2013_019_001_297341.pdf.
107. **Lily Hay Newman, Wired.** The Anatomy of a Cisco Counterfeit Shows Its Dangerous Potential. [Online] July 17, 2020. <https://www.wired.com/story/counterfeit-cisco-switch-teardown/>.
108. **Goodin, Dan.** Microsoft discovers critical SolarWinds zero-day under active attack. [Online] July 12, 2021. <https://arstechnica.com/gadgets/2021/07/microsoft-discovers-critical-solarwinds-zero-day-under-active-attack/>.
109. **Hope, Alicia.** PHP Team Averted a Supply Chain Attack After Hackers Compromised Their Self-Hosted Git Server and Inserted a Backdoor. [Online] April 9, 2021. <https://www.cpomagazine.com/cyber-security/php-team-averted-a-supply-chain-attack-after-hackers-compromised-their-self-hosted-git-server-and-inserted-a-backdoor/>.

110. **Musil, Steven.** Trump signs bill barring US government use of Kaspersky. [Online] December 12, 2017. <https://www.cnet.com/news/privacy/trump-signs-law-barring-us-government-from-using-kaspersky/>.
111. **Popov, Nikita.** Changes to Git commit workflow. [Online] March 28, 2021. <https://news-web.php.net/php.internals/113838>.
112. **Samuel H. Russ, Jacob Gatlin.** Three Ways to Hack a Printed Circuit Board. [Online] August 21, 2020. <https://spectrum.ieee.org/three-ways-to-hack-a-printed-circuit-board>.
113. **Office of the Assistant Secretary of Defense for Sustainment.** Supply Chain Risk Management Framework Project Report - Phase 1. [Online] February 15, 2023. https://www.acq.osd.mil/log/LMR/scrm_report.html/DoD_SCRM_Framework_Report_Phase_I.pdf.