



National Security Agency
Cybersecurity Technical Report

**DoD Microelectronics:
Levels of Assurance Definitions and
Applications**

JULY 2022

U/OO/173659-22

PP-22-1079

Version 1.0



For additional information, guidance or assistance with this document, please contact the Joint Federated Assurance Center (JFAC) at <https://jfac.navy.mil>.



Notices and history

Document change history

Date	Version	Description
JULY 2022	1.0	Initial Publication

Disclaimer of warranties and endorsement

The information and opinions contained in this document are provided "as is" and without any warranties or guarantees. Reference herein to any specific commercial products, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government, and this guidance shall not be used for advertising or product endorsement purposes.

Publication information

Author(s)

National Security Agency
Cybersecurity Directorate
Joint Federated Assurance Center

Contact information

Joint Federated Assurance Center: <https://jfac.navy.mil>

Defense Industrial Base Inquiries and Cybersecurity Services: DIB_Defense@cyber.nsa.gov

Media Inquiries / Press Desk: Media Relations, 443-634-0721, MediaRelations@nsa.gov

Purpose

This document was developed in furtherance of NSA's cybersecurity missions. This includes its responsibilities to identify and disseminate threats to National Security Systems, Department of Defense information systems, and the Defense Industrial Base, and to develop and issue cybersecurity specifications and mitigations. This information may be shared broadly to reach all appropriate stakeholders.



Executive summary

This document describes a consistent and measurable approach to addressing assurance risks in the fabrication of custom microelectronic components (CMC), comprised of Application Specific Integrated Circuits (ASIC), Field Programmable Gate Arrays (FPGA), and other microelectronic devices whose function is custom or configurable. This document defines three levels of hardware assurance and the steps necessary to apply them in the protection of custom microelectronic parts used in DoD systems.

For the purpose of this document the Defense Acquisition University (DAU) definition for hardware assurance (HwA) has been adapted to the following:

“An evidence-supported level of confidence that a CMC device and its configuration do not contain unexpected characteristics or exhibit unintended behaviors due to the influence of an adversary or known vulnerabilities that will enable an adversary to influence the system’s behavior. These characteristics or behaviors could range from degraded reliability to denial of service or to complex functional changes.”

Consistent with this definition, the Joint Federated Assurance Center (JFAC) has identified three levels of HwA to be applied by DoD programs to their top-level system and its critical components.

Level of assurance	Typical criteria
	<p>If the system fails, U.S. Government (USG) capability will be reduced in a meaningful way. If the system is subverted, it can cause harm to U.S. personnel, property, or interests. However:</p> <ul style="list-style-type: none"> • Essential operational capabilities for the DoD will remain available even during a system failure.
	<p>If the system fails, the consequences will be grave. If the system is subverted, it can cause serious harm to U.S. personnel, property, or interests. However:</p> <ul style="list-style-type: none"> • Essential operational capabilities for the DoD may be degraded during a system failure, and • Redundant capabilities can be brought online as part of a continuity of operations plan, and • The failure of the system will not cause cascade effects across many DoD or allied systems.



	<p>If the system fails, the consequences will be extremely grave. If the system is subverted, it can cause exceptionally grave harm to U.S. personnel, property, or interests. A failure or subversion of this system:</p> <ul style="list-style-type: none"> • May represent an existential risk to the USG, and • May cascade across many DoD systems in a way that impacts total operational readiness in an immediate way, and • Will interrupt essential operational capabilities of the DoD.
---	--

Once the system is categorized at the appropriate level of assurance (LoA), the respective CMC is further analyzed to determine potential threats to the manufacturing process. The threats are defined by two characteristics at each level: cost and utility. The following table documents the cost and utility characteristics at each LoA.

Criteria		LoA1	LoA2	LoA3
Attack cost	Access			
	A single available point of access	●	●	●
	A difficult point of access <i>or</i> multiple available points of access		●	●
	Multiple points of difficult access			●
	Technology			
	Existing public technology	●	●	●
	Low implementation risk technology		●	●
	Technologically feasible			●
	Investment			
	Minimal investment of resources	●	●	●
	A large multidisciplinary team		●	●
	A nation scale directed priority			●
Attack utility	Value of effect			
	Disable or subvert a system	●	●	●
	Establish vulnerabilities (for future exploitation)		●	●
	Degrade system performance			●
	Targetability			



Criteria		LoA1	LoA2	LoA3
	Inherently targetable and controllable	●	●	●
	Affects only a subset of systems		●	●
	Blind attacks (<i>Difficult to precisely target or control the outcome</i>) ¹			●

After CMC LoAs have been determined by the program, the program should utilize JFAC best practice guides for the relevant assurance level to identify the threats present at that level and effective techniques for mitigating each.

These mitigations can be incorporated directly into a Program Protection Plan (PPP). In this document, CMC products are defined to include the full range of devices containing reprogrammable digital logic that can implement arbitrary digital functions and fully custom integrated circuits. This includes devices marketed as field programmable gate arrays (FPGAs), such as complex programmable logic devices (CPLD), and system-on-a-chip (SoC) FPGAs.

The following is a list of the available and anticipated best practice guides.

- Level of Assurance 1 FPGA Best Practices
- Level of Assurance 2 FPGA Best Practices
- Level of Assurance 3 FPGA Best Practices
- Level of Assurance 1 CIC Best Practices
- Level of Assurance 2 CIC Best Practices
- Level of Assurance 3 CIC Best Practices
- Level of Assurance 1 COTS Best Practices
- Level of Assurance 2 COTS Best Practices
- Level of Assurance 3 COTS Best Practices

¹ As will be discussed later in this document, LoA3 systems are best approached with a full risk analysis to identify which blind attacks are of concern to a given system. Realistic risks in this space are often idiosyncratic, and the most concerning blind attacks are typically not the most expensive. LoA3 is the least "one size fits all" of all the categories specifically because many such systems are judged to need to concern themselves with such unpredictable effects.



Contents

DoD Microelectronics: Levels of Assurance Definitions and Applications	i
Executive summary	iv
1 Levels of assurance	1
1.1 Identification criteria for components.....	2
1.1.1 Final determination of LoA by DoD.....	3
1.1.2 LoA selection example.....	4
1.2 Component levels of assurance threat characteristics.....	6
1.2.1 Cost and utility of threats	6
1.2.2 Cost characteristics.....	8
1.2.3 Relationship of cost and utility to LoA.....	9
2 Levels of assurance	11
2.1 LoA1 threats	11
2.2 LoA2 threats	13
2.3 LoA3 threats	15
3 Implementation strategy	17
4 Summary	19
Appendix A: JFAC FPGA Documents Overview	20
Appendix B: Standardized Terminology	22

Figures

Figure 1: Example LoA determination process for subcomponents	5
---	---

Tables

Table 1: System and component LoA criteria as determined by national impact.....	3
Table 2: Mapping critical components to levels of assurance.....	5
Table 3: LoA threats defined by cost and utility	7
Table 4: Cost characteristics and descriptions	8
Table 5: Utility characteristics and descriptions	9
Table 6: Threat criteria that an LoA must protect against	10



1 Levels of assurance

For the purpose of this document the Defense Acquisition University (DAU) definition for hardware assurance (HwA) has been adapted to the following:

“An evidence-supported level of confidence that a CMC device and its configuration do not contain unexpected characteristics or exhibit unintended behaviors due to the influence of an adversary or known vulnerabilities that will enable an adversary to influence the system’s behavior. These characteristics or behaviors could range from degraded reliability to denial of service or to complex functional changes.”

Consistent with this definition, the Joint Federated Assurance Center (JFAC) has set forth three levels of hardware assurance (HwA) that can be applied by DoD programs to their top-level system and its critical components. These are known as levels of assurance (LoA).

In these documents, each level of assurance identifies three things:



The seriousness of the **consequence** to national security in the event that the system and device fail or are subverted.



The specific types and characteristics of the **threats** that must be addressed for the given level.



A list of JFAC approved **mitigations** for protecting the respective device from the list of given threats.

Programs should first identify the appropriate LoA for the top-level system. Once the system is categorized at the appropriate LoA, the program should analyze the respective CMC to determine potential threats to the manufacturing process. The threats are defined by two characteristics at each level: cost and utility.

Once the program determines the LoAs, it should use a JFAC best practice guide for the relevant assurance level and type of CMC to identify the threats present at that level



and effective techniques for mitigating each. These mitigations can be incorporated directly into a Program Protection Plan (PPP).

The following are the available and anticipated best practice guides:

- *Level of Assurance 1 FPGA Best Practices*
- *Level of Assurance 2 FPGA Best Practices*
- *Level of Assurance 3 FPGA Best Practices*
- *Level of Assurance 1 CIC Best Practices*
- *Level of Assurance 2 CIC Best Practices*
- *Level of Assurance 3 CIC Best Practices*
- *Level of Assurance 1 COTS Best Practices*
- *Level of Assurance 2 COTS Best Practices*
- *Level of Assurance 3 COTS Best Practices*

In this report, CMC products include the full range of devices containing reprogrammable digital logic that can implement arbitrary digital functions and fully custom integrated circuits. This includes devices marketed as field programmable gate arrays (FPGAs), such as complex programmable logic devices (CPLD), and system-on-a-chip (SoC) FPGAs.

1.1 Identification criteria for components

Within this framework and as part of the development of a PPP, a program must determine which LoA is appropriate for each CMC or for each subsystem containing a CMC. That determination is dependent on two elements:

- National impact caused by the failure or subversion of the top-level system, and
- Criticality of the component to that system.

The larger contributor of the two in this determination is the national impact caused by the failure of the system. That is, systems where:

- Failure can be **mitigated** using alternative capabilities and options in real time belong in LoA1.
- The consequence of failure is **dramatic**, but which do not represent existential threats, belong in LoA2.
- Failure represents an **existential threat** to the United States belong in LoA3.



This guidance is not applicable to systems that do not minimally meet LoA1 criteria, but programs should consider implementing the appropriate LoA1 mitigations depending on the their needs.

The following table summarizes the criteria for a program to determine the appropriate LoA for their top-level system.

Table 1: System and component LoA criteria as determined by national impact

Level of assurance	Typical criteria
	<p>If the system fails, U.S. Government (USG) capability will be reduced in a meaningful way. If the system is subverted, it can cause harm to U.S. personnel, property, or interests. However:</p> <ul style="list-style-type: none"> • Essential operational capabilities for the DoD will remain available even during a system failure.
	<p>If the system fails, the consequences will be grave. If the system is subverted, it can cause serious harm to U.S. personnel, property, or interests. However:</p> <ul style="list-style-type: none"> • Essential operational capabilities for the DoD may be degraded during a system failure, and • Redundant capabilities can be brought online as part of a continuity of operations plan, and • The failure of the system will not cause cascade effects across many DoD or allied systems.
	<p>If the system fails, the consequences will be extremely grave. If the system is subverted, it can cause exceptionally grave harm to U.S. personnel, property, or interests. A failure or subversion of this system:</p> <ul style="list-style-type: none"> • May represent an existential risk to the USG, and • May cascade across many DoD systems in a way that impacts total operational readiness in an immediate way, and • Will interrupt essential operational capabilities of the DoD.

1.1.1 Final determination of LoA by DoD

The DoD program Milestone Decision Authority (MDA) makes the final determination of the appropriate LoA for top-level systems based on projected national impact. Once that determination is made, the program can select the required LoA of each component by its criticality, as identified through the PPP’s Trusted Systems and Networks (TSN) analysis.



The program's TSN analysis will result in the assignment of a level of criticality commensurate with the consequence of the component's failure to the system. The TSN levels of criticality are:

- Level 1 - Total Mission Failure
- Level 2 - Significant/Unacceptable Degradation
- Level 3 - Partial/Acceptable
- Level 4 - Negligible

Some components may have their LoA lowered from the system LoA, because they are insufficiently critical to the system itself. The program reduces the LoAs only after a thorough criticality analysis, which takes into account all dependencies, including from the perspective of all trusted relationships and connected devices². Components within a system are often implicitly trusted and relied upon by other components, albeit in subtle ways that only reveal themselves with thorough analysis.

It is not sufficient to demonstrate that the role of a sub-system is less important. Instead, its compromise must not allow an adversary access or influence over other, more critical functions. For example, a component that manages power distribution through a system may be of low complexity and low sensitivity. However, at the same time, its subversion might be catastrophic. For networked components, no component should be assigned an LoA lower than another component with which it shares a trusted relationship or unfiltered connection.

1.1.2 LoA selection example

The following figure illustrates the process of selecting the system-and-component-level LoAs. In the illustration, the program is creating an airplane and stepping through the assurance LoA decision points.

1. The MDA and the program decide that the top-level system will require LoA3 assurance protections.
2. The program identifies all the system's critical components and initially assigns them the system-level LoA; LoA3.
3. The program performs a TSN analysis on each of the critical components resulting in some being downgraded to lower levels of assurance mitigations.

² Trusted relationships exist between devices whenever one device provides input to another device that is given privileges to cause changes and effects. Two devices that share a connected bus or network are a risk, even when they are not intended to communicate with one another in normal operation.



- Finally, the program applies the mitigations to each component appropriate for its level of assurance.




Figure 1: Example LoA determination process for subcomponents

Once this analysis is completed, the analysts on the program can refer to the following table to determine the appropriate LoA for a given component.

Table 2: Mapping critical components to levels of assurance

System LoA	Criticality of component to the system			
	Negligible	Partial / Acceptable	Significant / Unacceptable	Total Mission Failure
LoA1			LoA1	LoA1
LoA2	LoA1	LoA1	LoA2	LoA2
LoA3	LoA1	LoA2	LoA3	LoA3



1.2 Component levels of assurance threat characteristics

Once a component's LoA has been selected, the program is encouraged to utilize the JFAC-authored best practice guides for threats at the relevant LoA. For instance, a program building an FPGA-based system at LoA1 should consult "*Level of Assurance 1 FPGA Best Practices Overview*." Each best practice guide provides the attack and mitigations descriptions designed to satisfy the TSN attack/mitigation analysis in HwA for CMCs.

The following sections define the criteria for identifying what threats will be mitigated at each LoA. They detail the criteria used to measure the likelihood of each threat and under which LoA it should be mitigated. These definitions inform the evaluation of the effectiveness of relevant mitigation approaches and do not represent the mitigations themselves.

1.2.1 Cost and utility of threats

To define the criteria, a program must understand a threat's cost and utility to an adversary. The *JFAC FPGA Assurance Best Practices* guides include the evaluation of each threat category in these two areas:

- The full cost to the adversary to implement the attack and
- The utility of the attack to the adversary.

In the context of the JFAC Assurance Best Practices, **cost** refers to the entire spectrum of resources required to carry out the respective threat. These costs include investment in research and development, personnel requirements, tooling, processes, and funding. Cost also includes opportunity costs associated with the attack. In particular, the opportunity costs include the accesses the adversary must gain and risk to carry out an attack.

Utility is evaluated on two aspects: the value of the effects for the adversary and the degree to which the adversary is in control of the attack (referred to as "targetability"). These two characteristics help programs prioritize mitigations where an adversary could achieve success with reliable high-value effects.

Considered together, cost and utility measure the "likelihood" value in a TSN analysis. This likelihood refers to the potential of an adversary to carry out a specific attack. When graphed against these two axes, FPGA assurance threats are binned into the



LoA at which they are first relevant. The following table lists and defines the type of threats addressed at each LoA according to these two criteria:

Table 3: LoA threats defined by cost and utility

Level	Definitions
	Threats that have a low cost to the adversary and provide them with high utility. They represent the most likely CMC threats.
	Threats that have high or moderate utility for an adversary. For example, introducing a compromised element that is only effective when combined with other compromised components or security failures. LoA2 may include both low- and moderate-cost threats. These represent threats that are likely: <ul style="list-style-type: none">• High-utility effects with moderate costs,• Moderate-utility effects with low costs, and• Moderate-utility effects with moderate costs. Plus: <ul style="list-style-type: none">• Includes all threats from LoA1
	Threats that have a high cost to implement or have low utility. Specifically, this covers threats that are feasible but may require very high costs. It also covers threats that have marginal or difficult-to-control effects. This LoA represents those threats that are of concern, but which are least likely to be enacted: <ul style="list-style-type: none">• High-utility and moderate-utility threats with high costs,• Low-utility threats with low or moderate costs, and• Low-utility threats with high costs (subject to a system-specific risk evaluation). Plus: <ul style="list-style-type: none">• Includes all threats form LoA1 and LoA2

The three LoAs address the threat categories beginning with the low-cost and high-utility threats and continue across the spectrum.

The mitigation requirements for each level include the requirements of the previous levels and provide cascading protections. That is, a system that requires LoA2 will provide mitigations to all threats that are relevant to LoA2 and LoA1. Accomplishing LoA3 mitigations requires a program to provide mitigations for the threats in LoA2 and LoA1.



The type of threats that would be addressed by each LoA have specific areas of evaluation for the attack cost and attack utility.

1.2.2 Cost characteristics

For HwA, JFAC defines attack cost characteristics in three specific areas as captured in the following table:

Table 4: Cost characteristics and descriptions

Cost characteristic	Description
Access 	The level of access to a vendor, a network, a design center, or a fabrication facility that an adversary requires to conduct an attack. This also includes shipment and storage in the supply chain; during assembly, integration, and testing; and in the field. This includes the number of access points and the amount of time that access is required.
Technology 	The level and complexity of technology required by an adversary to conduct an attack. This includes the types of tools/software, techniques, and level of expertise.
Investment 	The volume of resources necessary to carry out the attack. These resources include funding, labor, and time. All resource costs are specified as relative to a nation state actor.

JFAC defines attack utility characteristics in two specific areas as described in the following table:



Table 5: Utility characteristics and descriptions

Utility characteristic	Description
<p>Value of Effect</p> 	<p>The value of the effect of an attack measures the positive outcome of the attack for an adversary, given success. This measurement is relative to the criticality of the particular system.</p> <p>A high-value effect might disable or subvert a system in arbitrary ways. Moderate-value effects might establish vulnerabilities that can be combined with future operations in useful ways. The lowest value effects degrade performance in a non-specific way, leaving the adversary to hope that the overall degradation establishes an advantage.</p>
<p>Targetability</p> 	<p>The measurement of the attack’s ability to be directed to a specific target for a specific effect at a specific time. This category includes evaluating the reliability and predictability of the outcome. Attacks with high targetability can be controlled by the adversary to occur on demand and to a specific device or system. Attacks with a lower targetability may be unreliable, happen at times that are hard to control, or even happen in cases where an adversary does not want them to occur, leading to possible coincidental detection.</p>

1.2.3 Relationship of cost and utility to LoA

An important factor in categorizing the threat to the appropriate LoA is that the cost-related criteria increases as the cost to the adversary increases. Therefore,

- LoA1 protects against the simplest and least costly threats.
- LoA3 protects against the most costly threats.

Conversely, the criteria related to the utility of a given attack decreases as the LoA increases, indicating that LoA3 must protect against threats that achieve lower utility for the adversary. Less critical USG systems protect against attacks with low cost and high utility to the adversary. The most critical USG systems must concern themselves with all threats including those with high costs and/or attacks that have lower utility effects.

Each LoA must defend against attacks relevant to lower LoAs (i.e., LoA2 must defend against all relevant LoA1 **and** LoA2 threats). This overlapping coverage relationship is captured in the following table:



Table 6: Threat criteria that an LoA must protect against

Criteria		LoA1	LoA2	LoA3	
Attack cost	Access				
	A single available point of access	●	●	●	
	A difficult point of access or multiple available points of access		●	●	
	Multiple points of difficult access			●	
	Technology				
	Existing public technology	●	●	●	
	Low implementation risk technology		●	●	
	Technologically feasible			●	
	Investment				
	Minimal investment of resources	●	●	●	
	A large multidisciplinary team		●	●	
	A nation scale directed priority			●	
	Attack utility	Value of effect			
		Disable or subvert a system	●	●	●
Establish vulnerabilities (for future exploitation)			●	●	
Degrade system performance				●	
Targetability					
Inherently targetable and controllable		●	●	●	
Affects only a subset of systems			●	●	
Blind attacks (<i>Difficult to precisely target or control the outcome</i>) ³				●	

³ As will be discussed later in this document, LoA3 systems are best approached with a full risk analysis to identify which blind attacks are of concern to a given system. Realistic risks in this space are often idiosyncratic, and the most concerning blind attacks are typically not the most expensive. LoA3 is the least "one size fits all" of all the categories specifically because many such systems are judged to need to concern themselves with such unpredictable effects.



2 Levels of assurance

Levels of assurance are achieved when the protections implemented on a system mitigate specific levels of threats. The threats are categorized according to the five criteria outlined in the preceding section:

- Access,
- Technology,
- Investment,
- Value of effect, and
- Targetability.



2.1 LoA1 threats

To meet LoA1, a system must protect against the following threat categories of LoA1 attacks.

LoA1 attacks:

- Exploit a single available point of access,
- Use existing public technology,
- Require minimal investment of resources,
- Disable or subvert system capabilities, and
- Are inherently targetable and controllable.

LoA1 threats present an attack that is relatively inexpensive for a nation-state adversary to carry out and has both high-value and targetable outcomes. Using the criteria, these attacks require only a single available point of access and a minimal investment of resources; are carried out using existing public technology; and are inherently targetable with high utility to an adversary.

Threats that exceed any of these LoA1 criteria must be resolved at a higher LoA.



Access – A single available point of access to some portion of the device supply chain is defined as the following:

- An Internet connected network, regardless of other security measures⁴;

⁴ The final version of NIST 800-172 (currently a draft) may provide a standard by which an Internet connected network could be considered sufficiently secure.



- Any single uncleared U.S. person⁵;
- A group of associated foreign nationals within a U.S. organization, such as a corporate office operated in a foreign country;
- A foreign-owned company servicing part of the supply chain; or
- Any number of foreign nationals from a high-threat country or its allies with access to parts of the CMC supply chain.



Technology – Existing public technology means that an attack can be conducted using tools that are already available in the public or commercial domain or are straightforward advances of public technology. Examples would include:

- Development tools provided by Engineering Design Automation (EDA) vendors,
- Internal debugging features that are capable of changing device configurations,
- Laboratory or fabrication equipment used as intended,
- Publicly available open-source projects,
- Published academic research, and
- Results of USG Research & Development (R&D) investment at the unclassified level, even when protected by International Traffic in Arms Regulations (ITAR).



Investment – Minimal investment of resources means that an attack requires a team with existing knowledge and skills as well as individuals with domain knowledge in the technology area of the system being assured. For the LoA analysis, minimal resources is defined as any effort consisting of approximately the equivalency of six person-years of specific CMC/domain expertise or less, focused solely on attacking the target of interest.



Value of Effect – Attacks that **disable or subvert a system capability** enable an adversary to remove a capability from service or cause it to perform specific deleterious actions. When combined with high targetability, these represent the worst-case scenario for a failure of hardware assurance, enabling an adversary to take over or disable capabilities on command.



Targetability – Inherently targetable and controllable threat operations are executed in a way that provides straightforward means to understand and

⁵ The use of a clearance in this context is complex, as a program may be unclassified. As of now, the USG does not have a specific means of assessing individuals for risk of involvement in sabotage that can cause harm to the security of the USA. Should such a system be developed, it would be a better choice for this role. As it is, redundancy serves as an equally valid method to meet the requirements that can be inferred from this statement.



predict the effect of an attack and provide a mechanism to control or time the attack. For example, an adversary who introduces new code into a system design can implement a broad number of malicious functions.

A denial-of-service attack falls in this category, if and only if, it is possible for the adversary to control when it takes effect after the device is fielded. However, a simple reduction in reliability not tied to any trigger⁶, which therefore cannot be controlled or timed in a planned way, does not fall in this category.

Programs with systems required to achieve LoA1 based on the national impact of their compromise can refer to the corresponding *Assurance Threats Catalog* to determine the categories of threats against which it will protect. Accompanying that document are LoA mitigations packages, which specify mitigations that resist attacks against the trustworthiness of the device. The *Level of Assurance 1 Best Practices Overview* specifies a pre-evaluated list of options that can sufficiently mitigate each threat of interest. JFAC can provide custom guidance if the suggested mitigations prove infeasible or uniquely costly for some applications. While these mitigations are in place, the CMC can be considered to have achieved LoA1.



2.2 LoA2 threats

To meet LoA2, a system must protect against the threat of LoA2 attacks.

LoA2 attacks:

- Exploit a difficult point of access,
- Use technology with low risks of implementation,
- Require a large multidisciplinary team,
- Establish vulnerabilities, and
- Affect only a subset of systems.

Compared to LoA1 threats, LoA2 threats require an increased level of cost, limited utility for the adversary, or both.

⁶ A sufficient trigger can be as simple as a change to an operational environment, when such an environment has unique characteristics.



Using the following criteria, LoA2 threats require either a single difficult point of access or multiple points of simple access in the supply chain to carry out. They may also require known nation-state developed techniques and use a large multidisciplinary team to implement. In addition to the high-utility threats included in LoA1, LoA2 threats include attacks that open the door to future attacks, without being useful in isolation. For example, attacks that pre-position vulnerabilities or access for future attacks are of specific interest. As such, these attacks may not necessarily be targeted or predictable at the time of their execution.



Access – A difficult point of access requires the adversary to compromise a system or an individual to circumvent extensive practices taken to protect that access. Difficult points of access include the following:

- A single air-gapped computer network,
- A single cleared U.S. person,
- A group of uncleared U.S. persons, such as a small corporate office operated in the U.S., and
- Shipping practices that are approved for transport of information or equipment classified Secret or Top Secret.

In addition, these threats may take advantage of multiple available points of access.



Technology – Low implementation risk technology includes any technology which, while not publicly available now, could be implemented with sufficient effort and minimal risk of outright failure. This includes:

- Capabilities that public academic research has identified or internal USG R&D has shown to be practical; and
- Techniques that have, according to substantial amounts of open research, been performed successfully, but for which commercial tools are not available.

In addition, these threats may take advantage of existing public technology.



Investment – A large multidisciplinary team indicates a team conducting the operation that may draw on multiple skills not necessarily associated with either the CMC or the application domain. For the LoA analysis, a large multidisciplinary team is defined as any effort consisting of at most 50 work-years of a wide range of technical expertise, focused solely on attacking the device of interest.



For example, physics and materials science experts may have suitable technical skills. This level also accounts for a more substantial amount of effort, potentially a sizeable organization working on the attack for a year or more.



Value of Effect – Attacks that introduce a vulnerability that is not by itself a complete attack may still **establish vulnerabilities**. These attacks would also require additional access and technical development to develop into a complete attack. In addition, attacks that disable or subvert capabilities are included at LoA2.



Targetability – Attacks that **affect only a subset of systems** where the adversary has no control over that subset. For systems that are affected, the behavior must still be inherently targetable and controllable. In addition, any inherently targetable and controllable attacks from LoA1 are relevant at this level.

Programs with systems required to achieve LoA2 based on the national impact of compromise should refer to the corresponding *Assurance Threats Catalog* to determine the specific attacks against which it will need to protect. Accompanying that document are CMC LoA mitigations packages, which specify mitigations that resist attacks against the trustworthiness of the device.

Level of Assurance 2 Best Practices specifies a pre-evaluated list of options that can sufficiently mitigate each threat of interest. Further, custom guidance from JFAC can also be made available if the suggested mitigations prove infeasible or uniquely costly in some applications. While these mitigations are in place, the CMC can be considered to have achieved LoA2.



2.3 LoA3 threats

Similar to the threats in LoA2, threats at LoA3 have a high cost and/or a low utility to the adversary. In the extreme case of high cost **and** low utility attacks, a system-specific threat analysis is appropriate to eliminate all threats that may be present but are not necessarily relevant to the system under consideration.

LoA3 attacks:

- Exploit multiple points of difficult access,



- Use technology an adversary could feasibly develop with some investment,
- Require the coordination of many resources and prioritized direction from a nation-state,
- Degrade the behavior or performance of a system, and
- Blindly impact parts, devices, or users while offering the adversary limited control.

Using the previously described criteria, these threats may require multiple difficult points of access, are not technologically realized but feasible, and are generally blind attacks. These may be difficult to use for precision targeting or may exist solely for pre-positioning.



Access – Multiple points of difficult access in different areas of the CMC supply chain could include:

- Multiple people working on different elements of the CMC or government design teams. These people can be cleared or uncleared, or
- Multiple people performing different functions in the fabrication process.



Technology – Technologically feasible threats are those where existing research indicates that the technology could be developed with an investment that would be feasible for a known adversary. These threats might not be associated with existing or known tools and might not have associated reporting indicating adversary activity. Moreover, while all of the threats are possible, it may be that there is no known or ongoing adversary investment in the capability.



Investment – A directed nation-scale priority refers to a substantive program conducted by a nation-state that coordinates resources from many specialties and organizations across a wide scope to facilitate an attack.



Value of Effect – Those that **degrade the behavior or performance of a system** without fully disabling any specific feature or a reliable and specific planned effect.⁷ For instance, a communications link might be “degraded” in a way that prevents all meaningful communication. Such an attack would fall under disabling a

⁷ The term *degradation* may be used in some domains in a different way.



capability. In addition, this LoA must include all higher value effects described in LoA1 and LoA2.

Targetability – Blind attacks are attacks that impact large numbers of parts, whole device families, or users without effort to impact a specifically targeted part and in a way that has a significant likelihood of discovery.



Blind attacks are those where it is hard to predict the interaction between what the adversaries do and the consequence of the attack. This could include attacks that are performed against far more targets than expected or, without an outside trigger or without foreknowledge of the attack outcome that would inform the adversary of its execution. These attacks can also include activation times that cannot be controlled once fielded; that is to say, a pre-determined time at which devices will fail. In addition, this LoA must consider all higher targetability attacks described in LoA1 and LoA2.

When a system is required to achieve LoA3, the program management can refer to the corresponding *Assurance Threats Catalog* document to determine the categories of threats against which it will need to provide protections. Typically, LoA3 will require a specific threat analysis to determine which attacks are relevant to that system. Specifically, such an analysis may eliminate some attacks that are blind, of low value, and of extremely high cost to the adversary while preserving attacks that do not share all of those criteria, such as low cost blind attacks, or high utility blind attacks. Such an analysis must be system-specific.

One way to access subject matter expertise and develop such a plan is direct engagement with a JFAC lab. An appropriate JFAC lab can conduct a tailored threat assessment to evaluate critical mitigations on a system-by-system basis. While these mitigations are in place, the CMC can be considered to have achieved LoA3.

3 Implementation strategy

As part of the PPP process, each program must develop a plan for assurance. This report establishes a specific approach designed to meet this requirement for CMC systems by assigning the system or sub-system to one LoA. The mitigation packages available from JFAC represent pre-evaluated sets of mitigations that, if implemented correctly, achieve the desired LoA and satisfy the assurance requirements for the identified threat. In cases where an existing mitigation package is not sufficient,



programs may engage directly with JFAC to develop alternatives. As with other elements of the PPP, the program office provides requirements to performers and auditing to validate that the corresponding requirements were met.

To enable this process, ongoing strategic technology development is required to guarantee that appropriate technology is both available and cost-effective. The paragraphs below describe a way of evaluating when additional R&D in assurance technology is necessary.

LoA1 should be achievable through the use of purely commercial, non-protected technology or other approved standards-based guidance when programs are implemented using available technology that is currently in production. This means programs must stay apprised of technology and standards to protect against the relevant threats. Specialized capabilities may still be required to provide assurance in cases where legacy CMC are in use.

LoA2 and LoA3 may require the use of government IP or specialized commercial IP. It may also require additional design processes at the program level; specialized screening of devices by commercial anti-counterfeiting labs; and evaluation of platforms and practices by JFAC labs or others on a per CMC platform basis.

At LoA3, programs will likely be required to work directly with the JFAC labs or other domain experts to ensure their procurement processes and test methods are sufficient to validate their use in national security systems.

Programs should document their LoA analysis, determinations, and mitigations in the PPP. Programs without experience selecting and/or implementing technical mitigations for CMC HwA are strongly encouraged to engage their appropriate JFAC representatives as early in the program lifecycle as possible. The JFAC representative will guide the program through the process to ensure efficient and effective mitigation strategies. Furthermore, it is highly recommended that programs implementing LoA2 and LoA3 protections consult JFAC prior to entering each MDA milestone and before each major design review.

In preparation for the preliminary design review (PDR), JFAC can be consulted on the appropriate LoA for any component when the program is uncertain of the correct designation. In preparation for the critical design review (CDR), JFAC can be consulted



on mitigation plans. At later stages in the design and manufacturing, JFAC can be consulted for HwA recommendations or investigations. At a minimum, all programs should engage with JFAC for a review to prepare for the CDR to obtain a JFAC evaluation of the assurance risk analysis and risk plan so the appropriate information is incorporated into the PPP.

4 Summary

In practice, each program must identify and implement the LoA required to resolve specific risks to that program. Achieving these LoAs will depend on both the CMC platform's assurance (both the hardware and software tools) and the actions of the application developers using the platform. This includes all steps of implementation of the system using the CMC.

Conducted properly, this LoA framework can minimize the impact of hardware assurance CMC threats. Additionally, it equips the program with the understanding necessary to implement a response in the case of a compromise. JFAC is available to guide programs through this process. Additional information for JFAC may be found at <https://jfac.navy.mil>.



Appendix A: JFAC FPGA Documents Overview

This appendix describes the JFAC FPGA set of documents as well as a brief content description.

1. **Levels of Assurance Definitions and Applications** – This document describes the three levels of assurance. It also provides instructions on how to select the LoA for a given system or mission.
2. **Threat Catalog** – This document details and defines the hardware assurance threats within the FPGA space. It identifies which threats are of concern at a given LoA. It provides context to understand the reason for the proposed mitigations in the subsequent documents. Reviewing this document is not required to implement the JFAC FPGA hardware assurance best practices, but it may provide context that explains the rationale behind certain decisions.
3. **Level of Assurance 1 Best Practices** – This series of documents provides a set of mitigations that must be applied to reach LoA1 in the JFAC FPGA assurance flow. For each threat of interest at LoA1, it provides one or more proposed mitigations. Some mitigations are sufficiently straightforward and stable to be described within the document itself. Other mitigations are more detailed and are described in additional documents, as listed below:
 - Standards of Counterfeit Screening for LoA1
 - Design Flow Assurance for LoA1
 - Platform Design Review for LoA1
 - Cryptographic IDs for Counterfeit Discovery for LoA1
 - Third Party IP Review for LoA1
 - Built-In Configuration Authentication for LoA1
4. **Level of Assurance 2 Best Practices** – This series of documents provides a set of mitigations that must be applied to reach LoA2 in the JFAC FPGA assurance flow. The series includes the following set of documents:
 - Standards of Counterfeit Screening for LoA2
 - Approved Second Order Effects Screening Methods for LoA2
 - Post Assembly Analysis for LoA2



- Third Party IP Review for LoA2
 - Platform Design Review for LoA2
5. **Level of Assurance 3 Best Practices** – This series of documents provides a set of mitigations that must be applied to reach LoA3 in the JFAC FPGA assurance flow.



Appendix B: Standardized Terminology

The following terms are used in the Joint Federated Assurance Center Field Programmable Gate Array Best Practices documents. These terms are modified from Defense Acquisition University definitions to support common understanding.

Application design – The collection of schematics, constraints, hardware description language (HDL), and other implementation files developed to generate an FPGA configuration file for use on one or many FPGA platforms.

Application domain – This is the area of technology of the system itself, or a directly associated area of technology. For instance, the system technology domain of a radar system implemented using FPGAs would be "radar" or "electronic warfare."

Configuration file – The set of all data produced by the application design team and loaded into an FPGA to personalize it. Referred to by some designers as a "bitstream", the configuration file includes that information, as well as additional configuration settings and firmware, which some designers may not consider part of their "bitstream."

Controllable effect – Program-specific, triggerable function allowing the adversary to attack a specific target.

Device/FPGA device – A specific physical instantiation of an FPGA.

External facility – An unclassified facility that is out of the control of the program or contractor.

Field programmable gate array (FPGA) – In this context FPGA includes the full range of devices containing substantial reprogrammable digital logic. This includes devices marketed as FPGAs, complex programmable logic devices (CPLD), system-on-a-chip (SoC) FPGAs, as well as devices marketed as SoCs and containing reprogrammable digital logic capable of representing arbitrary functions. In addition, some FPGAs incorporate analog/mixed signal elements alongside substantial amounts of reprogrammable logic.

FPGA platform – An FPGA platform refers to a specific device type or family of devices from a vendor.



Hard IP – Hard IP is a hardware design captured as a physical layout, intended to be integrated into a hardware design in the layout process. Hard IP is most typically distributed as Graphic Design System II (GDSII). In some cases, Hard IP is provided by a fabrication company and the user of the IP does not have access to the full layout, but simply a size and the information needed to connect to it. Hard IP may be distributed with simulation hardware description language (HDL) and other soft components, but is defined by the fact that the portion that ends up in the final hardware was defined by a physical layout by the IP vendor.

Level of assurance (LoA) – A Level of Assurance is an established guideline that details the appropriate mitigations necessary for the implementation given the impact to national security associated with subversion of a specific system, without the need for system-by-system custom evaluation.

Physical unclonable function (PUF) – This function provides a random string of bits of a predetermined length. In the context of FPGAs, the randomness of the bitstring is based upon variations in the silicon of the device due to manufacturing. These bitstrings can be used for device IDs or keys.

Platform design – The platform design is the set of design information that specifies the FPGA platform, including physical layouts, code, etc.

Soft IP – Soft IP is a hardware design captured in hardware description language (HDL), intended to be integrated into a complete hardware design through a synthesis process. Soft IP can be distributed in a number of ways, as functional HDL or a netlist specified in HDL, encrypted or unencrypted.

System – An aggregation of system elements and enabling system elements to achieve a given purpose or provide a needed capability.

System design – System design is the set of information that defines the manufacturing, behavior, and programming of a system. It may include board designs, firmware, software, FPGA configuration files, etc.

Target – A target refers to a specific deployed instance of a given system, or a specific set of systems with a common design and function.



Targetability – The degree to which an attack may have an effect that only shows up in circumstances the adversary chooses. An attack that is poorly targetable would be more likely to be discovered accidentally, have unintended consequences, or be found in standard testing.

Third-party intellectual property (3PIP) – 3PIP is a functional unit designed by a different organization than the principal design team. Typically, 3PIP is a pre-existing design, offered for sale by a commercial organization.

Threat category – A threat category refers to a part of the supply chain with a specific attack surface and set of common vulnerabilities against which many specific attacks may be possible.

Utility – The utility of an attack is the degree to which an effect has value to an adversarial operation. Higher utility effects may subvert a system or provide major denial of service effects. Lower utility attacks might degrade a capability to a limited extent.

Vulnerability – A flaw in a software, firmware, hardware, or service component resulting from a weakness that can be exploited, causing a negative impact to the confidentiality, integrity, or availability of an impacted component or components. (MITRE, <https://cve.mitre.org/about/terminology.html>)