ATARC

Implementation Guide: Transitioning from NIST SP 800-63-3 to SP 800-63-4

ATARC Identity Management Working Group

July 2026

Table of Contents

Acknowledgments

Adam McBride, U.S. Department of Health and Human Services, ATARC Identity Management Working Group Government Co-Chair

Kelvin Brewer, Ping Identity, ATARC Identity Management Working Group Industry Chair

David Treece, Yubico, ATARC Identity Management Working Group Industry Co-Vice Chair

Lisa Palma, LC&J Security Solutions LLC, ATARC Identity Management Working Group Industry Co-Vice Chair

Christopher Grant, U.S. Army Enterprise ICAM

Chris Oakleaf, U.S. Department of Veterans Affairs

Don Woodard, Ping Identity

Alyssa Feola, Seventeen Sierra, LLC

Implementation of NIST 800-63-4

1.0 A Strategic Evolution: From Compliance Checklists to Holistic Risk Management

The transition from NIST Special Publication (SP) 800-63-3 to SP 800-63-4 represents a fundamental strategic shift in digital identity management. This is not a simple version update; it is a doctrinal shift from a static control-selection exercise to a continuous, risk-adapting lifecycle. This new standard moves beyond a compliance-oriented model toward a dynamic, lifecycle-based approach called the Digital Identity Risk Management (DIRM) process. This updated framework is designed to be more adaptive to evolving threats while formally incorporating customer experience and privacy considerations into the core risk management process. Instead of treating identity assurance as a one-time selection of levels, the DIRM process establishes a continuous cycle of definition, assessment, tailoring, and evaluation to ensure identity solutions remain effective, usable, and secure over time.

The new DIRM framework restructures the risk assessment process to be more comprehensive and iterative. The following table contrasts the key characteristics of the previous approach with the new framework.

Feature	SP 800-63-3 Approach	SP 800-63-4 DIRM Framework
Core Philosophy	A risk-based selection process focused on mitigating identity, authentication, and federation errors as distinct events.	A holistic, five-step lifecycle process that integrates security, privacy, and customer experience into a continuous improvement cycle.
Key Stages	A linear process of assessing potential impact across six harm categories and selecting the appropriate assurance levels (xALs).	A five-step iterative process: Define, Assess, Select, Tailor, and Continuously Evaluate.
Risk Factors	Focused primarily on the potential negative impacts of an authentication or proofing failure on the organization and the user.	Expands risk considerations to include impacted entities beyond the user and formally mandates assessments for privacy, customer experience, and threat resistance.
Primary Output	A Digital Identity Acceptance Statement documenting the selected assurance levels and rationale.	A comprehensive Digital Identity Acceptance Statement (DIAS) that serves as the capstone output of the entire DIRM process, formally documenting all tailoring decisions, compensating controls, and supplemental controls.

The new DIRM framework is structured around five core steps, creating a clear roadmap for managing digital identity risk throughout the lifecycle of an online service.

1. **Define the Online Service** This step establishes the functional scope, user groups, and impacted entities for an online service to provide context for all subsequent risk decisions.
2. **Conduct Initial Impact Assessment** This step assesses the potential harms of a compromise for each user group across a standard set of impact categories.
3. **Select Initial Assurance Levels and Baseline Controls** This step maps the assessed impact level to an initial set of assurance levels and their corresponding baseline controls.
4. **Tailor and Document Assurance Levels** This step adjusts the baseline controls based on specific assessments of privacy, customer experience, and threat resistance, documenting all decisions.
5. **Continuously Evaluate and Improve** This step establishes a program to monitor performance metrics and user feedback to ensure the identity solution remains effective and to drive ongoing improvements.

The following sections provide a practical, step-by-step breakdown of how to implement this new DIRM framework within your organization.

2.0 Executing the Digital Identity Risk Management (DIRM) Process

This section serves as the core of the implementation guide, providing actionable instructions for IT and security managers to execute each of the five steps of the new Digital Identity Risk Management (DIRM) process. Following these steps will ensure a thorough, documented, and defensible approach to managing digital identity risk in alignment with NIST SP 800-63-4.

2.1 Step 1: Define the Online Service

This initial step is a new formal requirement in SP 800-63-4 and is critical for establishing the context for all subsequent risk decisions. It requires organizations to create a comprehensive description of the online service, its purpose, and its users. This documentation ensures that risk assessments are grounded in the specific operational realities of the service.

At a minimum, the service definition must include the following components:

- **Organizational Mission and Objectives:** A clear statement of the business or mission objectives the online service supports.

- **Mission and Business Partner Dependencies:** A notation of any reliance on external organizations or dependencies by mission partners on the online service.
- **User Groups:** A breakdown of the distinct groups of users who will access the service, including the specific online transactions and access privileges available to each group.
- **Impacted Entities:** A list of all entities that could be negatively affected by a failure of the digital identity system, including users, the organization, mission partners, and the populations served.
- **Data Processed:** A description of the functionality of the online service and the data it is expected to process through its interfaces.
- **Legal and Regulatory Requirements:** An inventory of any legal, regulatory, or contractual requirements, including privacy obligations, that apply to the service.
- **Preexisting Assessments and Technologies:** Documentation of the results of any prior DIRM assessments and the current state of identity technologies in use.
- **Availability of Identity Evidence:** An estimation of whether the target user groups can provide the types of evidence required for identity proofing.

2.2 Step 2: Conduct the Initial Impact Assessment

SP 800-63-4 provides more detailed guidance for assessing the potential harms that could result from a digital identity failure. The framework moves from the six categories in the previous standard to five mandatory impact categories, with a clearer focus on harms to a wider range of impacted entities beyond just the user and the organization. For each user group defined in Step 1, organizations must assess the potential adverse impacts of an identity system compromise.

The table below summarizes the five mandatory impact categories and provides examples of potential harms to different entities.

Impact Category	Examples of Potential Harms
Degradation of Mission Delivery	**To individuals:** Inability to access eligible government services or benefits. **To the organization:** Inability to perform primary mission/business functions effectively.
Damage to Trust, Standing, or Reputation	**To individuals:** Damage to image or reputation as a result of impersonation. **To the organization:** Deterioration of public trust in the organization.

Impact Category	Examples of Potential Harms
Unauthorized Access to Information	**To individuals:** Breach of personal information, leading to identity theft or other secondary harms. **To the organization:** Unauthorized disclosure of organizational information assets or intellectual property.
Financial Loss or Liability	**To individuals:** Debts incurred or assets lost as a result of fraud. **To the organization:** Costs related to fraudulent activity or loss of business.
Loss of Life or Danger to Human Safety, Human Health, or Environmental Health	**To individuals:** Death or physical injury. **To the organization:** Damage to the workforce or the surrounding environment, potentially making it uninhabitable.

2.3 Step 3: Select Initial Assurance Levels and Baseline Controls

This step involves a mapping process where you take the "Combined Impact Level" (Low, Moderate, or High) derived from the impact assessment in Step 2 and map it to an initial assurance level for identity (IAL), authentication (AAL), and, if applicable, federation (FAL). For instance, a Moderate combined impact level would initially map to IAL2, AAL2, and FAL2. While this mapping process is conceptually similar to SP 800-63-3, it is critical to understand that the definitions and requirements for the assurance levels themselves have been significantly updated to address evolving threats and technologies.

Key Changes in Identity Assurance Level (IAL) Definitions

The most critical change is at IAL1, which no longer permits self-asserted identities. The new IAL1 raises the baseline for all identity proofing by requiring evidence that proves the real-world existence of the claimed identity.

Assurance Level	SP 800-63-3 Definition	SP 800-63-4 Definition
IAL1	Self-asserted. No requirement to link the applicant to a specific real-life identity.	Requires validation against sources to prove the **real-world existence** of the claimed identity and provides some assurance the applicant is associated with it.
IAL2	Evidence supports the real-world existence of the identity and verifies the applicant is associated with it. Allows remote or in-person proofing.	Requires collecting **additional evidence** and a more rigorous process for validating the evidence and verifying the identity.

Assurance Level	SP 800-63-3 Definition	SP 800-63-4 Definition
IAL3	Requires in-person identity proofing with a trained representative.	Adds the requirement for a trained representative to interact directly with the applicant and collect **at least one biometric.**

Key Changes in Authenticator Assurance Level (AAL) Definitions

The updates at the AAL level reflect the increased prevalence of sophisticated phishing attacks, mandating stronger, phishing-resistant authenticators at the highest level.

Assurance Level	SP 800-63-3 Definition	SP 800-63-4 Definition
AAL1	Single-factor or multi-factor authentication using a wide range of technologies.	Requires either single-factor or multi-factor authentication. Verifiers are expected to make MFA options available and encourage their use.
AAL2	Requires multi-factor authentication with two distinct factors and approved cryptographic techniques.	Requires multi-factor authentication with two distinct factors. Approved cryptographic techniques are required.
AAL3	Requires a hardware-based authenticator that provides verifier impersonation resistance.	Requires the use of a **public-key cryptographic authenticator with a non-exportable private key that provides phishing resistance.**

2.4 Step 4: Tailor and Document Assurance Levels

Tailoring is the most significant procedural addition in SP 800-63-4. It formally allows organizations to move beyond a one-size-fits-all approach and adjust the baseline controls selected in Step 3 based on a detailed analysis of specific operational risks. This change empowers organizations to stop over-provisioning or under-provisioning security, allowing for a more cost-effective and user-friendly approach that is still defensible. This step ensures the final identity solution is optimized for its unique environment.

The tailoring process is informed by three mandatory assessments:

- **Privacy Assessment:** Identifies potential unintended consequences to the privacy of individuals resulting from the selected controls.
- **Customer Experience Assessment:** Identifies potential barriers, undue burdens, or friction for individuals seeking to access the service.

- **Threat Resistance Assessment:** Determines if the baseline controls are sufficient to address specific threats, threat actors, and known tactics in the organization's operational environment.

Based on these assessments, organizations can introduce compensating or supplemental controls.

Control Type	Purpose	Example
Compensating Control	A control used **in lieu of** a required baseline control when the baseline is not feasible or another control sufficiently mitigates the risk.	Using a federal background investigation to compensate for some identity evidence validation requirements.
Supplemental Control	A control added in **addition to** baseline controls to address a specific threat not covered by the baseline.	Requiring all users of an AAL2 system to use only phishing-resistant authenticators due to a high risk of targeted attacks.

These controls are not loopholes. Compensating controls require rigorous justification of their equivalence, and supplemental controls must be driven by specific, documented threat intelligence. Both must be formally captured in the DIAS to withstand an audit.

The final output of the DIRM process is the **Digital Identity Acceptance Statement (DIAS)**. This formal document records all decisions made during the risk management process and must include, at a minimum:

- Initial impact assessment results.
- Initially assessed assurance levels (xALs).
- The final, tailored xALs with rationale for any changes.
- All compensating controls, including a justification of their comparability.
- All supplemental controls implemented.

2.5 Step 5: Continuously Evaluate and Improve

Under SP 800-63-4, you are now required to establish an ongoing evaluation program to ensure that digital identity solutions remain effective against evolving threats and continue to meet user needs. This requires organizations to collect and analyze performance metrics to drive continuous improvement.

Key performance metrics to monitor include:

Proofing Success & Failure

- Pass Rate (Overall and by proofing type)
- Fail Rate (Overall and by proofing type)
- Abandonment Rate (Overall and by proofing type)

Authentication & Account Management

- Authentication Failures
- Account Recovery Attempts
- Authenticator Type Usage

Fraud & Customer Support

- Confirmed Unauthorized Access or Fraud
- Reported Unauthorized Access or Fraud
- Help Desk Calls (by type)
- Redress Requests

Beyond the core DIRM process, SP 800-63-4 introduces several other new operational requirements that managers must address to ensure a compliant and trustworthy program.

2.6 DIRM, DIRA, and DIAS form a clear hierarchy in NIST SP 800-63-4

DIRM (Digital Identity Risk Management) is the **overarching framework/process**.

DIRA (Digital Identity Risk Assessment) is the **practical method** (playbook) agencies use to execute DIRM.

DIAS (Digital Identity Acceptance Statement) is the **formal documented output** of the process.

2.6.1 How They Work Together (Step-by-Step Flow)

- **DIRM (The Framework)** This is the normative requirement in NIST SP 800-63-4 (Section 3). It defines a structured, risk-based process to:
    - Identify risks to the online service (mission, users, organization).
    - Assess potential harms (e.g., privacy loss, financial loss, reputation damage).
    - Determine initial **assurance levels** (IAL, AAL, FAL) for different user groups/transactions.

    - Tailor those levels with compensating or supplemental controls if needed.
    - Continuously monitor and improve.

DIRM is mandatory for federal agencies and recommended more broadly. It applies to services you manage **and** external services you use (e.g., SaaS tools).

- **DIRA (The Practical Implementation)** The **GSA Digital Identity Risk Assessment Playbook** provides the step-by-step, agency-friendly way to perform DIRM. It includes worksheets, data collection guidance, and integration with existing processes like RMF/FISMA. Key activities in a DIRA:
    - Define the service and user groups.
    - Collect application/system data.
    - Analyze impacts and determine initial assurance levels.
    - Evaluate tailoring needs.
    - Document everything.
- **DIAS (The Deliverable)** This is the **formal record** produced at the end of the DIRA/DIRM process. Organizations **SHALL** create a DIAS for every online service. Relying Parties must review CSP/IdP DIAS documents and incorporate them into their own.

Minimum contents of a DIAS:

- Initial impact assessment results.
- Initially assessed IAL/AAL/FAL.
- Any tailored assurance levels + rationale.
- All compensating controls and residual risks.
- All supplemental controls.

Visual Relationship

```
DIRM (NIST Framework)
   ↓ (implemented via)
DIRA (GSA Playbook - practical assessment)
   ↓ (produces)
DIAS (Formal Acceptance Statement - documented output)
```

- **DIRM** tells you *what* must be done (risk management requirements).
- **DIRA** tells you *how* to do it operationally.
- **DIAS** proves *you did it* and records the decisions for compliance, audits, and sharing with partners.

This flow ensures risk-based, tailored, and documented digital identity decisions instead of one-size-fits-all levels. The process is iterative — especially the continuous evaluation step in DIRM.

Resources:

- NIST SP 800-63-4 DIRM section: https://pages.nist.gov/800-63-4/sp800-63/dirm/
- GSA DIRA Playbook (includes DIAS template): https://www.idmanagement.gov/playbooks/dira/

3.0 Key Operational Requirements: Redress and AI/ML Governance

In addition to the DIRM process, SP 800-63-4 introduces several new operational requirements that strengthen the overall integrity, fairness, and trustworthiness of a digital identity program. This section highlights the most critical new mandates that organizations must implement to be compliant with the new standard.

3.1 Establishing a Formal Redress Process

Organizations must now provide a formal, transparent, and accessible process for individuals who have been negatively impacted by the identity system. This ensures that users who are unable to be proofed, cannot authenticate, or have other grievances have a clear path to seek remedy.

The redress process must, at a minimum:

- Enable individuals to convey grievances through a process that is documented, trackable, and easy to find on a public-facing website.
- Be implemented as a “dedicated function” for impartially reviewing evidence and resolving issues.
- Make human support personnel available to intervene and override outputs generated by automated or algorithmic mechanisms.
- Educate support personnel on issue handling procedures and available alternatives for users.
- Incorporate findings from the redress process into continuous evaluation and improvement activities.

3.2 Assessing Artificial Intelligence (AI) and Machine Learning (ML) Systems

With the increasing use of AI and ML in identity solutions—from biometric matching to fraud detection—SP 800-63-4 introduces specific requirements to manage their associated risks.

Organizations using AI/ML in their identity systems must:

- Document all uses of AI/ML and disclose this information to relying parties that use the system.
- Perform risk assessments for AI/ML systems, referencing frameworks such as the NIST AI Risk Management Framework (AI RMF).
- Perform and document privacy risk assessments for personal information and data processed by such systems.

This action plan provides the final roadmap for putting these new requirements into practice.

4.0 Action Plan: Your Transition Checklist

This final section provides a high-level checklist to guide managers through the transition from SP 800-63-3 to the new framework in SP 800-63-4. Use this as a practical tool for planning and executing the necessary changes to your digital identity program.

Assemble a Cross-Functional Team: Include representatives from IT, security, privacy, legal, and business units to manage the DIRM process.

Inventory and Define All Online Services: Formally document each service's scope, user groups, partner dependencies, and impacted entities as required by Step 1 of the DIRM process.

Re-assess All Services with the New 5-Category Impact Model: Transition from the six categories in SP 800-63-3 to the five in SP 800-63-4, focusing on the expanded view of harms to all impacted entities, not just the user and the agency.

Map Impacts to New Assurance Level Definitions: Determine initial IAL, AAL, and FAL for each service based on the updated definitions, paying close attention to the new IAL1 requirement for "real-world existence" and AAL3 for phishing resistance.

Perform Tailoring Assessments: Conduct and document the required privacy, customer experience, and threat resistance assessments for each service to refine baseline controls.

Develop and Finalize the Digital Identity Acceptance Statement (DIAS): Create the formal DIAS document for each online service, capturing the full DIRM lifecycle including all tailoring decisions, compensating controls, and final xALs.

Establish a Continuous Monitoring Program: Define and begin collecting the performance metrics (e.g., pass/fail rates, help desk calls, fraud reports) required for the new continuous evaluation process.

Implement a Formal Redress Mechanism: Document and operationalize a user redress process that meets the requirements of SP 800-63-4, including dedicated functions and trained support personnel.

Review, Document, and Assess AI/ML Usage: Identify any use of AI/ML in your identity systems, perform the required NIST AI RMF and privacy risk assessments, and complete all necessary documentation.

Identity Assurance Level (IAL) Implementation Guide

Executive Summary

The transition from NIST SP 800-63-3 to NIST SP 800-63-4 for Identity Assurance Levels (IALs) preserves the foundational three-level structure (IAL1–IAL3) but fundamentally redefines the rigor required to establish and maintain confidence in a claimed identity. Revision 4 completes the move from a static, decision-tree selection process to a **Digital Identity Risk Management (DIRM)** model, where identity proofing is treated as a continuous lifecycle rather than a one-time enrollment event.

From an IAL perspective, three critical shifts define the new baseline:

- **The End of Self-Assertion at IAL1:** Revision 4 eliminates the "self-asserted" identity. Even at the lowest assurance level (IAL1), agencies are now required to validate core attributes against authoritative or credible sources. This change is specifically designed to mitigate the proliferation of synthetic identities and automated "bot" enrollments.
- **Mandatory Human Oversight and Redress:** IAL is no longer an algorithmic-only process. Agencies must provide a documented, trackable redress path for users who fail automated proofing. This includes the requirement for human support personnel with the authority to override automated adjudication outputs when supplemental evidence is provided.
- **Rigorous Biometric Binding at IAL3:** For the highest assurance level, the standard now mandates an on-site, attended session with a proofing agent and the collection of at least one biometric characteristic. This ensures a high-confidence, non-repudiable link between the validated identity and the issued credential.

At a policy level, agencies should treat **IAL2 as the baseline for any service providing access to non-public PII**. Under SP 800-63-4, this requires a shift toward "inclusive security"—leveraging the DIRM process to tailor proofing workflows so they remain accessible to diverse populations without compromising the integrity of the validation.

Introduction and Background

This document provides practical guidance to U.S. federal agencies on transitioning Identity Assurance Level (IAL) implementations from NIST SP 800-63-3 to NIST SP 800-63-4. It assumes that agencies already operate SP 800-63-3–aligned proofing infrastructures, including remote document verification, in-person enrollment centers, and existing integrations with authoritative data sources (e.g., DMV, SSA). The goal is to provide a roadmap for re-tuning these existing investments to meet the heightened validation, privacy, and equity requirements of Revision 4.

Scope

The scope of this whitepaper is focused on IAL-related changes (governed by SP 800-63A-4) and their impact on federal environments already broadly compliant with SP 800-63-3. Specifically, it addresses:

- **Risk-based tailoring of IALs** using the DIRM process, aligning proofing strength with the potential harms of an identity failure.
- **Validation requirements for IAL1**, moving beyond self-assertion to source-based confirmation.
- **Human-in-the-loop redress and AI/ML governance**, ensuring that automated identity evidence validation (including biometric matching) is transparent, auditable, and subject to human override.
- **The evidence tiering system**, clarifying the use of FAIR, STRONG, and SUPERIOR evidence in the proofing lifecycle.

Intended Audience

- **Chief Information Officers (CIOs):** Accountable for identity risk strategy and the resource allocation required for modernizing proofing workflows.
- **IAM/ICAM Teams:** Responsible for the technical implementation of evidence collection, validation APIs, and biometric matching systems.
- **Privacy and Civil Liberties Officers:** Charged with ensuring that the increased data collection required by Revision 4 is balanced with PII minimization and equitable access.
- **Customer Experience (CX) Leads:** Focused on minimizing friction in the proofing journey and managing the “warm hand-off” to Trusted Referees or redress agents.

Assumptions

- The agency aligns with SP 800-63-3 for its current enrollment and proofing workflows.
- The agency has existing contracts or integrations with Credential Service Providers (CSPs) or authoritative data sources for attribute validation.
- Existing help desk or support infrastructures can be adapted to support the mandatory human-driven redress requirements.

Key Changes: IAL Side-by-Side Comparison

Feature	NIST SP 800-63-3 (Old)	NIST SP 800-63-4 (New)
IAL1 Foundation	Attributes are **self-asserted**; there is no requirement to link the applicant to a specific real-life identity.	Must support the **real-world existence** of the identity; core attributes must be validated against authoritative/credible sources and linked to the person.
IAL2 Process	Evidence supports real-world existence; requires remote or physically-present identity proofing	Requires a **rigorous process** for collecting and validating evidence to verify the identity.
IAL3 Process	**Physical presence** is mandatory; attributes must be verified by an authorized representative examining physical documents.	Requires an **on-site attended** session with a proofing agent, physical inspection of evidence, and the collection of at least one **biometric**.
AI/ML Governance	Not explicitly addressed.	**Mandatory disclosure** and privacy risk assessments for AI identity functions.
Redress	Recommended general support.	**Mandatory, trackable process** with human override authority.

IAL1 Changes

The Raised Bar for IAL1

In the previous version (SP 800-63-3), IAL1 was effectively the default level for systems that did not need to verify a user's identity. At that time, IAL1 allowed for "self-asserted" identities, meaning there was no requirement to actually link an applicant to a specific, real-life identity or validate the information they provided.

However, SP 800-63-4 fundamentally redefines this baseline by eliminating the "self-asserted" identity at IAL1. Under the new guidelines, even at the lowest assurance level (IAL1), organizations must now validate a user's core attributes against authoritative or credible sources. This change was specifically introduced to raise the bar against highly scalable automated "bot" enrollments and synthetic identity fraud.

The Introduction of "No IAL"

Because IAL1 now requires actual evidence validation, SP 800-63-4 explicitly introduces the option to bypass the Identity Assurance Level framework altogether so organizations aren't forced to validate data unnecessarily.

Before selecting an initial assurance level, organizations must first determine if identity proofing is even needed. According to the updated guidelines, identity proofing (and thus selecting an IAL) is not required if:

- The online service does not need any personal information to execute digital transactions.
- The system needs personal information, but self-asserted attributes are acceptable because the potential harms from accepting unverified information are insignificant.

If either of these conditions are met, the organization's impact assessment may determine that no proofing is required, resulting in "no IAL" being selected. In these scenarios, the identity proofing processes described in the guidelines are simply deemed not applicable to the system.

Why This Matters: The DIRM Shift

Transitioning to IAL under Revision 4 is less about changing the "lock" and more about establishing a "Security Operations Center" for identity. Agencies must move from a static assessment of what a user has (documents) to a continuous evaluation of the validity of the identity claim.

IAL Selection and Tailoring Guidance

Agencies **SHALL** follow the 5-Step DIRM process for IAL selection:

- **Step 1: Define the Online Service.** Document functional scope, user groups, and all entities (including communities and mission partners) that may be impacted by service failures.
- **Step 2: Conduct Initial Impact Assessment.** Evaluate risks for each user group across five new impact categories: Mission Delivery; Trust/Reputation; Information Access; Financial Loss; and Human/Environmental Safety.
- **Step 3: Select Initial xALs.** Map the combined impact level (Low, Moderate, High) to the initial IAL (1, 2, or 3).
- **Step 4: Tailor Controls.** This is a new mandatory phase where agencies assess the initial IAL against **Privacy, Customer Experience (CX), and Threat Resistance**. If IAL2 remote proofing creates high abandonment for a specific group, document the rationale for **Compensating Controls** (e.g., using a Trusted Referee).
- **Step 5: Continuous Evaluation.** Establish a program to gather performance metrics and user feedback to adjust controls over time. Capture all decisions and tailoring rationales in the **Digital Identity Acceptance Statement (DIAS)**.

Upgrade IAL1 Technical Controls

If an agency currently uses IAL1 based on self-assertion, it **SHALL** update its technical processes to include **validation**. This includes confirming core attributes against authoritative or credible sources and implementing "linkage" steps to ensure the attributes belong to the applicant.

Agencies must now prepare an expanded **DIAS** for each online service. This document **SHALL** include:

- Initial impact assessment results and initially assessed xALs.
- A **rationale** for any tailoring that changed the IAL from the initial assessment.
- A detailed list of all **compensating controls** (used in lieu of normative requirements) and **supplemental controls** (added for specific threats).

Implement Redress and Issue Handling

Agencies **SHALL** provide a documented, accessible, and trackable **redress process** for users who experience proofing failures. This includes:

- Providing human support personnel who have the authority to **override** algorithmic or automated adjudication outputs.
- Educating personnel on alternative avenues for users to gain access to services.

Federal Specifics for IAL

- **Minimum IAL2 for PII:** Any service releasing non-public PII should be proofed at IAL2 or higher.
- **Human-in-the-loop Override:** For high-value transactions, an automated "fail" result from a biometric or document scan **SHALL NOT** be the final adjudication without a path to human review.
- **Evidence Validation:** Agencies should prioritize digital evidence (e.g., mDLs) that supports cryptographic validation over simple visual scans of physical documents.

Govern AI and Machine Learning (AI/ML)

If the agency utilizes AI/ML for IAL functions—such as automated identity evidence validation or biometric matching—it **SHALL**:

- Document and disclose the use of AI/ML to any relying parties.
- Provide information on **training data sets**, model update frequencies, and testing results.
- Perform and document specific **privacy risk assessments** for data processed by these AI/ML systems.

Establish Performance Metrics

Agencies must begin tracking specific **identity proofing metrics**, including overall pass/fail rates, abandonment rates per proofing type (remote vs. in-person), and help desk resolution times related to identity issues.

Transition Timeline

1. **Months 0–6:** Reassess all services via the DIRM process; update the DIAS for all PII-releasing systems.
2. **Months 6–18:** Update technical integrations to include IAL1 source validation and establish the formal redress/override workflows.

Analogy: Upgrading from SP 800-63-3 to SP 800-63-4 is like moving from a **static security manual** to a **live security operations centre**. While the old standard provided a set of rules for choosing a lock, the new standard requires you to constantly monitor who is at the door, how they feel about the entry process, and whether the lock needs to be changed because a new threat has been discovered.

Authenticator Assurance Levels (AAL) Implementation Guide

Summary

The transition from **NIST SP 800-63-3** to **NIST SP 800-63-4** for **Authentication Assurance Levels (AALs)** preserves the familiar three-level structure (AAL1–AAL3) but materially changes how agencies select, implement, and monitor authenticators. Revision 4 moves from a checklist orientation to a **Digital Identity Risk Management (DIRM)** model: agencies are expected to determine AALs based on mission impact, threat environment, population characteristics, and equity and privacy considerations, and then show that those decisions are continuously validated over time.

From an AAL perspective, three themes define the new baseline:

- **Phishing-resistant MFA becomes central, not optional.** AAL2 implementations must offer at least one phishing-resistant option, and AAL3 requires phishing-resistant, hardware- or platform-bound authenticators with strong verifier impersonation resistance (for example, FIDO2 passkeys, PIV/CAC, or equivalent cryptographic authenticators). Weak factors such as SMS and email OTP are no longer acceptable for higher-assurance use cases.
- **Risk-based tailoring is explicit.** Agencies must use the DIRM process in the base volume to justify AAL selections per digital service, rather than defaulting to uniform levels across systems. This includes mapping threats like phishing, credential stuffing, and session hijacking to concrete authenticator requirements and compensating controls.
- **Continuous performance monitoring is required.** AAL is no longer a "set once and forget" decision. Agencies are expected to track authenticator usage, failures, fraud indicators, recovery events, and help desk trends, and adjust AAL decisions, methods, or policies as those metrics change over time.

At a policy level, agencies should treat **AAL2 as the practical floor for any system that releases or manipulates personally identifiable information (PII)**, consistent with the intent of **Executive Order 13681**, which mandates multi-factor authentication for such services. In practice under SP 800-63-4, that means pairing MFA with phishing-resistant options for high-value and high-risk transactions, and reserving AAL1 for truly low-risk, low-impact use cases.

To navigate this transition, agencies should:

- **Reassess AAL selections via DIRM in the first 3–6 months**, focusing on high-value services, PII exposure, and mission-critical functions. This should produce a defensible mapping from systems and transactions to required AALs, plus a prioritized list of gaps (for example, reliance on SMS OTP where phishing-resistant options are needed).

- **Plan and execute a phased migration over 12–18 months**, emphasizing:
    - Replacement or containment of non–phishing-resistant authenticators;
    - Introduction of phishing-resistant MFA options for all AAL2 services and mandatory use at AAL3;
    - Instrumentation and reporting needed to demonstrate continuous monitoring of authenticator performance and user impact.

Framed this way, the shift from SP 800-63-3 to SP 800-63-4 is less about renaming assurance levels and more about **operationalizing AALs as a living, risk-based control**, anchored in phishing-resistant authentication and measurable outcomes.

Introduction and Background

This document provides practical guidance to U.S. federal agencies on transitioning **Authenticator Assurance Level (AAL)** implementations from **NIST SP 800-63-3** to **NIST SP 800-63-4**. It assumes that agencies already operate **SP 800-63-3–aligned infrastructures**, including **PIV-based authentication**, multi-factor authentication (MFA) for privileged and remote access, and established identity governance processes. The goal is not to replace existing investments, but to **re-evaluate and tune them** against the revised requirements in SP 800-63-4.

Scope

The scope of this whitepaper is intentionally narrow: it focuses on **AAL-related changes** in NIST SP 800-63-4 and what they mean for federal environments that are already broadly compliant with SP 800-63-3. Specifically, it addresses:

- **Risk-based tailoring of AALs** using the **Digital Identity Risk Management (DIRM)** process introduced in the base volume of SP 800-63-4, including how agencies should align authenticator strength with mission impact, data sensitivity, and threat environment rather than applying a single, uniform level across all systems.
- **Continuous improvement and performance metrics**, including expectations to measure authenticator usage, failure and recovery rates, fraud indicators, and help desk trends over time, and to use those metrics to refine AAL selections and authenticator portfolios on an ongoing basis rather than on a one-time “go-live” basis.
- The **heightened emphasis on usability, accessibility, and equity** within authentication, including requirements that MFA options be usable by people with disabilities, that biometric systems perform equitably across demographic groups, and that authenticator choices not create disproportionate barriers for specific communities or user populations.

Topics such as **Identity Assurance Levels (IAL), Federation Assurance Levels (FAL)**, and the detailed mechanics of identity proofing, federation protocols, or wallet-based credentials are referenced only where they directly influence AAL selection or authenticator design. The primary focus remains on **how agencies should think about, select, and operate authenticators** under the new revision.

Intended Audience

This document is written for a cross-functional set of federal stakeholders:

- **Chief Information Officers (CIOs)** who are accountable for overall digital identity strategy, portfolio-level risk decisions, and major authentication modernization investments.
- **Identity and Access Management (IAM) teams** responsible for designing, implementing, and operating authenticators, MFA policies, and access management platforms across hybrid and cloud environments.
- **Information security officers** and broader cybersecurity staff who must ensure that authentication controls keep pace with evolving threats such as phishing, credential stuffing, and session hijacking, and who map those controls into broader zero trust strategies.
- **Privacy officers and civil rights/civil liberties stakeholders** who are charged with ensuring that authentication approaches respect privacy, support equity, and avoid disparate impacts across different user populations.

While technical in nature, this paper is structured to support **joint decision-making** across these roles, with the intent that risk, usability, privacy, and operations are considered together when selecting and migrating to AAL-compliant authenticators.

Assumptions

The guidance and examples in this document are based on several key assumptions about the reader's current environment:

- **The agency already aligns with SP 800-63-3** for major external-facing and high-value internal systems, including:
    - **Use of PIV/CAC** or equivalent hardware-backed authenticators for high-risk or high-value use cases.
    - **Deployment of MFA** for remote access, administrative access, and systems that process or expose personally identifiable information (PII), consistent with prior federal mandates.
- **Existing IAM platforms support standards-based authentication** (e.g., SAML, OAuth 2.0, OIDC) and can integrate additional authenticators (for example, FIDO2/WebAuthn, one-time passwords, or push-based methods).

- **Agencies have at least basic monitoring and reporting** for authentication events (for example, logs of success/failure, device or network attributes), even if those data are not yet systematically used for DIRM or continuous improvement.

Given these starting conditions, the remainder of the whitepaper concentrates on **what must change**—in authenticator choices, policy, monitoring, and governance—to satisfy the **AAL expectations of NIST SP 800-63-4** without discarding the foundational work done under SP 800-63-3.

Key Changes

The table below summarizes the four most significant shifts in authentication requirements from SP 800-63-3 to SP 800-63B-4. These updates focus on modern threats, usability, and flexibility while maintaining or improving assurance levels.

Change Area	SP 800-63-3 (Previous)	SP 800-63B-4 (Current)	Why It Matters (Impact)
Phishing Resistance	Required only at AAL3; recommended at AAL2	Must be offered/ available to all users at AAL2, Required at AAL3	Major Security upgrade. Directly counters the most common attack vector. Strong alignment with EO 13681 for PII systems.
AAL3 Authenticator Requirements	Strictly hardware based cryptographic authenticators	Non-exportable cryptographic key in secure element/TPM/ TEE allowed FIPS 140 Level 1 sufficient if key are non-exportable	Makes AAL3 more practical on modern commercial devices (smartphones, laptops) all while preserving strong security.
Session Management & Usability	Relatively strict reauthentication timeouts	Relaxed timeouts (SHOULD: 12 hours overall / 15 - 30 minutes inactivity at AAL2) Device-Bound Session Credentials (DBSC) recommended	Significant usability improvements, reduces user friction, and supports higher adoption of stronger authenticators.
Biometrics & Presentation Attack Detection (PAD)	Biometrics allowed in MFA; PAD recommended but not strictly required	PAD SHALL for facial recognition; strongly recommended for other modalities, Biometrics must always be paired with another factor	Stronger protection against deepfakes and presentation attacks. Improves security and equity (requires non-biometric alternatives).
Redress	Recommended general support.	Mandatory, trackable process with human override authority.	

For a comprehensive row-by-row comparison, refer to the **Detailed Side-by-Side Mapping Table** in the Appendix. Agencies should pay special attention to phishing resistance and PAD requirements, as these represent the largest shifts in security posture.

Authenticator Assurance Level	SP 800-63-3 Implementation (Old)	SP 800-63-4 Implementation (New)
AAL1	Provides some assurance of control; requires single-factor or multi-factor authentication (MFA).	Provides basic confidence; verifiers are now **expected to make MFA options available** and encourage their use.
AAL2	High confidence; requires two distinct authentication factors and **approved cryptographic techniques**.	High confidence; requires two factors and approved cryptography. Verifiers are now **expected to offer phishing-resistant options** at this level.
AAL3	Very high confidence; requires a hardware-based authenticator and **verifier impersonation resistance**.	Very high confidence; requires a **public-key cryptographic authenticator with a non-exportable private key** that provides phishing resistance.

Another way to look at the major changes between the old and new standards and providing a brief explanation.

Change Area	Brief Explanation (63-4 vs 63-3)	Impact/Why it Matters
Authenticator types permitted	Modern cryptographic authenticators (e.g., passkeys) are favored; some weak factors (SMS/Email OTP) are disallowed for AAL2/AAL3.	Shifts the ecosystem toward more secure, modern, and usable authentication methods.
Phishing resistance - new definition	Required at AAL3; must be offered at AAL2. Specifically defined via verifier impersonation resistance.	Directly addresses the top threat vector in modern attacks.
FIPS 140 validation relaxed	FIPS 140 Level 1 is sufficient for non-exportable keys in secure elements, simplifying AAL3 on commercial devices.	Makes AAL3 more achievable on common platforms like smartphones and laptops.
Session management and reauthentication intent redefined and relaxed	Inactivity timeouts can be extended (30-60 min) and maximum session time is up to 12 hours (SHOULD), with recommendations for Device-Bound Session Credentials (DBSC).	Major usability improvement by reducing reauthentication friction, while DBSC maintains security.

Change Area	Brief Explanation (63-4 vs 63-3)	Impact/Why it Matters
Replay resistance, authentication intent, non-exportable keys	Strengthened requirements, especially for AAL3, mandating non-exportable cryptographic keys.	Ensures high assurance authenticators are resilient against key export and device compromise.
Password requirements (renamed from "memorized secrets")	Simplified, focusing only on minimum length (8 characters); complexity rules discouraged when MFA is used.	Improves usability and memorability, recognizing that MFA provides the primary security barrier.
Account recovery substantially revised	Multiple paths aligned to AAL; high-assurance recovery processes are now clearly defined.	Provides a robust, risk-based approach to account recovery, a common point of failure.

Highlight of Major Changes: AAL3 no longer strictly requires “hardware-based” but mandates non-exportable keys in secure elements/TPMs. Biometrics require Presentation Attack Detection (PAD) and must always be paired with another factor.

Why this matters: The changes operationalize AALs as a dynamic, risk-based control, shifting focus from static compliance checklists to continuous security and usability.

AAL Selection and Tailoring Guidance

Quick Overview of DIRM Process (Base Volume Sec. 3)

The Digital Identity Risk Management (DIRM) process in the base volume of SP 800-63-4 is mandatory for selecting and justifying AALs. It moves away from a one-size-fits-all model by requiring agencies to tailor AAL based on context.

Possible Decision Flowchart:

1. **Start with Impact Assessment:** Determine the impact of a compromised account/ transaction (e.g., loss of PII, financial loss, mission disruption).
2. **Determine Baseline AAL:** Map impact to a baseline AAL (e.g., High Impact $\to$ AAL3, Moderate Impact with PII $\to$ AAL2, Low Impact $\to$ AAL1).
3. **Tailor Based on Context:** Adjust the AAL based on four key factors:
    - **Threats:** Increase AAL if high-risk threats (e.g., targeted phishing) are present.
    - **Privacy/Equity:** Ensure the selected AAL methods are equitable and meet privacy requirements (SP 800-53, SP 800-52).

- **Usability:** Choose methods that minimize user friction where assurance permits.
- **Population Characteristics:** Consider the user base (e.g., public users vs. federal employees).

4. **Document in Digital Identity Acceptance Statement (DIAS):** Justify the final AAL choice, any compensating controls, and continuous monitoring plans.

Federal Specifics

- **Minimum AAL2 for PII:** Consistent with EO 13681 and modern risk, AAL2 should be the practical minimum for any system that processes or releases personally identifiable information (PII).
- **Phishing Resistant Mandatory for Staff/Contractors:** For federal employees and contractors accessing federal systems, agencies should prioritize phishing-resistant options (e.g., PIV/CAC, FIDO2 passkeys) to meet the AAL2 requirement and align with Zero Trust mandates.

Current State Assessment and Gap Analysis

Agencies must first understand their current environment relative to the new requirements.

Inventory Authenticators

The following table provides a template for inventorying systems and their current authentication posture.

System / Application Name	HR Portal	Public Benefits Portal	Financial System
FISMA ID / System Owner	XYZ-123 / HR Director	ABC-456 / CIO	DEF-789 / CFO
Current AAL (1/2/3)	AAL2	AAL1	AAL3
Primary Authenticators in Use (e.g., PIV, password + TOTP, passkey)	Password + SMS OTP; some PIV cards	Password only	PIV/CAC only
Session Management (timeout, binding method, re-auth frequency)	30-min inactivity; no device binding	60-min timeout cookie-based	15-min inactivity; session tokens tied to device IP

Account Recovery Process (e.g., email OTP, knowledge based, admin reset)	Email magic link + security questions	Self Service email reset	Admin-initiated, in-person verification
Number of Users / User Groups	5,000 federal employees	50,000 public users	500 internal staff
Notes / Special Considerations (e.g., mobile only, legacy system)	Legacy app; no phishing resistant option yet	High usability priority	High-value transactions; critical system

Exhibit - Agency Inventory

Gap Analysis Checklist per AAL

The gap analysis identifies specific policy and implementation deviations from SP 800-63-4.

No.	1	2	3	4	5
Requirement (SP 800.63B-4 Reference)	Phishing resistant authenticator MUST be offered to all subscribers (Sec. 3.2.2)	At AAL3: Authenticator uses non-exportable cryptographic key in secure element/TPM (Sec. 3.3.3)	Memorized secrets (passwords): minimum 8 characters; no composition rules required when paired with another factor (Sec. 3.11.2)	Session binding and re-authentication: relaxed inactivity timeout allowed; device-bound session credentials (DBSC) recommended (Sec. 5.2)	Presentation Attack Detection (PAD) SHALL be used for facial recognition biometrics (Sec. 3.4.1)
Applies to AAL	AAL2 and AAL3	AAL3 only	All	All	AAL2 and AAL3
Current State (Yes/No/Partial + details)	No - only password + SMS (HR Portal)	Yes (Financial System PIV/ CAC)	8+ chars but still enforce complexity	15 min timeout, no DBSC	N/A (no biometrics in use)
Gap?	Yes	No	Partial	Partial	N/A
Risk Rating (H/M/L)	High	–	Medium	Medium	–

Recommended Action / Compensating Control	Add WebAuthN / passkey or PIV option to HR Portal	–	Remove special character rules	Implement DBSC or extend to 30 - 60 min where safe (HR Portal)	–
Notes	Critical for EO 13681 and modern threats	PIV/CAC already meets this requirement	Improves usability	Usability win	Future consideration if biometrics are adopted

Exhibit - Gap Analysis

Risk Level	Definition	Action Needed	Why It Matters (Impact)
High	Blocks compliance with SP 800-63B-4 or exposes to major modern threats (phishing, credential stuffing, spoofing)	Immediate remediation (0-6 months); document in DIAS	Major Security upgrade. Directly counters the most common attack vector. Strong alignment with EO 13681 for PII systems.
Medium	Partial compliance or usability/equity impact; minor threat increase	Plan remediation (6-12 months); add compensating controls	Makes AAL3 more practical on modern commercial devices (smartphones, laptops) all while preserving strong security.
Low	Minor deviation or fully addressed by existing controls	Monitor only; no immediate action	Significant usability improvements, reduces user friction, and supports higher adoption of stronger authenticators.

Exhibit - Risk Rating

Phased Implementation Roadmap

The transition should be broken down into five distinct phases to manage risk and minimize user disruption.

Phase	Timeline	Primary Goal	Key Deliverables	Responsibilities	Success Criteria
1: Preparation	0-3 Months	Policy update & procurement	Updated AAL policy, DIRM completed for high-value systems, procurement of FIDO2 security keys/passkey support.	CIO, IAM Team, Privacy Officer	Policy officially approved; new authenticators procured.
2: Migration & Binding	3-9 Months	Introduce phishing-resistant MFA	Phishing-resistant options (e.g., WebAuthN) deployed for all AAL2/AAL3 systems; legacy SMS/Email OTP phased out where appropriate.	IAM Team	Phishing-resistant MFA adoption rate > 50% for AAL2 systems.
3: Session & Recovery	9-12 Months	Implement revised session rules	Device-Bound Session Credentials (DBSC) pilot; re-authentication timeouts relaxed; revised account recovery paths deployed.	IAM Team, InfoSec	User help desk tickets related to session timeouts reduced by 25%.
4: Testing & Rollout	12-15 Months	User acceptance and accessibility	Pilot testing on key user groups; independent accessibility checks completed; full-scale rollout.	IAM Team, End-User Support	Zero major user-reported issues post-rollout; full system migration achieved.

5: Monitoring & Continuous Improvement	Ongoing (Quarterly)	Operationalize continuous monitoring	Establish metrics dashboard for authenticator performance; Quarterly DIRM review cycles established.	InfoSec, IAM Team	Automated reporting on authenticator failure rates, fraud indicators, and recovery times.

AAL-Specific Compliance Checklists and Requirements

Agencies should use detailed checklists to ensure full compliance for each AAL. Cross-references to SP 800-63B-4 sections (e.g., Sec. 3 for authenticators, Sec. 5 for sessions) are essential for the most accurate guidance.

AAL1 Checklist Example

Requirement	SP 800-63B-4 Reference	Compliance Status	Implementation Notes
Single factor allowed	Sec. 3.1	[ ] Yes	Used only for low-impact, low-risk systems.
Verifier expected to make MFA available	Sec. 3.1.2	[ ] Yes	MFA is an available option for all users.
Memorized secrets (passwords) meet minimum 8 chars	Sec. 3.11.2	[ ] Yes	Password policy updated to remove complexity rules.
Account recovery path is available	Sec. 4	[ ] Yes	Email OTP or self-service options in place.

AAL2 Checklist Example

Requirement	SP 800-63B-4 Reference	Compliance Status	Implementation Notes
Two distinct authentication factors required	Sec. 3.2	[ ] Yes	Must combine two different factor types (e.g., knowledge + possession).
Phishing-resistant authenticator MUST be offered	Sec. 3.2.2	[ ] Yes	Passkeys, FIDO2, or PIV/CAC are options for users.
Authenticator uses approved cryptography	Sec. 3.2.3	[ ] Yes	Cryptographic modules used are approved.
Re-authentication (SHOULD) be relaxed (up to 12h/30-60m inactivity)	Sec. 5.2	[ ] Yes	Implemented DBSC to extend session while maintaining security.

AAL3 Checklist Example

Requirement	SP 800-63B-4 Reference	Compliance Status	Implementation Notes
Public-key cryptographic authenticator required	Sec. 3.3	[ ] Yes	Only PIV/CAC or high-assurance FIDO2 authenticators allowed.
Authenticator provides Verifier Impersonation Resistance (phishing resistance)	Sec. 3.3.2	[ ] Yes	All authenticators meet this standard.
Non-exportable cryptographic key in secure element/TPM/TEE	Sec. 3.3.3	[ ] Yes	Key material cannot be extracted from the authenticator.
All AAL2 requirements also apply		[ ] Yes	Inherit all base requirements from AAL2.

Sheet example with references to NIST for the most accurate guidance URL here: NIST SP 800-63-4 Digital Identity Guidelines

Best Practices, Considerations, and Common Pitfalls

Usability and Equity

- **Encourage Multi-Factor Options:** While AAL1 permits a single factor, encourage MFA options to improve baseline security and support user choice.
- **Local Biometric Verification:** Favor biometric implementations where verification occurs locally on the device (e.g., Face ID, fingerprint) rather than server-side, as this preserves privacy and typically enables stronger security.
- **Accessibility:** Ensure all MFA options are usable by people with disabilities (e.g., avoid time-sensitive challenges for cognitive disabilities).

Security

- **Prioritize Phishing Resistant (e.g., WebAuthN, PIV):** Make these the default and most encouraged options for all federal and high-value public users.
- **Use Fraud Indicators without Changing AAL:** Leverage continuous monitoring and fraud detection systems to dynamically challenge users or block sessions without permanently lowering or raising the static AAL designation.

Privacy

- **Tailor SP 800-52 Controls:** Ensure that the selection and operation of authenticators align with the identity assurance privacy controls from SP 800-52 (Please refer to section 3.5 as an example).
- **Redress Mechanisms:** Provide clear, accessible, and timely mechanisms for users to address issues like mistaken rejection, account lockout, or biometric performance errors.

Migration Tips

- **Phased Approach and Rollout:** Use a phased roadmap (as described above) to avoid sudden, disruptive changes. Start with the most security-aware user groups (e.g., IT staff).
- **Test Synced Authenticator Carefully:** While synced authenticators (like FIDO2 passkeys stored in cloud wallets) are allowed at AAL2, they are explicitly forbidden at AAL3 due to the non-exportable key requirement.

Pitfalls

- **Over-reliance on Legacy OTPs (SMS/Email):** These are easily phishable and should be phased out, especially for AAL2 systems.
- **Ignoring Session Monitoring:** Failing to implement continuous performance monitoring (Phase 5) means the agency is not compliant with the AAL-as-a-living-control mandate of SP 800-63-4.

Resources and References

- Official NIST Links:
  - NIST SP 800-63-4 Digital Identity Guidelines Suite
  - NIST Cryptographic Module Validation Program (CMVP) for FIPS validation
- Tools/Consortia:
  - FIDO Alliance (Passkeys and WebAuthn standards)
- Related Guidance:
  - NIST SP 800-53 (Security and Privacy Controls)
  - Executive Order 13681 (Federal Identity, Credential, and Access Management)
  - CISA guidance if available

Appendix

Glossary

Term	Definition (NIST SP 800-63-4 Context)
Phishing Resistant	An authenticator or authentication process that provides high assurance of **verifier impersonation resistance**, meaning a phishing site cannot collect the necessary information to authenticate to the legitimate service.
Non-Exportable Key	A private cryptographic key that is stored in a secure element (e.g., TPM, TEE, security key) such that it cannot be extracted or copied from the physical device. A requirement for AAL3.
Device-Bound Session Credentials (DBSC)	A session management technique where the session token is cryptographically bound to the user's device, significantly reducing the risk of session hijacking.
Presentation Attack Detection (PAD)	Techniques (e.g., liveness detection) used to detect and defeat attempts to spoof a biometric authentication system using artifacts like photos, videos, or deepfakes. Required for AAL2/AAL3 facial recognition.
Digital Identity Risk Management (DIRM)	The continuous process (introduced in SP 800-63-4 Base Volume) used to assess threats, tailor assurance levels, and monitor performance of digital identity services.
Digital Identity Acceptance Statement (DIAS)	A formal document (or record) created as part of the DIRM process to record and justify the final assurance level selection and associated controls for a digital service.

Templates

Digital Identity Acceptance Statement (DIAS) Excerpt for AAL

Field	Description	Justification/Reference
Digital Service Name	HR Portal	
Impact Level (High, Mod, Low)	Moderate (due to PII exposure)	Loss of PII (names, SSNs, financial data)
Baseline AAL	AAL2	Alignment with PII mandates (EO 13681).
Tailoring Applied	Yes: Threat (Phishing) and Usability	Mitigate phishing risk; improve user experience.
Final Approved AAL	AAL2	
Required Authenticators	Password + Phishing Resistant MFA (WebAuthn/PIV)	Phishing resistance must be offered at AAL2.
Compensating Controls	Implement DBSC to extend session time.	Mitigates Medium risk from session timeout.
Continuous Monitoring Metrics	Authenticator Failure Rate, Phishing Success Rate, Recovery Time.	Required by SP 800-63-4.

Flowcharts (AAL Selection)

AAL Selection Process Flow

1. **START:** Identify Digital Service / Transaction.
2. **Impact Assessment:** Determine impact level (Low/Moderate/High) of compromise.
3. **Baseline AAL:**
    - Low Impact $\to$ AAL1
    - Moderate Impact $\to$ AAL2
    - High Impact $\to$ AAL3

4. **Risk Tailoring (Threats):** Are high-level threats (e.g., targeted phishing) present?
    - **YES:** Increase AAL (e.g., AAL2 $\to$ AAL3) OR implement strong Compensating Control (e.g., DBSC, anti-phishing gateway).
    - **NO:** Proceed.
5. **Risk Tailoring (Privacy/Equity/Usability):** Do selected authenticators meet privacy/equity rules and support usability goals?
    - **NO:** Select alternative authenticators or implement Redress Mechanism.
    - **YES:** Proceed.
6. **FINAL AAL Determination:** Record final AAL and controls.
7. **DOCUMENT:** Complete Digital Identity Acceptance Statement (DIAS).
8. **END**

Federation Assurance Levels (FAL) Implementation Guide

The transition from **NIST SP 800-63-3** (the old standard) to **NIST SP 800-63-4** (the new standard) regarding **Federation Assurance Levels (FAL)** shifts focus from a rigid mapping of impact levels to a dynamic **Digital Identity Risk Management (DIRM)** process. While the fundamental three-tier structure of FAL remains, the new standard introduces more nuanced selection criteria for high-impact systems, new federated models like **subscriber-controlled wallets**, and mandatory **continuous evaluation**.

Implementation Guide: FAL 1, 2, and 3 Comparison

The following table compares the technical implementation and control objectives for federation across both standards.

Federation Assurance Level	SP 800-63-3 (Old Standard)	SP 800-63-4 (New Standard)
FAL1	Allows for bearer assertions signed by the Identity Provider (IdP) using approved cryptography.	Focuses on basic protection against **forged assertions**. Requires IdP signatures and audience restrictions.
FAL2	Adds the requirement that the assertion be **encrypted** so only the Relying Party (RP) can decrypt it.	Specifically aims to protect against **injection attacks** and forged assertions. Required whenever **personal information** is passed in an assertion.
FAL3	Requires the subscriber to present **proof of possession of a cryptographic key** referenced in the assertion.	Explicitly designed to protect against **IdP compromise**. Establishes very high confidence that communicated information matches CSP/IdP records.

Changes Required for a Compliant US Federal Agency

To maintain compliance with the new SP 800-63-4 standard, a federal agency must evolve its federation strategy from a one-time assessment to a continuous risk-management lifecycle.

1. Adopt the Five-Step DIRM Selection Process

Agencies can no longer rely on a simple high-water mark of impact to select an FAL. They **SHALL** follow the new **DIRM process**:

1. **Step 1: Define the Service:** Identify not just users, but all **impacted entities** (e.g., mission partners or communities).
2. **Step 2: Initial Impact Assessment:** Assess risks across five updated categories: Mission Delivery; Trust/Reputation; Information Access; Financial Loss; and Human/Environmental Safety.
3. **Step 3: Initial Selection:** Map the combined impact (Low/Moderate/High) to an initial FAL. Notably, high-impact services no longer default automatically to FAL3; agencies **SHALL** conduct a specific assessment of **IdP compromise risk** to determine if FAL2 or FAL3 is more appropriate.
4. **Step 4: Mandatory Tailoring:** Agencies **SHALL** assess the initial FAL against **Privacy, Customer Experience (CX), and Threat Resistance**. This may lead to implementing **compensating controls** (e.g., stricter auditing) or **supplemental controls** (e.g., phishing-resistant requirements).
5. **Step 5: Continuous Improvement:** Establish a program to monitor federation performance using data-driven metrics.

2. Implement New Performance Metrics

Agencies are now expected to track the health of their federated transactions. Required or recommended metrics include:

- **Unauthorized Access/Fraud:** Tracking confirmed, suspected, and reported fraud per authentication type.
- **Redress Resolution:** Monitoring the number of redress requests and the average time taken to resolve them.
- **IdP Integration:** Documenting reporting requirements for any integrated third-party IdP to ensure consistent data exchange for security and fraud prevention.

3. Support Modern Federated Models

The new standard explicitly addresses the rise of **Subscriber-Controlled Wallets** and **Attribute Bundles**.

- Agencies **SHOULD** prepare to support models where the user's device (wallet) acts as the IdP, presenting **CSP-signed attribute bundles** directly to the RP.
- If using these models, the RP **SHALL** establish a trust agreement, potentially through a **federation authority**, to accept assertions without needing a direct relationship with every individual wallet.

4. Enhanced Redress and Human Intervention

Federal agencies **SHALL** provide a formal issue-handling process for federation failures (e.g., a rejected assertion or account linking error).

- This process **SHALL** be accessible, trackable, and documented on a public-facing website.
- Agencies **SHALL** ensure human personnel are available to **override automated adjudication** outputs from algorithmic support mechanisms.

5. AI and Machine Learning Governance

If the agency's federation or fraud-detection systems use AI/ML:

- The use of AI/ML **SHALL** be disclosed to all RPs.
- The agency **SHALL** provide documentation on model training data sets, model update frequency, and testing results.
- A specific **privacy risk assessment** for all data processed by AI/ML systems must be performed and documented.

Analogy: Transitioning from the old FAL standard to the new one is like moving from a **physical passport check** to a **global "trusted traveler" network**. Under the old rules, you simply showed the right level of ID at the border (the FAL level). Under the new rules, the system is constantly communicating with other countries to share threat data (continuous evaluation), providing a digital "wallet" for your credentials (new models), and ensuring there is always a human officer available to help if the electronic gate fails (redress).