

SHARED DEVICES, INDIVIDUAL IDENTITIES: SECURE SSO AND SINGLE-LOGOUT ARCHITECTURE FOR MULTI-USER FRONTLINE MOBILE FLEETS

Mallikarjuna Yadav Bachu
Advisory Innovation Engineer, USA.

ABSTRACT

Frontline operations in retail, healthcare, warehousing, and field service increasingly equip teams with shared handheld devices rather than individually assigned hardware, driven by cost, device management overhead, and the practical reality that a given shift role, not a given person, is often what a device is provisioned for. This shared device model creates an identity and access management problem not well addressed by conventional single sign on, SSO, architectures, which were designed around the assumption that a device has a single, relatively persistent user. On a shared frontline device, a dozen or more workers may authenticate and deauthenticate across a single hardware unit in one day, and any failure to fully terminate the previous worker's session before the next worker begins theirs constitutes a session leakage event with direct privacy and compliance consequences. This article examines architecture patterns for secure multi user identity on shared frontline mobile fleets, with particular attention to single logout, SLO, the often neglected counterpart to single sign on that governs how reliably a session is torn down. Drawing on established identity protocol standards and security research published through early 2023, the analysis compares shared PIN, badge or near field communication tap in, device bound passkey, and biometric fast switch approaches across security, usability, implementation cost, and offline resilience, and proposes a hybrid design that combines a fast front channel identity switch with a verified back channel logout propagation step. The article presents comparative security and usability results, analyzes the specific failure modes that produce session leakage and incomplete logout, and closes with a decision framework and glossary for architecture and security teams building identity systems for shared device fleets.

Keywords:

single sign on, single logout, shared device identity, frontline mobile fleets, session management, FIDO2, OpenID Connect, zero trust architecture, mobile device management, workforce identity.

1. INTRODUCTION

The architecture of enterprise identity systems has, for most of its history, been built around a stable pairing between a person and a device. Single sign on, SSO, protocols such as Security Assertion Markup Language and OpenID Connect were designed primarily to let one person move between many applications without repeated logins, not to let many people move through one device in rapid succession. Frontline operations, including retail point of sale, warehouse scanning, home health visits, and delivery fleets, increasingly invert this assumption: a single ruggedized handheld or tablet is shared across a shift roster, handed from one worker to the next at shift change, break coverage, or task reassignment, often multiple times within a single working day.

This shared device model is attractive for straightforward operational reasons. Provisioning and maintaining one device per role rather than one per person reduces hardware and mobile device management cost, simplifies device lifecycle management, and avoids the logistical burden of distributing and recovering personal hardware across a workforce that may include high turnover, part time, or seasonal staff. The tradeoff is that identity, authentication, and especially logout must now operate correctly many times per device per day, under time pressure, frequently without reliable network connectivity, and without the luxury of a slow, deliberate login flow that a once a day personal device login can tolerate.

Single logout, SLO, is the specific capability this article treats as under examined relative to its operational importance. Most identity architecture discussion focuses on sign on: how to authenticate a user quickly, securely, and with minimal friction. On a shared device, sign on quality matters, but sign out quality matters just as much, because an incompletely terminated session from the prior worker is not merely an inconvenience, it is a live credential and a live data exposure sitting inside a device that has just been handed to someone else. A logout that only closes the top level application, while leaving background services, cached tokens, or other installed applications still authenticated as the prior worker, creates exactly this exposure, and is a failure mode this article returns to repeatedly in the security and error analysis sections.

The article is organized around four questions. First, what makes the shared device identity problem structurally different from the personal device SSO problem that most identity architecture guidance assumes. Second, what architecture patterns are available for fast, secure multi user identity switching on a shared handheld device, and how do they compare on security and usability. Third, how should single logout be designed so that it reliably propagates across every application and service session on the device, not just the foreground application, including under intermittent connectivity. Fourth, what does this mean in practice for teams building or operating identity systems for frontline fleets. The remainder of the article develops an evaluation framework around these questions and closes with a decision framework for practitioners.

2. BACKGROUND AND RELATED WORK

2.1. Single Sign On Protocol Foundations

Modern enterprise SSO rests on a small number of widely deployed standards. The Security Assertion Markup Language, SAML 2.0, formalized by OASIS in 2005, defines an assertion based protocol in which an identity provider vouches for a user to a service provider, historically dominant in large enterprise and education deployments. The OAuth 2.0 Authorization Framework, defined by Hardt in RFC 6749 in 2012, defines a delegation protocol for issuing scoped access tokens, and OpenID Connect Core 1.0, published by Sakimura and colleagues through the OpenID Foundation in 2014, layers an identity and authentication model on top of OAuth 2.0, becoming the dominant modern foundation for consumer and increasingly enterprise identity. JSON Web Tokens, standardized by Jones, Bradley, and Sakimura in RFC 7519 in 2015, provide the compact, signed token format most OpenID Connect deployments use to carry identity and session claims.

Native and mobile specific guidance followed as these protocols moved beyond browser based web applications. OAuth 2.0 for Native Apps, published by Denniss and Bradley as RFC 8252 in 2017, established best practices for running OAuth flows inside installed mobile applications rather than purely in a browser, a pattern directly relevant to the multi application session propagation problem central to this article's single logout analysis. The OAuth 2.0 Threat Model and Security Considerations document, published by Lodderstedt, McGloin, and Hunt as RFC 6819 in 2013, remains a widely referenced catalog of the attack classes, including token leakage and replay, that this article's error and risk analysis in Section 9 draws on when framing shared device specific failure modes.

2.2. Single Logout Mechanics

Front Channel and Back Channel Logout

OpenID Connect's logout family of specifications defines two complementary propagation mechanisms. Front channel logout notifies relying party applications of a logout event through the user's browser or in app web view, requiring the user agent to be active and able to load each relying party's logout endpoint in turn. Back channel logout instead has the identity provider directly call each relying party's logout endpoint server to server, independent of whether the user's device or browser remains active, which matters directly for a shared device scenario where the device may be handed to the next worker before every front channel notification has had a chance to complete. RP Initiated Logout defines the complementary flow by which a relying party application can request that the identity provider terminate a user's broader session, not merely the local application session.

2.3. Device Bound Authentication Standards

The Web Authentication specification, published as a W3C Recommendation in 2019, and the FIDO Alliance's Client to Authenticator Protocol, CTAP2, published in 2020, together define the modern passkey model: a device bound or synchronized public key credential that avoids transmitting a shared secret, such as a password or PIN, over the network or storing one on a server. These standards were designed primarily around the single user personal device model, and this article's architecture discussion in Section 4 examines how device bound credentials can be adapted to a multi user shared device without undermining the non shared secret property that is their core security advantage. A subsequent update, Web Authentication Level 2, reached W3C Recommendation status in 2021 and refined resident credential and user verification semantics in ways directly relevant to the multi credential storage question in Section 4.3. Underlying key management practice for the asymmetric credentials these standards depend on follows general purpose guidance such as the National Institute of Standards and Technology's key management recommendations, published as Special Publication 800-57 Part 1 Revision 5 in 2020.

2.4. Identity Assurance and Zero Trust Guidance

The United States National Institute of Standards and Technology's Digital Identity Guidelines, published by Grassi, Garcia, and Fenton as Special Publication 800-63-3 in 2017, establish widely referenced authenticator assurance levels that inform how security conscious organizations classify the strength of a given sign on method. The companion Zero Trust Architecture guidance, published by Rose, Borchert, Mitchell, and Connelly as Special

Publication 800-207 in 2020, reframes access decisions around continuous verification of a specific session and device state rather than a one time perimeter check, a framing this article's single logout analysis depends on, since a zero trust model treats an un terminated prior session as an active and continuing trust violation rather than a historical event.

A separate body of research addresses continuous and transparent authentication on personal mobile devices, including the survey by Patel, Chellappa, Chandra, and Barbello published in 2016, which catalogs behavioral and biometric signals usable to continuously verify a device's current user without repeated explicit logins. This article's biometric fast switch approach in Section 4 draws on this continuous authentication framing but applies it specifically to the problem of distinguishing which of several known workers is currently holding a shared device, rather than the single user anomaly detection framing most of that literature addresses.

3. THE SHARED DEVICE PROBLEM

Table 1 contrasts the shared frontline device model against the conventional personal device model along the dimensions most relevant to identity architecture. The contrast is not simply a matter of degree, several of these dimensions represent qualitatively different requirements that a personal device SSO deployment is never forced to satisfy.

Dimension	Personal device model	Shared frontline device model
Logins per device per day	Typically one, persisting across the day	Often five to twenty or more across a shift roster
Consequence of incomplete logout	Low; same user resumes next session	High; next user may access prior user's data or identity
Acceptable login latency	Several seconds is tolerable	Must often be under two to three seconds to avoid shift disruption
Device possession assumption	Device usually associated with one person	Device explicitly provisioned for a role, not a person
Connectivity assumption	Generally online	Frequently intermittent, especially in warehouses or rural routes
Biometric enrollment model	One enrolled user per device	Multiple enrolled users per device, selected at runtime

Table 1. Table 1 compares identity relevant dimensions between the personal device model most SSO guidance assumes and the shared frontline device model this article addresses.

3.1. Fleet Composition and Heterogeneity

Frontline fleets are rarely homogeneous. Figure 1 shows the device composition of a representative multinational frontline operation spanning retail, warehouse, and field service functions, illustrating that any identity architecture proposal must function across several device platforms and operating systems simultaneously rather than targeting a single ecosystem.

Figure 1. Composition of a Representative Frontline Device Fleet




Figure 1. Composition of a representative frontline device fleet by device category.

This heterogeneity constrains the architecture choices available in Section 4. A design that depends on a specific hardware security module or biometric sensor present only on one platform cannot serve the full fleet without a fallback path for devices lacking that capability, a constraint that recurs throughout the comparative analysis in this article.

3.2. Session Leakage as the Central Risk

The specific risk this article treats as central is session leakage: any condition in which a session, credential, cached data, or active authentication state belonging to one worker remains accessible after the device has been handed to a different worker. Session leakage can occur at several layers simultaneously, including the operating system level session, the identity provider level session, and the per application local session state, and a robust single logout design must address all three rather than assuming that terminating one automatically terminates the others.

4. ARCHITECTURE APPROACHES FOR MULTI USER IDENTITY

4.1. Shared PIN

The simplest and most widely deployed approach assigns a short numeric PIN to each worker, entered at the start of a session to establish identity without necessarily involving the identity provider directly for every switch. This approach requires no additional hardware and switches quickly, but depends on a shared secret that can be observed, shared informally between coworkers, or reused across a long employment tenure, and provides comparatively weak assurance of who is actually holding the device at any given moment.

4.2. Badge or Near Field Communication Tap In

This approach pairs each worker with a physical badge or card, already common in frontline environments for building access or time and attendance, and extends it to device identity switching through a near field communication, NFC, tap against the device. This leverages hardware many frontline organizations have already deployed for other purposes, offers meaningfully stronger assurance than a shared PIN since it requires physical possession of a specific credential, and switches quickly, but introduces a dependency on every device having NFC hardware and every worker having their badge available and undamaged.

4.3. FIDO2 Passkey Per User

Adapting the device bound passkey model from Section 2.3 to a shared device requires storing multiple resident credentials on a single device, one per enrolled worker, and presenting a worker selection step before or during the authenticator ceremony, since the standard flow was designed around a single resident user per device. This approach offers the strongest cryptographic assurance of the options considered, since it avoids any shared secret entirely, but the selection and ceremony steps typically add latency relative to a tap in or PIN entry, and multi credential resident key management is not uniformly well supported across the authenticator and operating system combinations present in a heterogeneous fleet.

4.4. Biometric Fast Switch

A biometric fast switch approach uses a face or fingerprint match against a small, locally stored set of enrolled worker templates to identify which worker is currently holding the device, without requiring an explicit PIN, tap, or ceremony step at all. This offers the fastest switch time of the options considered and requires no action beyond picking up the device, but requires biometric sensor hardware present on the device, raises distinct data protection considerations around storing biometric templates for an entire shift roster rather than a single owner, and performs less reliably in some frontline conditions, such as gloved hands for fingerprint sensors or personal protective equipment covering part of the face for facial recognition.

4.5. Comparative Scoring Across Approaches

Figure 2 scores all four approaches, plus the hybrid design developed in Section 5, across five dimensions: security, switch speed, user friction, implementation cost, and offline resilience. No single approach dominates every dimension. The shared PIN scores well on implementation cost and offline resilience but poorly on security. The FIDO2 passkey approach scores well on security but comparatively poorly on implementation cost given the multi credential management burden described in Section 4.3. This pattern of mutually exclusive strengths is the central motivation for the hybrid design presented in Section 5.

Figure 2. Multi User Identity Approaches Scored Across Five Dimensions (scale of 1 to 5, higher is better unless noted)




Figure 2. Comparative scoring of multi user identity approaches across five dimensions.

5. SINGLE LOGOUT DESIGN PATTERNS

5.1. The Hybrid Design

The hybrid design proposed in this article pairs a fast front channel identity switch, typically badge tap in or biometric match, with a verified back channel logout step that must complete, or visibly fail over to a safe fallback state, before the new worker's session is considered fully established. In practice, this means the device displays the new worker's interface quickly, satisfying the switch speed requirement from Section 3, while the back channel logout call to each relying party application's session store runs concurrently, and any application that has not confirmed logout of the prior session within a short grace period is forcibly suspended rather than left running with ambiguous session state.

5.2. Session Lifecycle Across a Shift Handoff

Figure 6 illustrates the resulting session lifecycle across a representative shift handoff under the hybrid design, decomposing the handoff into its constituent phases. The combined logout propagation and device handoff window occupies well under a minute even though it comprises several distinct steps, each of which is individually fast but which must complete in the correct order to avoid session leakage.




Figure 6. Session lifecycle timeline across a representative shift handoff event.

The idle timeout window shown in Figure 6 is a deliberate design element, not merely an incidental delay: requiring a short period of detected inactivity before a handoff is accepted as legitimate reduces the chance that a brief, accidental device transfer, such as a worker momentarily passing a device to a colleague to show something on screen, triggers a full logout and re authentication cycle unnecessarily.

5.3. Handling Intermittent Connectivity

Back channel logout, as described in Section 2.2, depends on the identity provider being able to reach each relying party's logout endpoint, which cannot be assumed on a device operating in a warehouse basement or a rural delivery route with unreliable connectivity. The hybrid design addresses this by maintaining a local, signed logout intent record on the device itself the moment a handoff begins, independent of network reachability, and enforcing that record at the operating system and application sandboxing level immediately, with the network confirmed back channel logout call queued to complete and update central session records as soon as connectivity returns. This ensures that the device itself never presents the prior worker's session to the new worker regardless of connectivity, even though the central identity provider's record of that session's termination may be briefly delayed.

6. Evaluation Methodology

6.1. Simulated Fleet and Shift Patterns

The comparative results in Sections 7 and 8 are drawn from a simulated frontline fleet constructed to approximate shift patterns typical of retail, warehouse, and field service operations, including scheduled shift changes, unscheduled break coverage handoffs, and occasional device reassignment mid shift. Each simulated session included realistic variation in connectivity, ranging from continuous connectivity to extended offline periods, and in handoff circumstances, including both orderly handoffs where the prior worker explicitly signs out and disorderly handoffs where a device is simply picked up by the next worker without an explicit prior sign out action.

6.2. Metrics

Security outcomes are measured primarily as the session leakage rate, the percentage of handoffs in which some element of the prior worker's session remained accessible after the new worker's session began, broken down by the specific layer at which leakage occurred. Usability outcomes are measured as identity switch time, the elapsed time from device pickup to a fully usable, correctly attributed session, and as the rate of failed or retried switch attempts. Table 2 summarizes the full metric set used across the evaluation.

Metric	Definition	Primary section
Session leakage rate	Share of handoffs where prior worker session data remained accessible	Section 7
Logout propagation latency	Time from handoff initiation to confirmed logout across all sessions	Section 7
Identity switch time	Time from device pickup to a fully usable, correctly attributed session	Section 8
Switch failure or retry rate	Share of handoffs requiring a retry or manual fallback	Section 8
Offline resilience	Share of handoffs completed correctly without network connectivity	Section 8

Table 2. Table 2 defines the metrics used throughout the evaluation in Sections 7 and 8, and the section in which each is primarily reported.

6.3. Approaches and Device Tiers Compared

All five approaches introduced in Sections 4 and 5, shared PIN, badge or NFC tap in, FIDO2 passkey, biometric fast switch, and the hybrid design, were evaluated across the three device tiers summarized in Table 3, approximating the range of ruggedized and consumer grade hardware found in a typical frontline fleet.

Device tier	Representative hardware	Biometric sensor availability	NFC availability
Legacy ruggedized	Older generation handheld scanners and terminals	Rare	Common
Current ruggedized	Current generation handheld scanners and terminals	Common, fingerprint	Common
Consumer grade	Current generation smartphones and tablets	Common, face or fingerprint	Common

Table 3. Table 3 defines the three device tiers used in the comparative evaluation, approximating fleet hardware heterogeneity consistent with Figure 1.

7. SECURITY RESULTS

7.1. Session Leakage Rate by Approach

Table 4 reports session leakage rate by approach, aggregated across all device tiers and both orderly and disorderly handoff conditions. The shared PIN approach shows the highest leakage rate, driven primarily by workers who stay logged in deliberately to avoid repeated PIN entry, a behavior that undermines the approach regardless of its underlying protocol design. The hybrid design shows the lowest leakage rate of any approach, including the FIDO2 passkey approach alone, because its enforced local logout intent record, described in Section 5.3, closes the specific gap that a pure credential based approach does not address by itself: a strong credential for the new worker says nothing about whether the prior worker's session was actually torn down.

Approach	Session leakage rate	Most common leakage layer
Shared PIN	11.4 percent	Application level, worker stays logged in across handoff
Badge or NFC tap in	4.8 percent	Background service session not tied to tap event
FIDO2 passkey	3.1 percent	Background service session not tied to ceremony
Biometric fast switch	2.6 percent	Operating system session not re-evaluated on match
Hybrid design	0.6 percent	Rare back channel confirmation timeout under poor connectivity

Table 4. Table 4 reports aggregate session leakage rate and most common leakage layer by approach, across all device tiers and handoff conditions in the evaluation.

7.2. Authentication and Session Handoff Funnel

Figure 5 traces a representative handoff funnel under the hybrid design across ten thousand simulated shift changes, showing where attempts are lost at each stage. The largest single drop occurs between login attempt and credential verification, driven primarily by mistyped PINs or missed badge taps rather than by any back end failure, while losses in the later stages, device trust check and prior session clearance, are smaller and driven mostly by connectivity related delays rather than outright failures.

Figure 5. Authentication and Session Handoff Funnel per 10,000 Shift Changes




Figure 5. Authentication and session handoff funnel per ten thousand shift changes under the hybrid design.

7.3. Logout Propagation Latency

Logout propagation latency, the time from handoff initiation to confirmed logout across every relying party session on the device, averaged under three seconds for the hybrid design under normal connectivity, and remained under ten seconds even under the degraded connectivity conditions included in the simulation, since the local logout intent enforcement described in Section 5.3 removes the network round trip from the critical path for preventing session leakage, leaving only the central record update, rather than the leakage prevention itself, dependent on connectivity.

8. Usability and Operational Results

8.1. Identity Switch Time Distribution

Figure 4 shows the distribution of identity switch time across the five approaches. The badge or NFC tap in approach produces the fastest and most tightly clustered switch times, since a tap is a single, low variance physical action. The FIDO2 passkey approach shows both the slowest median switch time and the widest spread, driven by the added worker selection step described in Section 4.3, which varies in duration depending on how many workers are enrolled on a given device and how the selection interface is implemented. The hybrid design's switch time distribution closely tracks whichever front channel method, badge tap in or biometric match, it is paired with, since the back channel logout verification runs concurrently rather than adding to the visible switch time in the common case.

Figure 4. Distribution of Identity Switch Time by Approach




Figure 4. *Distribution of identity switch time by approach.*

8.2. Feature Support Profile

Figure 3 summarizes the hybrid design's support across six properties that recur throughout this article's requirements discussion. The hybrid design scores highest on sub two second identity switch and resistance to session replay after handoff, reflecting its core design goals, and scores lowest, though still positively, on fully offline operation, since the central session record update, as noted in Section 7.3, still depends on eventual connectivity even though local leakage prevention does not.

Figure 3. Feature Support Profile of the Recommended Hybrid De




Figure 3. Feature support profile of the recommended hybrid design across six operational properties.**8.3. Switch Failure and Retry Rate**

Table 5 reports switch failure or retry rate by approach and device tier. Older ruggedized hardware shows a materially higher retry rate across every approach, consistent with worn physical buttons, degraded NFC antennas, and lower resolution biometric sensors common on aging frontline hardware, reinforcing that identity architecture decisions cannot be evaluated independently of the fleet's actual hardware condition, which tends to skew older and more worn than a vendor's reference specifications would suggest.

Approach	Legacy ruggedized retry rate	Current ruggedized retry rate	Consumer grade retry rate
Shared PIN	3.1 percent	2.4 percent	2.0 percent
Badge or NFC tap in	6.8 percent	2.9 percent	2.2 percent
FIDO2 passkey	Not applicable	5.4 percent	4.1 percent
Biometric fast switch	Not applicable	7.2 percent	3.8 percent
Hybrid design	6.5 percent	3.1 percent	2.4 percent

Table 5. Table 5 reports switch failure or retry rate by approach and device tier. Legacy ruggedized hardware in this evaluation lacked biometric sensors entirely, hence the not applicable entries.**9. Error and Risk Analysis****9.1. Specific Failure Modes**

Beyond the aggregate session leakage rate reported in Section 7, four specific failure modes recur across approaches and merit individual attention from architecture and security teams.

1. Stale session reuse. A background application retains an authenticated session from the prior worker because it was never notified of the handoff, and the next worker unintentionally acts, or appears to act, under the prior worker's identity within that application.
2. Partial logout propagation. The front channel switch for the primary application completes correctly, but one or more secondary applications or background services do not receive, or do not act on, the corresponding logout signal, leaving an inconsistent identity state across the device.
3. Token replay after handoff. A cached access or refresh token issued to the prior worker remains valid and is reused, deliberately or accidentally, after a handoff that did not include a true revocation at the identity provider, as opposed to a purely local, device side logout.
4. Credential sharing workaround. Workers, under time pressure, share a PIN, badge, or in rare observed cases attempt to register each other's biometric data to the same enrollment slot, as an informal workaround to a switch process perceived as too slow, undermining the architecture's assurance properties regardless of its technical design.

9.2. Risk Weighting by Failure Mode

Table 6 assigns a qualitative likelihood and impact rating to each failure mode, synthesizing the quantitative results in Section 7 with the operational context in Section 9.1. Token replay after handoff is rated as comparatively low likelihood but high impact, since it requires a specific combination of a cached token and a motivated actor, but can expose data belonging to a worker who is no longer even present at the device if it does occur. Credential sharing workaround is rated as comparatively higher likelihood specifically for approaches with poor usability under time pressure, underscoring that usability and security are not independent concerns in this domain, a usability failure directly produces a security failure through this specific mechanism.

Failure mode	Likelihood	Impact	Most effective mitigation
Stale session reuse	Moderate	Moderate	Enforced local logout intent record, Section 5.3
Partial logout propagation	Moderate	High	Back channel logout with grace period enforcement
Token replay after handoff	Low	High	Short lived tokens with true revocation, not local clearing alone
Credential sharing workaround	Moderate to high under poor usability	High	Fast, low friction front channel switch method

Table 6. Table 6 presents a qualitative risk rating for each failure mode identified in Section 9.1, alongside the mitigation most directly effective against it.

10. Case Studies Across Frontline Fleets

The following case studies synthesize representative deployment patterns reported across published vendor material and industry commentary on shared device identity management through early 2023. They are presented as composite illustrations of recurring design patterns rather than as single named deployments.

10.1. Retail Point of Sale Shift Changes

A multi site retail operation sharing point of sale handhelds across cashier shifts found that its original shared PIN approach suffered from exactly the stay logged in behavior described in Section 7.1, since cashiers under customer facing time pressure frequently chose not to log out at the end of a shift. Moving to a badge tap in front channel, paired with the hybrid design's enforced back channel logout, reduced session leakage incidents substantially without slowing down the already fast tap based switch workers had grown accustomed to, since the back channel verification runs invisibly in the background.

10.2. Home Health and Mobile Clinical Visits

A home health provider issuing shared tablets to visiting clinicians, rotated across a roster larger than the device pool, treated token replay after handoff as a specifically unacceptable risk given the clinical and regulatory sensitivity of patient data, and prioritized the FIDO2 passkey approach's strong revocation properties over its slower switch time, accepting the usability cost documented in Section 8.1 as a deliberate tradeoff given the data sensitivity involved, consistent with the general pattern in Table 7 of weighting regulatory sensitivity above raw switch speed for clinical workflows.

10.3. Warehouse and Last Mile Delivery Fleets

A logistics operator running scan intensive warehouse and last mile delivery operations across a large fleet of aging ruggedized handhelds, consistent with the legacy ruggedized tier in Table 3, found that biometric sensors were simply not present on most of its deployed hardware, and adopted badge tap in as its primary front channel method by necessity rather than preference, while still layering the hybrid design's back channel logout verification on top, demonstrating that the hybrid pattern's back channel component is separable from any specific front channel choice and can be deployed incrementally across a heterogeneous, aging fleet.

11. Discussion

The results in Sections 7 through 9 converge on a consistent theme: on a shared frontline device, sign on strength and sign off reliability are separable engineering problems that must both be solved well, and solving only one, however thoroughly, leaves the other as an open attack surface. A strong front channel credential, such as a FIDO2 passkey, meaningfully reduces the risk that an unauthorized person can sign in as a given worker, but does nothing by itself to guarantee that the previous worker's session has actually been torn down, which is precisely why the FIDO2 passkey approach alone still shows a non trivial session leakage rate in Table 4, concentrated specifically in the background service layer that a credential ceremony does not touch.

This separation argues for treating single logout as a first class architecture concern with its own explicit design, evaluation metrics, and failure mode analysis, rather than as an assumed consequence of a strong sign on method. The hybrid design's core contribution in this article is not a novel credential mechanism, it deliberately reuses existing front channel methods, but a logout enforcement layer that functions correctly under the connectivity and time pressure constraints specific to frontline shared device operation, which the underlying SSO and passkey standards reviewed in Section 2 were not originally designed to address.

A secondary theme concerns the coupling between usability and security identified in Section 9.2. Because credential sharing workaround risk rises specifically when a switch method is perceived as slow or unreliable, an architecture decision that optimizes purely for cryptographic assurance while neglecting switch speed can produce a net security regression in practice, as workers route around a slow but theoretically strong mechanism. This suggests that usability testing under realistic time pressure should be treated as a security requirement for shared device identity systems, not merely a user experience nice to have evaluated separately from the security review.

12. Limitations

1. Simulated rather than live fleet data. The quantitative results in Sections 7 and 8 are drawn from a simulated frontline fleet constructed to approximate realistic shift and connectivity patterns, since no single public benchmark combining shared device handoff behavior, session leakage incidents, and connectivity conditions was available at the time of writing, which introduces some uncertainty into the absolute figures reported.
2. Hardware condition variability. Table 5's device tier results approximate representative hardware condition within each tier, but actual sensor wear, battery health, and antenna condition vary considerably across real deployed fleets and may shift retry rates beyond what this article's tiers capture.
3. Credential sharing workaround is difficult to measure directly. The risk weighting in Section 9.2 relies on qualitative industry commentary and operational judgment rather than directly observed incident data, since credential sharing is, by its nature, a behavior organizations rarely instrument or disclose in detail.
4. Regulatory variation not modeled. Specific compliance obligations around biometric data collection, session logging, and data residency vary by jurisdiction and industry in ways this article does not attempt to model exhaustively, and organizations should treat Section 10's case studies as illustrative rather than as compliance guidance for any specific regulatory regime.
5. Protocol specification evolution. Several of the OpenID Connect logout specifications referenced in Section 2.2 continued to evolve toward final status around the time of this article's preparation, and implementers should verify the current status and exact semantics of these specifications against the OpenID Foundation's published documents rather than relying solely on this article's summary.

13. Recommendations for Practitioners

Table 7 consolidates this article's findings into a decision framework organized by fleet characteristic, intended as a starting point for teams selecting a multi user identity approach for a shared frontline device fleet.

Fleet characteristic	Recommended front channel	Back channel requirement
Mixed age fleet with inconsistent NFC or biometric hardware	Badge or NFC tap in, with PIN fallback	Hybrid back channel logout, mandatory
Modern fleet with consistent biometric sensors	Biometric fast switch	Hybrid back channel logout, mandatory
High regulatory sensitivity, for example clinical or financial data	FIDO2 passkey per user	Hybrid back channel logout with short lived token revocation
Frequent extended offline operation	Badge or NFC tap in, lowest connectivity dependency	Local logout intent enforcement prioritized over central revocation

Table 7. Table 7 recommends a front channel and back channel configuration by fleet characteristic, synthesizing the security, usability, and risk findings in Sections 7 through 9.

- Treat single logout as a first class design requirement with its own explicit metrics, not an assumed byproduct of a strong sign on credential, consistent with the separation argued in Section 11.
- Enforce logout locally on the device at the moment a handoff begins, independent of network reachability, and treat central session revocation as an eventually consistent follow up rather than a precondition for safe local behavior.
- Measure switch speed under realistic time pressure as part of the security review, not solely the usability review, given the credential sharing workaround risk documented in Section 9.2.
- Select front channel methods based on actual fleet hardware condition, not vendor reference specifications, given the aging hardware retry rate gap documented in Table 5.
- Prefer short lived tokens with true revocation over purely local session clearing wherever the data sensitivity of the workflow justifies the added implementation cost, per the token replay mitigation in Table 6.

14. Future Work

Several directions follow from the gaps identified in this article. Standardized extensions to the FIDO2 and WebAuthn credential model specifically addressing multi resident credential selection on a shared device, rather than treating multi user support as an application layer workaround, would reduce the implementation burden documented in Section 4.3. Formal verification of back channel logout enforcement under adversarial connectivity

conditions, beyond the simulated connectivity patterns used in this article's evaluation, would strengthen confidence in the hybrid design's offline behavior guarantees. Field studies directly instrumenting credential sharing workaround behavior, rather than relying on qualitative industry commentary as in Section 9.2, would sharpen the risk weighting in Table 6 considerably. Finally, as passkey synchronization across devices becomes more widely deployed, research into how synchronized, as opposed to purely device bound, credentials interact with the shared device model would be valuable, since synchronization introduces session state considerations beyond the single device scope this article addresses.

15.CONCLUSION

Shared frontline mobile devices invert a core assumption behind most enterprise single sign on architecture: that a device has one relatively persistent user rather than a rapidly rotating shift roster. This inversion elevates single logout from a secondary concern to a primary architecture requirement, since the dominant risk in this environment is not unauthorized sign on but incomplete sign off, which silently exposes one worker's session and data to the next person holding the device. Across the approaches examined in this article, no single front channel credential method dominates every requirement simultaneously, and the hybrid design's core insight, pairing a fast front channel switch with an enforced, connectivity resilient back channel logout, addresses the specific gap that credential strength alone cannot close. For practitioners, the most actionable conclusion is to select a front channel method based on actual fleet hardware and operational conditions following the framework in Table 7, while treating enforced single logout as a non negotiable architectural layer regardless of which front channel method is chosen.

REFERENCES

- Barker, E. (2020). *Recommendation for Key Management, Part 1: General*. NIST Special Publication 800-57 Part 1 Revision 5.
- Denniss, W., and Bradley, J. (2017). *OAuth 2.0 for Native Apps*. RFC 8252, Internet Engineering Task Force.
- FIDO Alliance (2020). *Client to Authenticator Protocol, CTAP2, Specification*.
- Grassi, P. A., Garcia, M. E., and Fenton, J. L. (2017). *Digital Identity Guidelines*. NIST Special Publication 800-63-3.
- Hardt, D. (2012). *The OAuth 2.0 Authorization Framework*. RFC 6749, Internet Engineering Task Force.
- Hodges, J., Jones, J. C., Jones, M. B., Kumar, A., and Lindemann, R. (2021). *Web Authentication: An API for Accessing Public Key Credentials, Level 2*. W3C Recommendation.
- Jones, M., Bradley, J., and Sakimura, N. (2015). *JSON Web Token, JWT*. RFC 7519, Internet Engineering Task Force.
- Lodderstedt, T., McGloin, M., and Hunt, P. (2013). *OAuth 2.0 Threat Model and Security Considerations*. RFC 6819, Internet Engineering Task Force.
- OASIS (2005). *Security Assertion Markup Language, SAML, V2.0 Technical Overview*.
- OpenID Foundation (2022). *OpenID Connect RP Initiated Logout 1.0*.
- OpenID Foundation (2022). *OpenID Connect Back Channel Logout 1.0*.
- Patel, V. M., Chellappa, R., Chandra, D., and Barbell, B. (2016). *Continuous User Authentication on Mobile Devices: Recent Progress and Remaining Challenges*. IEEE Signal Processing Magazine, volume 33.
- Rose, S., Borchert, O., Mitchell, S., and Connelly, S. (2020). *Zero Trust Architecture*. NIST Special Publication 800-207.
- Sakimura, N., Bradley, J., Jones, M., de Medeiros, B., and Mortimore, C. (2014). *OpenID Connect Core 1.0*. OpenID Foundation.
- W3C (2019). *Web Authentication: An API for Accessing Public Key Credentials, Level 1*. W3C Recommendation.

Appendix A: Combined Metric Summary

Table 8 consolidates session leakage rate, identity switch time, and offline resilience for all five evaluated approaches at the current ruggedized device tier, providing a single reference table for the detailed per metric figures presented separately in Sections 7 and 8.

Approach	Session leakage rate	Median switch time	Offline resilience
Shared PIN	11.4 percent	1.8 s	High
Badge or NFC tap in	4.8 percent	1.2 s	High
FIDO2 passkey	3.1 percent	2.6 s	Moderate
Biometric fast switch	2.6 percent	1.0 s	Moderate
Hybrid design	0.6 percent	1.3 s	High

Table 8. Table 8 consolidates key metrics across all evaluated approaches at the current ruggedized device tier.

Appendix B: Glossary of Terms

- **Back channel logout.** A logout propagation mechanism in which the identity provider notifies relying party applications of a logout event through a direct server to server call, independent of the user's browser or device being active.
- **Device bound credential.** A cryptographic credential, such as a FIDO2 passkey, generated and stored such that its private key never leaves the specific device or authenticator that created it.
- **Front channel logout.** A logout propagation mechanism that notifies relying party applications through the user's active browser or in app web view.
- **Hybrid design.** The architecture proposed in this article, pairing a fast front channel identity switch with an enforced, connectivity resilient back channel logout step.
- **Identity switch time.** The elapsed time from device pickup to a fully usable, correctly attributed session for the incoming worker.
- **Relying party.** An application or service that depends on an identity provider to authenticate users rather than managing credentials itself.
- **Session leakage.** Any condition in which a session, credential, or cached data belonging to a prior device user remains accessible after the device has been handed to a different user.
- **Shared frontline device.** A handheld or tablet device provisioned for a shift role rather than an individual person, used by multiple workers across a working day.
- **Single logout, SLO.** The capability to reliably and completely terminate a user's session across every application and service on a device, as a counterpart to single sign on.
- **Single sign on, SSO.** The capability for a user to authenticate once and gain access to multiple independent applications or services without repeated logins.
- **Zero trust architecture.** A security model that treats every access request as requiring continuous verification of current session and device state, rather than relying on a one time perimeter check.

Appendix C: Implementation Checklist

Table 9 provides a condensed, pre deployment checklist synthesizing the architecture and risk guidance developed across Sections 5, 7, 9, and 13, organized as discrete, independently verifiable items for an implementation or security review team.

Checklist item	Verifies	Related section
Local logout intent is enforced at the operating system or sandbox level before network confirmation	Session leakage is prevented even if connectivity is unavailable	Section 5.3
Back channel logout calls every relying party application, not only the foreground application	Partial logout propagation cannot occur	Section 9.1
Access and refresh tokens issued per worker session are short lived with true revocation support	Token replay after handoff cannot succeed even if a token is cached	Section 9.2
Front channel switch method is tested under realistic time pressure, not only in a controlled test environment	Credential sharing workaround risk is minimized	Section 9.2
Device tier and hardware condition are assessed before selecting a front channel method	Switch failure or retry rate matches planning assumptions	Section 8.3

Idle timeout window is tuned to avoid accidental triggering during brief, legitimate device transfers	Unnecessary re authentication cycles are avoided	Section 5.2
---	--	-------------

Table 9. Table 9 provides a condensed pre deployment checklist for teams implementing the hybrid design or evaluating an existing shared device identity architecture against this article's findings.