POLITECNICO DI TORINO

Master's Degree in Computer Engineering

Master's Degree Thesis

Home Automation Applications for Older Adults (DORA)

Supervisor

Prof. Luca ARDITO

Candidate

YUJIE MU

December 2025

Abstract

Population ageing increases demand for trustworthy, usable home support. Many older adults wish to remain independent, yet face risks from falls, chronic conditions and social isolation, while commercial smart-home offerings are siloed, cloud-dependent, or not designed for accessibility. We present DORA, a voice-first, edge-executed well-being platform that integrates four capabilities—Home Safety (HS), Health Monitoring (HM), Smart-TV Interaction (TV) and Tele-Assistance (TA)—within a privacy-preserving microservice architecture.

We contribute three results. (i) We elicit requirements from elder, caregiver and remote-operator personas and derive domestic specifications that prioritise edge responsiveness, accessibility, data minimization and auditability. (ii) We design a microservice architecture with clear, domain-oriented service boundaries (notification, HM, HS, TV, TA), versioned MQTT event schemas and a defence-in-depth model that keeps biometric data on premises by default and, under explicit authorization with purpose and bounded time, supports standards-based exports (e.g., HL7 FHIR). (iii) We implement a working prototype that realizes all four functional modules: HS (environmental hazard detection, fall detection via knee-angle analysis, and alert escalation), HM (vital signs ingestion, trend analysis, daily summaries, and FHIR export), TV (voice-first interaction with on-device ASR, high-contrast overlays, and TTS synchronization), and TA (operator-initiated WebRTC video calls with consent gating and session logging). The implementation uses Home Assistant for device orchestration, MQTT for decoupled events, a local time-series store for vitals, and an Android TV application coupled with a WebRTC tele-assistance backend.

Architecturally, DORA separates sensing from decision and presentation. Device and wearable signals are normalised into stable JSON payloads on versioned topics; domain services subscribe, apply explainable analytics (rules, thresholds, rolling windows) and emit actions—text-to-speech prompts, TV overlays, caregiver notifications. This decoupling enables independent evolution, back-pressure handling and graceful restart. Privacy is addressed through local-first storage, consent-bound and purpose-limited exports, explicit data-quality indicators (e.g., signal confidence, motion flags) and immutable audit logs. Accessibility centres on a TV-first interaction model with short, confirmable prompts.

We evaluate the implemented system under controlled and realistic conditions. All four modules operate as designed: HS processes environmental events and fall detections, triggering sub-second median alerting and recovering cleanly after broker/service restarts. HM sustains 1 Hz ingestion, serves on-demand health

summaries and performs consent-scoped FHIR exports with only hundreds of milliseconds of local overhead. TA establishes operator-to-TV sessions in 1–3 s. TV interaction supports voice commands and visual feedback. Usability studies on three representative TV tasks—request health summary, acknowledge alert, join call—show low step counts and short completion times; accessibility reviews report readable overlays and clear recovery prompts.

These results suggest that decoupled events and idempotent consumers with retained state enable robust recovery without duplicate prompts. Interpretable analytics suffice for actionable home behavior, and TV-centric, short, confirmable prompts reduce cognitive load. Privacy and interoperability can coexist: local time-series storage, consent-scoped exports, and audit trails support integration with caregivers and clinical systems.

DORA shows that a modular, voice-first, local-first platform can improve perceived safety and connectedness for older adults without compromising data protection, while offering a reproducible path toward interoperable, explainable and trustworthy smart-home care.

Acknowledgements

I would like to express my heartfelt gratitude to my family for their unwavering support throughout this journey. To my father Jianmin MU, my mother Qiuhong Zhang, your encouragement and love have meant the world to me.

I am deeply grateful to my fidanzata, Wenwen Xian, who has been a constant companion in my life, providing immense support and understanding.

A special thanks goes to my lifelong friend Zhiyu Zhang, who I am sure will continue to be part of my life. I would also like to extend my gratitude to all my friends, including those not mentioned here, whose unwavering support has been invaluable during this time.

I would like to express my appreciation to the tutors, Luca Ardito, who supported me in the creation of this project.

Yujie Mu

Contents

List of Tables

List of Figures

Chapter 1

Introduction

Aging in place is a widely expressed preference among older adults, yet home environments can become challenging due to mobility limitations, chronic conditions, and the risk of emergencies. Social isolation and loneliness are serious public health concerns for older adults, with studies showing that these conditions increase mortality risk by 26–32% [1]. Meanwhile, family caregivers and healthcare providers often lack timely, trustworthy information and convenient channels to support residents. Many commercial smart-home products address isolated functions (security, lighting, voice control) and rely primarily on cloud processing, which raises privacy concerns and introduces latency. Interfaces are frequently not inclusive for people with reduced vision, limited dexterity, or low digital confidence. These gaps motivate an integrated, privacy-preserving, voice-first platform that improves safety and well-being while remaining simple to use.

We present DORA, a home-assistance platform designed around older adults' needs. DORA integrates four capabilities: Home Safety (environmental hazards and fall-related monitoring), Health Monitoring (wearable vitals and trends), Smart-TV Interaction (large-screen feedback and cognitive/social activities), and Tele-Assistance (operator-initiated calls with escalation). The platform follows an edge-first, microservice architecture: safety-critical logic runs locally on a home server to guarantee responsiveness and data sovereignty. Voice interaction is the primary interface, and the TV serves as the main visual surface.

1.1 Problem Statement

The problem addressed in this work can be stated as follows:

> Design and validate a privacy-preserving, voice-first home assistance

platform. Detect risks and trends, and connect the resident to human support in a timely manner, while remaining deployable on modest hardware and usable by older adults.

Achieving this requires balancing four forces: (i) responsiveness for safety-critical functions, (ii) data protection and transparency, (iii) modularity and maintainability, and (iv) inclusive interaction for non-technical users.

1.2 Objectives and Research Questions

We pursue three objectives:

- **O1 – Requirements and use cases.** Elicit needs of older adults, caregivers and remote operators; derive functional and non-functional requirements.
- **O2 – Architecture and design.** Define service boundaries, event schemas, data flows and security model for an edge-first platform.
- **O3 – Prototype implementation.** Implement an end-to-end system integrating device orchestration, rules/trend analysis, TV application and tele-assistance.

We examine three research questions: RQ1—Is a voice-first interaction sufficient to complete common tasks with minimal steps for older adults? RQ2—Does the model adequately ensure older adults' safety? RQ3—Does the system enable administrators (caregivers) to effectively observe and support the resident's status?

1.3 Contributions

Our contributions are fourfold:

1. We present a requirements set and use-case model for safety, health, TV interaction and tele-assistance, with access-rights and acceptance criteria aligned with domestic settings.
2. We design an architectural blueprint for DORA: domain-oriented microservices, MQTT-based events, time-series storage for vitals, and a defence-in-depth security approach that keeps biometric data on premises and exports FHIR-formatted subsets only on demand.[2, 3]
3. We implement a working prototype that realizes all four functional modules: HS (environmental hazard detection, fall detection, and alert escalation), HM (vital

signs ingestion, trend analysis, daily summaries, and FHIR export), TV (voice-first interaction with on-device ASR and visual overlays), and TA (operator-initiated WebRTC video calls with consent gating). The implementation integrates Home Assistant for device orchestration, a rules/trend engine, InfluxDB, an Android TV client, and a tele-assist service that establishes WebRTC sessions from an operator to the resident's TV with escalation policies.[4]

4. We evaluate the implemented system to verify responsiveness, availability, and usability, demonstrating that the prototype meets design objectives on modest hardware.

1.4 System Overview

DORA is organized into five logical layers that map to independently deployable services:

Device and Integration. Various smart home devices — such as environmental sensors, smart plugs, lighting systems, and connected appliances — are integrated through Home Assistant for unified control and monitoring.[5]

Service Layer. Home-safety rules evaluate environmental signals; the health-monitoring service ingests wearables' vitals, performs trend analysis and classifies states (normal/warning/critical); the notification service delivers voice and visual feedback; the tele-assist service manages operator-initiated sessions and escalation; the TV service renders prompts, summaries and video calls.

Intelligence. Simple, explainable analytics: rule evaluation and trend thresholds with configurable windows; anomaly detection across motion and vitals.

Interface. Voice-first commands and confirmations; caregiver dashboard is optional.

Security and Privacy. Local processing by default, encrypted storage for vitals, audit logs for calls and alerts, export of FHIR JSON only when explicitly authorized.[6, 3]

1.5 Personas and Interaction Scenarios

We consider four personas: the Older Adult (primary user), Family Caregiver and Remote Operator. Typical scenarios include: (i) real-time health monitoring and display vital data on the TV interface. In case of abnormal readings, it immediately issues a voice alert to notify the user or caregiver; (ii) trend-based alerts with caregiver notification; (iii) operator-initiated TV call; (iv) Displaying a calendar, today's tasks, and appointments on television helps prevent seniors from forgetting important events. These scenarios drive acceptance criteria such as maximum alert latency, call setup time, and the number of utterances required to complete tasks.

1.6 Requirements Summary

The functional requirements are grouped by module: Home Safety (fall-related heuristics, environmental alerts), Health Monitoring (ingestion, storage, trend analysis, daily summary, reminders, FHIR export), Tele-assistance (operator-initiated call, auto-answer policy, emergency escalation, session logging, TTS auto-reply). Non-functional requirements include: local execution for safety-critical paths, end-to-end alert delivery within a few seconds, availability targets.

1.7 Security Model and Compliance in DORA

Security in DORA follows a local-first, policy-driven model with defence-in-depth. Policies specify purpose- and time-bounded access, explicit consent, least privilege, data minimization, encryption in transit and at rest, and auditable decisions. Enforcement is distributed: each service exposes policy enforcement points that gate collection, processing, retention, and export. Attribute sources include Home Assistant entities, device and session state, user identity and role, time-of-day, and scenario flags (e.g., emergency).

Module-level mapping.

- **Home Safety.** Rules evaluate environmental signals (motion, smoke, water, temperature) and enforce obligations (notify, speak, escalate). Ongoing conditions (e.g., resident presence, quiet hours) allow revocation when context changes. The safety service acts as a trusted enforcement point mediating device operations and notifications.
- **Health Monitoring.** Ingestion and summaries are governed by consent, retention windows, and export scope. Exports (e.g., FHIR JSON) are permitted only with explicit authorization that binds subject, time range, and purpose. Local storage is the default; de-identification is mandatory on export.[3]
- **Tele-assistance.** Operator-to-TV sessions are gated by identity, role, and scenario (normal vs emergency). Auto-answer is allowed only under emergency policy; otherwise explicit consent is required. Session state is continuously checked; revocation terminates the call when conditions no longer hold.

Edge-first enforcement. Policy evaluation is performed on premises to preserve privacy and responsiveness. DORA services expose enforcement interfaces; attribute managers bridge Home Assistant entities, wearables, and runtime context into policy information points. Policies and event schemas are versioned, with audit logs for decisions and operator interactions.

Compliance and evaluation metrics. We evaluate security and usability using metrics within DORA's event-driven architecture:

- **End-to-end alert latency** (device → MQTT → rule → TTS/TV), reported by distribution (median and tail). Targets derive from domestic acceptance thresholds and accessibility constraints.
- **Authorization decision time** on modest hardware; **revocation propagation time** from attribute change to effective stop.
- **Call setup time** for operator → TV sessions under normal and emergency policies; **availability** under service restarts.
- **Usability indicators** for elder-centric flows: steps and confirmations to complete frequent tasks (e.g., daily summary, acknowledge alert, place/receive a call). As Corbett et al. note, "VHAs are perceived as companions and can improve social connectedness and reduce loneliness" [1], which underscores the importance of designing intuitive, elder-friendly interaction flows.

Pointers to detailed theory and implementation are provided in Chapter 2 (privacy-by-design, policy-driven access control, and enabling technologies), Chapter 3 (architecture and system design), Chapter 6 (TV client implementation and deployment notes), and Chapter 4 (methods and results).

1.8 Research Scope and Assumptions

The prototype targets a single dwelling with one resident, one TV device and an on-premises server. Wearables provide heart rate, SpO_2 and temperature as representative metrics. For pipeline validation, simulations also include additional placeholders (e.g., glucose and blood pressure). Computer-vision fall detection is limited to simple heuristics; advanced CV is out of scope. Cloud use is optional and limited to backups and remote operator access; no third-party data sharing is assumed.

1.9 Ethical and Privacy Considerations

Privacy is central to the design. The system minimizes data collection, stores vitals locally by default, and supports explicit, revocable consent for any export. Audit logs record operator calls and alert decisions. Interfaces avoid stigmatizing language and provide clear opt-out choices. [6]

1.10 Evaluation Plan

The evaluation focuses on: (i) rule execution latency on the edge node; (ii) device-to-caregiver alert delivery; (iii) operator- to-TV call setup; (iv) task completion with minimal utterances. Measurements are repeated to account for variability; results are reported with averages and dispersion and interpreted against target thresholds derived from domestic use.

1.11 Thesis Structure

Chapter 2 surveys enabling technologies and related work; Chapter 3 details the architecture and service design; Chapter 6 describes the TV client implementation and deployment notes; Chapter 4 presents methods and results; Chapter 7 concludes and outlines future work.

Chapter 2

State of the Art and Background for DORA

2.1 Overview

We survey the foundations that enable DORA—a Home Assistant–based well-being platform for older adults—with emphasis on four modules: Home Safety, Health Monitoring, Smart-TV Interaction, and Tele-Assistance. We review edge-first smart-home orchestration, event-driven messaging, health data interoperability, real-time audiovisual communication, and elder-centric interaction design [5, 3, 1]. We further analyze related platforms in the ambient assisted living (AAL) and telehealth domains, highlight integration challenges at the intersection of IoT and health data governance, and synthesize lessons learned from prior deployments. We conclude by identifying research gaps that motivate DORA's local-first, explainable, and privacy-preserving approach.

Domestic well-being systems must reconcile three competing forces: responsiveness for safety-critical interaction, inclusivity for users with diverse abilities and technology confidence, and protection of intimate household data. Central to the design is the decision to execute sensing, reasoning, and interaction primarily on the edge (i.e., within the home), using the cloud for optional backup, analytics, or remote operator access. The literature indicates that edge execution improves latency predictability while also minimizing exposure of biometric and behavioral traces [5, 6]. At the same time, achieving interoperability with clinical systems necessitates standards-based data models and auditable export mechanisms under explicit consent [3]. These dual imperatives shape the technology stack presented below.

2.2 Edge-first Orchestration in Smart Homes

2.2.1 Home Assistant as Integration Layer

Home Assistant (HA) is a widely adopted open-source platform that integrates a broad spectrum of sensors, actuators, and hubs via a modular integration model. Architecturally, HA follows an event bus pattern: device state updates and service calls are posted to a central bus, where automations, scripts, and external clients can subscribe and react. HA's design emphasizes local execution, YAML- and UI-driven configuration, and a pluggable component model that supports both LAN protocols (e.g., Zigbee, Z-Wave, MQTT, Philips Hue, LIFX) and cloud-backed APIs when necessary [5, 7].

From a systems viewpoint, HA plays three roles in DORA. First, it provides device discovery and state normalization, mapping vendor-specific representations (e.g., motion vs occupancy sensors, proprietary battery states) into a consistent entity model with standardized attributes (`state`, `last_changed`, `attributes`). Second, it acts as a broker between low-power domains (Zigbee/Z-Wave) and IP-based services, handling protocol translation and reliability semantics (e.g., retained states, reconnects). Third, it exposes a stable API surface (WebSocket and REST) and a persistent event stream that domain services can use to implement higher-level reasoning without re-implementing device adapters. This architectural separation allows safety and health services to evolve independently from the device layer while leveraging HA's mature integration ecosystem.

Several studies have analyzed HA's suitability for AAL environments, noting its extensibility, active community, and increasing support for formal configuration validation and sandboxing of custom components [5, 7]. However, challenges persist: (i) integration heterogeneity leads to inconsistent attribute naming; (ii) long-running automations can block the event loop unless carefully designed; and (iii) third-party custom components may lag behind platform updates. DORA mitigates these risks by offloading safety-critical logic to dedicated services, adopting versioned schemas at the event boundary, and validating device payloads against JSON schemas before ingestion.

Related Platforms. OpenHAB, Apple HomeKit, Google Home, and Amazon Alexa provide alternative orchestration approaches. OpenHAB similarly offers open-source, local-first control with an OSGi plugin system, but with a steeper learning curve for custom bindings [8]. HomeKit emphasizes privacy and local certification, but restricts device categories and developer extensibility. Google Home and Alexa excel in voice integration and commercial device coverage, yet depend heavily on cloud services and opaque data processing. Comparative analyses suggest that open, local-first platforms (HA/OpenHAB) offer better control over data and latency, at the cost of greater configuration effort and operator responsibility [8, 9]. DORA selects HA for its velocity of integration development, vibrant community, and event-bus model compatible with microservices.

Table 2.1: Comparison of representative approaches and DORA.

Approach (refs)	**Salient features**	**Advantages**	**Limitations vs DORA**
Voice-based home control [10, 11, 12, 1, 13]	Speech commands to control appliances	Natural interaction; low barrier for simple tasks	Cloud dependence common; limited governance; no TV-centric prompts or remote assist
IoT/WSN smart-home [8, 14, 9, 15, 16]	Sensor/actuator integration; rules	Saves time/energy; broad device coverage	Heterogeneous semantics; weak audit/consent; limited clinical interoperability
WebRTC telepresence [4, 17, 18]	Browser/device-native A/V	Low-latency calls; cross-platform	Often standalone; lacks integration with safety/health workflows
FHIR interoperability [3]	Standards-based health data exchange	Reusable with clinical systems	Rarely edge-first; consent/data-quality often under-specified
DORA (this work)	Edge-first HS/HM/TV/TA + MQTT; consent/audit; TV-centric UI	Sub-second prompts; 1–3 s calls; privacy by default; FHIR export on demand	Requires local node; operator training for escalation

2.2.2 Event-driven Messaging with MQTT

MQTT is a lightweight publish/subscribe protocol designed for unreliable networks and constrained devices. Its core features include topic-based routing, three quality-of-service (QoS) levels, retained messages for last-known values, and persistent sessions. In DORA, MQTT decouples producers (sensors, HA bridges) from consumers (safety, health, TV, and tele-assist services). Topics are organized by domain (e.g., `hazard/gas`, `vitals/ingest`, `fall/detected`) to support clear separation of concerns and facilitate independent service evolution.

From a reliability perspective, retained messages ensure that late subscribers (e.g., a restarted TV client) receive the latest state upon subscription, while QoS 1/2 guarantee at-least-once or exactly-once delivery where necessary. DORA adopts an idempotent consumer design: messages carry event identifiers and logical timestamps to enable deduplication and out-of-order handling. Security is provided by TLS transport, per-client credentials, and topic-level ACLs that separate read/write privileges across services [6, 2].

The literature on MQTT in AAL and telehealth underscores trade-offs among QoS, latency, and battery usage, especially for wearable gateways. Studies have explored adaptive QoS and batching strategies to cope with transient congestion and network variability [8]. DORA applies these lessons by reserving QoS 1 for safety-critical topics (e.g., alerts, call control) while using QoS 0 for high-rate telemetry with application-level smoothing. Back-pressure is handled through bounded processing queues and non-blocking consumers to avoid event loop stalls.

2.3 Health Data Interoperability

2.3.1 FHIR Resources and Domestic Time-series

HL7 FHIR (Fast Healthcare Interoperability Resources) defines modular resources (e.g., `Observation`, `Patient`, `Device`, `Encounter`) and exchange formats (JSON, XML) to support interoperable health data exchange. FHIR's profiling mechanism allows implementers to constrain and extend base resources for specific use cases (e.g., vital signs profiles that standardize codes and units). In DORA, locally stored vital signs (heart rate, SpO_2, temperature) are periodically summarized; when export is authorized, a FHIR Bundle is constructed containing the relevant `Observation` instances, device metadata, and provenance annotations (time window, purpose, consent reference) [3].

Domestic time-series databases (e.g., InfluxDB, TimescaleDB) efficiently handle append-only measurements with retention policies and downsampling. For DORA, this model supports rolling window analytics (median, variance, threshold crossings) without exposing raw streams externally. Research in edge analytics shows that pre-aggregation near the source reduces bandwidth and privacy risks while maintaining clinical utility for trend detection [6]. A practical design choice is to store raw samples for a short horizon (e.g., days) and keep downsampled aggregates over longer periods (e.g., weeks), consistent with minimization principles.

Mapping and Semantics. A common challenge is the mapping from consumer grade wearable metrics to clinical concepts. Device-specific smoothing, sampling biases during movement, and skin tone effects can skew readings [3]. The literature recommends surfacing data quality indicators (signal confidence, motion flags), time synchronization across sensors, and explicit unit coding (UCUM). DORA includes these indicators in both local summaries and FHIR exports to signal reliability to downstream systems.

2.3.2 Wearable Ingestion and Trend Analysis

Wearables expose data via vendor SDKs, proprietary cloud APIs, or local BLE/-GATT profiles. Cloud APIs simplify ingestion but raise privacy concerns and increase latency, while BLE requires local protocol handling but avoids remote dependencies. Prior work highlights BLE fragmentation (characteristics, pairing flows) and the value of gateway software that abstracts device differences [8]. DORA supports both approaches with a preference for local ingestion when feasible, buffering data during connectivity gaps and annotating samples with acquisition context (rest, walking) when available.

Trend analysis in domestic settings favors simple, explainable methods: moving averages, exponentially weighted smoothing, and threshold-based flags calibrated per user. While machine learning can detect subtle anomalies, research cautions against opaque models in safety-critical and high-variance environments, recommending hybrid pipelines that combine rules with interpretable features [15]. DORA adopts such a hybrid approach and reserves complex models for optional, offline analysis.

2.4 Real-time Communication for Tele-Assistance

2.4.1 WebRTC Fundamentals

WebRTC provides a browser- and device-native stack for secure, low-latency audio/video with SRTP media encryption, DTLS keying, and ICE to establish connectivity between endpoints. Architectures range from peer-to-peer to SFU (Selective Forwarding Unit) and MCU (Multipoint Control Unit). SFUs route encoded media streams without transcoding, striking a balance between scalability and resource use on edge devices; MCUs decode/mix/re-encode, offering layout control at higher cost [4, 17].

For a TV endpoint in a domestic environment, SFU-backed topologies reduce device CPU load and allow the operator console to connect reliably across NATs. Literature on telepresence for older adults emphasizes fast call setup (sub-3 seconds user-perceived), visual identity cues to establish trust, and clear escalation paths when a call is not answered. DORA incorporates identity overlays and operator role badges on the TV, with auto-answer permitted only under emergency scenarios consistent with household policy.

2.4.2 Operator Workflows and Escalation

Tele-assistance solutions in AAL commonly integrate call routing, escalation policies (retry, route to caregiver), and context overlays (e.g., last known activity, vital trends). Research highlights the importance of minimal operator steps, scripted prompts for clarity, and comprehensive session logging for duty-of-care audits. DORA's workflow includes pre-call checks (device availability, do-not-disturb), consent gating in non-emergencies, and post-call notes.

Integration challenges include codec compatibility (H.264/AV1/VP9) with TV hardware decoders and echo cancellation in living rooms. Lessons learned from deployments point to fallback codecs and conservative initial bitrates to avoid stalls during call establishment [4, 18].

2.5 Elder-centric Interaction and Usability

2.5.1 Voice-first and Large-screen Design

The HCI literature for older adults recommends high-contrast palettes (e.g., WCAG AA/AAA), large type suitable for living-room viewing, and sparse layouts with generous hit targets for remotes. Voice prompts should be short, confirmable, and tolerant of hesitations. Error messages should propose exactly one clear next step to reduce cognitive load [1, 10, 11, 13]. As Corbett et al. observe, "older adults desired more education and training on VHA use" [1], highlighting the need for intuitive interfaces. DORA's TV overlays use legible components and animated focus cues. Voice is the primary modality.

The home context poses additional constraints: lighting variability, background noise (TV, appliances), and multi-user households with differing preferences. Personalization (font size, contrast themes), quiet-hours policies, and confirmation levels improve acceptance. The literature suggests that perceived control and transparency are as critical as raw task performance in long-term adoption.

2.5.2 Usability Evaluation Methods

Evaluation in AAL emphasizes ecologically valid, low-burden protocols. Common quantitative measures include task steps, time to completion, and error counts; qualitative methods include think-aloud, brief semi-structured interviews, and diary notes. Standard instruments such as SUS and NASA-TLX can be adapted with simplified wording and shorter forms for older adults. DORA adopts repeated, short sessions aligned with daily routines (e.g., requesting a summary, acknowledging an alert), enabling longitudinal observation without fatigue.

2.6 Home Safety: Sensing and Local Reasoning

2.6.1 Environmental Signals and Alerts

Home safety sensing fuses motion, door/window contact, smoke/CO, water leaks, and temperature to form contextual hypotheses (e.g., "resident inactive unusually long", "water leak near kitchen"). The state of the art ranges from threshold heuristics to statistical models and lightweight activity recognition. Given the need for explainability and low false positives in domestic contexts, rule-based reasoning with hysteresis and cooldowns remains competitive, especially when combined with multi-sensor corroboration.

Alerting channels must balance intrusiveness and effectiveness. Text-to-speech (TTS) and TV overlays are appropriate for in-home prompts; caregiver notifications (push, SMS) should be rate-limited and deduplicated. Prior work reports that stacked alarms and repeated alerts contribute to alarm fatigue and disengagement; adopting suppression windows and explicit acknowledgment flows mitigates this risk.

2.6.2 Reliability and Graceful Degradation

Edge-first safety systems must remain functional during partial service failures. Design patterns include retained critical states in the broker, watchdog timers for device liveness, and circuit breakers that shed non-critical processing under load. Research on resilient home automation highlights the benefits of idempotent event handling and monotonic state machines for recovery after restarts. DORA's services are designed to restart quickly, rebuild subscriptions, and reconcile last-known states without human intervention.

2.7 Data Privacy and Protection in Domestic Settings

2.7.1 Minimization, Purpose Limitation, and Consent

Regulatory frameworks (e.g., GDPR) and ethical guidance emphasize minimization, purpose limitation, and consent for processing sensitive data within homes. The AAL literature warns against collecting rich raw streams (audio/video) by default; instead, it recommends on-device processing and ephemeral features when practicable. DORA adheres to a local-first strategy: vitals and logs are stored on premises; exports are explicitly initiated or pre-authorized with narrow scopes (time range, recipients, attributes) and recorded in audit trails.

2.7.2 Auditability and Transparency

Audit logs should record the rationale behind decisions (e.g., thresholds crossed, conditions met), the attributes evaluated, and the obligations executed (e.g., de-identification). Research shows that plain-language explanations improve user trust and support caregiver accountability. DORA attaches concise explanations to prompts and provides a dashboard view for caregivers or administrators to review historical decisions.

2.8 Position within the Literature

Three bodies of work intersect in DORA: (i) smart-home orchestration and automation; (ii) telehealth and wearable-driven monitoring; and (iii) human-centered interaction for older adults. Smart-home platforms such as Home Assistant and OpenHAB demonstrate the feasibility of local, event-driven control. Telehealth research validates the clinical value of trend tracking and remote consultations, but also documents data governance challenges and patient burden from complex devices. HCI work provides actionable patterns for accessible, low-friction interaction. Studies indicate that "VHAs reduced loneliness in many participants" and that "all researchers noted that VHAs provided a source of companionship" [1], supporting the value of voice-first interaction for older adults.

Comparable systems include AAL frameworks integrating fall detection, medication reminders, and caregiver portals; commercial ecosystems (e.g., Alexa Together) that offer check-ins and alerts; and research prototypes combining wearables with in-home displays. Many solutions rely on cloud processing, impeding operation during outages and raising privacy concerns. Others lack standardized health data exports or provide limited transparency into decision logic. DORA's contribution is the integration of these strands into a coherent, edge-first architecture with explainable reasoning, consent-bound FHIR exports, and TV-centric interaction designed for living rooms.

2.9 Illustrative Scenarios

Scenario 1: Night-time water leak. A water sensor near the kitchen reports a leak; temperature and motion indicate the resident is likely asleep. The safety service applies a quiet-hours policy: a soft chime on the TV and a low-volume TTS prompt request confirmation. If unacknowledged, the system escalates to a brighter overlay and notifies a caregiver after a suppression window. The broker retains the alarm state so the TV client, upon reconnection, re-renders the prompt rather than assuming normalcy.

Scenario 2: Morning health summary. On request, the health service summarizes heart rate, SpO_2, and sleep proxies from the previous day. The TV overlay presents a two-line summary with trend arrows; the voice prompt explains any warnings in plain language and proposes one next step (e.g., "Would you like me to remind you to hydrate?"). If the resident later authorizes sharing, a FHIR bundle covering the last seven days is exported to a clinician portal, annotated with purpose and consent.

Scenario 3: Operator-initiated welfare check. An operator initiates a call after an anomaly notification. The TV displays the operator's identity and role; if an emergency scenario is active, auto-answer engages, otherwise the resident confirms acceptance. The SFU negotiates media with conservative initial bitrates, then ramps up. Session logs record decisions, timestamps, and any escalation (caregiver ping) if unanswered.

2.10 Integration Challenges and Lessons Learned

Device heterogeneity. Wearables, sensors, and hubs expose inconsistent semantics and reliability. Lessons from deployments stress the need for versioned schemas at the service boundary, adapter isolation, and continuous validation of payloads to avoid propagation of malformed states.

Runtime variability. Event bursts and device resource limits (CPU/GPU) can degrade responsiveness. MQTT QoS tuning, bounded processing queues, and conservative initial bitrates reduce perceived instability during alerts and calls.

Explainability over opacity. Simple, interpretable rules and thresholds enable trust and troubleshooting. When machine learning is used, models should be auditable, bounded in scope, and paired with human-readable explanations.

Privacy by default. Local storage with explicit, scoped exports reduces risk and regulatory complexity. Users value clear records of what was shared, with whom, and why; therefore, audit views and consent history are first-class features rather than afterthoughts.

2.11 Research Gaps

Despite progress in AAL, three gaps persist. First, few systems demonstrate an end-to-end, local-first pipeline that spans device integration, event reasoning, elder-centric interaction, and standards-based health data export. Second, empirical characterizations of domestic end-to-end latency (sensor $\rightarrow$prompt), revocation timeliness, and call setup times are sparse in the literature. Third, transparency mechanisms remain limited: users and caregivers seldom receive concise, actionable explanations or auditability across alerts, calls, and exports. DORA addresses these gaps through an integrated architecture, a measurement-oriented evaluation plan, and explicit design commitments to explainability and privacy.

Chapter 3

Methodology and System Design for DORA

3.1 Scope and Objectives

DORA is conceived as a Home Assistant–based well-being platform for older adults that unifies four functional modules—Home Safety (HS), Health Monitoring (HM), Smart-TV Interaction (TV), and Tele-Assistance (TA)—under a privacy-preserving, edge-first architecture. This chapter specifies the architectural and methodological foundations required to implement a coherent, auditable, and inclusive system that can be deployed on modest home hardware while preserving data sovereignty, and defines the experimental design, metrics, and validation steps that bridge the transition from the technology landscape surveyed in Chapter 2 to the empirical findings reported in Chapter 4. In doing so, we articulate the rationale for design choices (e.g., selection of Home Assistant for device integration; the adoption of MQTT for event decoupling; the preference for edge processing and explainable analytics), detail the microservice interactions and data handling practices, and provide the procedures and instrumentation used to assess responsiveness, availability, usability, and privacy.

The high-level objectives are: (i) define a modular, event-driven, local-first system that minimizes exposure of sensitive data; (ii) establish measurable procedures spanning HS/HM/TV/TA to quantify end-to-end alerting, interactive latency, and recovery; (iii) operationalize consent-aware, standards-based exports (HL7 FHIR) alongside transparent audit trails; and (iv) ensure that interaction patterns—particularly on TV—support accessibility through short, confirmable prompts.

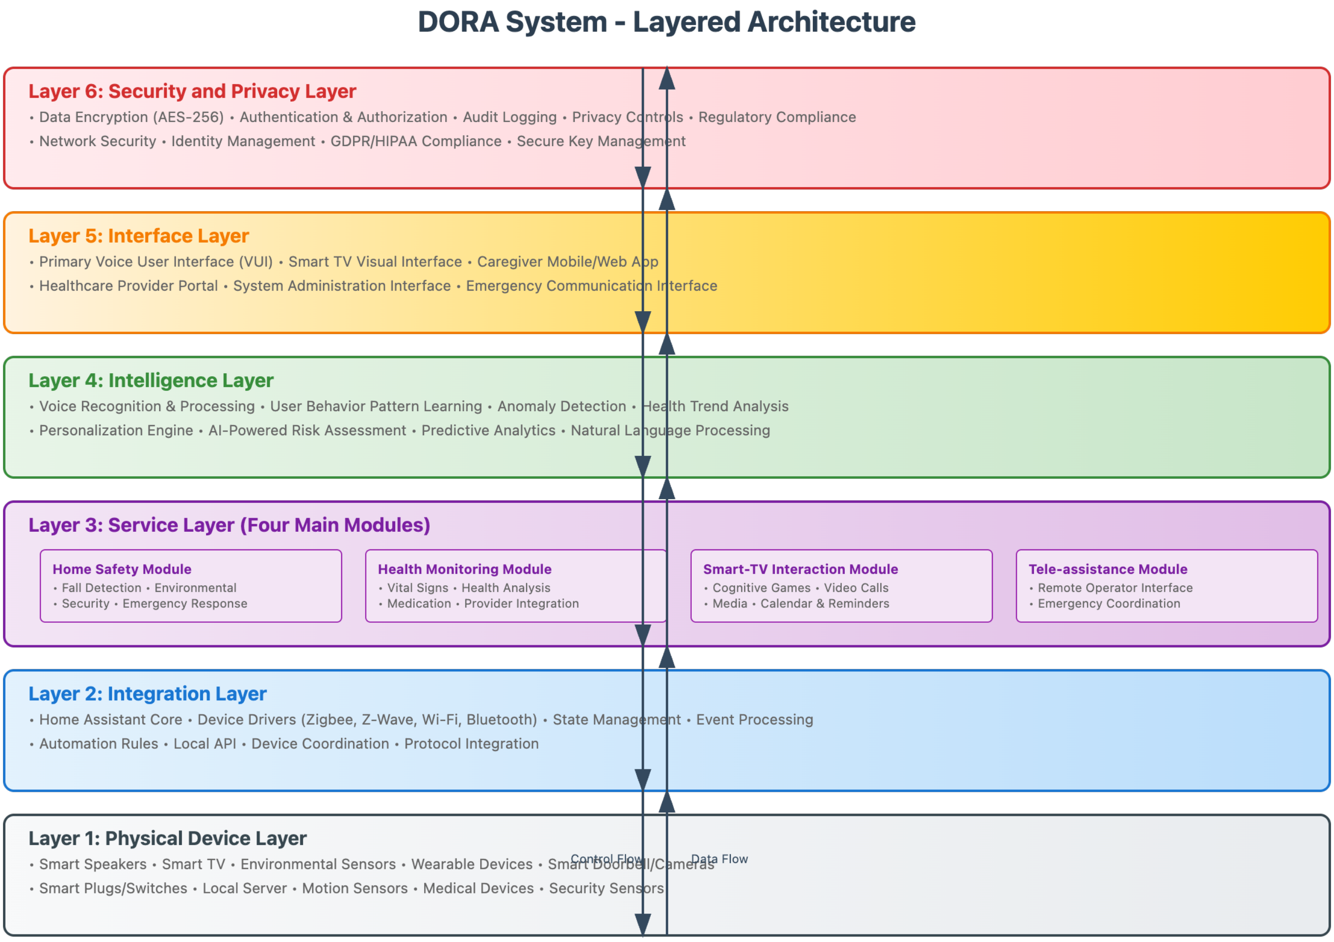

Figure 3.1: Layered view of DORA showing device, integration, service, intelligence, interface, and security/privacy layers.

3.2 System Architecture and Data Flow

3.2.1 Rationale for Edge-first, Microservice Composition

Domestic environments confront unique constraints: runtime variability, resource limits, diverse consumer devices and protocols, and stringent privacy expectations. Placing safety-critical and user-facing logic on an edge node provides predictable latency and continuity. A microservice composition separates concerns—device orchestration, safety reasoning, health trend summarization, TV presentation, and telepresence—so that each domain can evolve independently while limiting the blast radius of failures. This separation is particularly advantageous for accessibility and privacy: HS and HM can adopt explainable, low-false-positive analytics tuned for domestic settings; TV can focus on legibility, focus management, and error recovery; TA can optimize media setup and transparency (identity and role badges), gated by consent policies.

3.2.2 Home Assistant as Device Integration Layer

Home Assistant (HA) was selected as the integration layer for devices and hubs because it offers: (i) wide ecosystem coverage for sensors and actuators across LAN protocols (Zigbee, Z-Wave) and selected cloud APIs; (ii) a consistent entity model with state and attribute semantics; (iii) local execution by default; and (iv) mature interfaces (WebSocket, REST) that expose an event bus without binding domain logic to the automation engine. HA normalizes heterogeneous devices into a common representation (e.g., motion vs occupancy sensors, battery levels, contact states) and emits state changes that DORA services can subscribe to via translators that publish to MQTT. This separation allows DORA to treat HA as a robust adapter layer, while concentrating domain logic in small services with auditable behavior.

3.2.3 Event Decoupling with MQTT

MQTT provides a lightweight publish/subscribe substrate with retained messages and quality-of-service (QoS) levels appropriate for constrained devices and unreliable links. DORA maps HA state transitions, wearable readings, and domain decisions onto topic hierarchies organized by domain, such as `hazard/gas`, `hazard/water`, `vitals/ingest`, and `fall/detected`. Payloads are JSON documents that carry unique identifiers, timestamps, and correlation fields to facilitate independent service upgrades. Producers publish idempotent messages; consumers are idempotent and maintain bounded queues to absorb bursts without stalling event loops. Retained messages are employed judiciously for critical prompts so that reconnecting clients (e.g., the TV app) can re-render the latest alert state.

Data sources and simulation in Home Assistant In this prototype, physical devices were not required during development and tests. Instead, we used Home Assistant to host *simulated* entities (helpers/template sensors) whose states were driven by small scripts so that DORA would observe realistic value changes through the same APIs and event bus used for real devices. Concretely:

- Environment signals for HS (e.g., gas/CO/CO_2, temperature, water leak, smoke surrogate) were mapped to HA entities and exercised by scripted state transitions at controlled rates; HA emitted state $\rightarrow$ event updates that the HS service consumed via MQTT.
- Vitals for HM (heart rate, SpO_2, glucose, temperature, sleep proxies, steps, blood pressure) were exposed as HA sensor entities and updated by a Python simulator; DORA ingested them as time-series just like physical wearables.

This approach keeps the integration faithful while enabling deterministic, repeatable scenarios for evaluation.

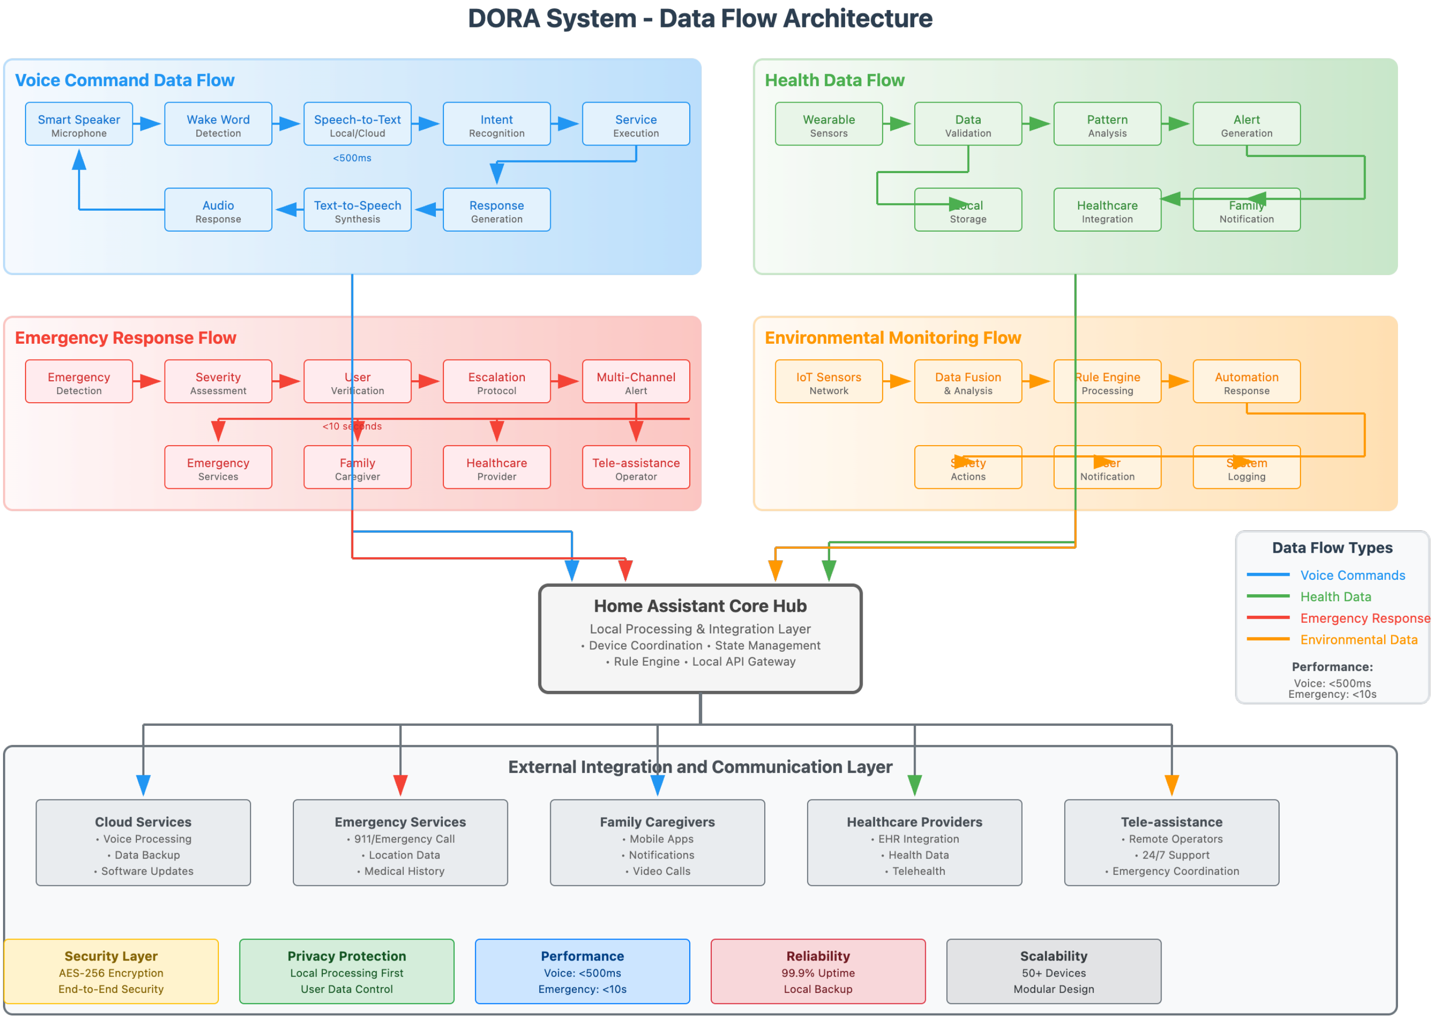

Figure 3.2: End-to-end data-flow across voice, emergency response, health, and environmental pipelines, integrated via the Home Assistant core hub.

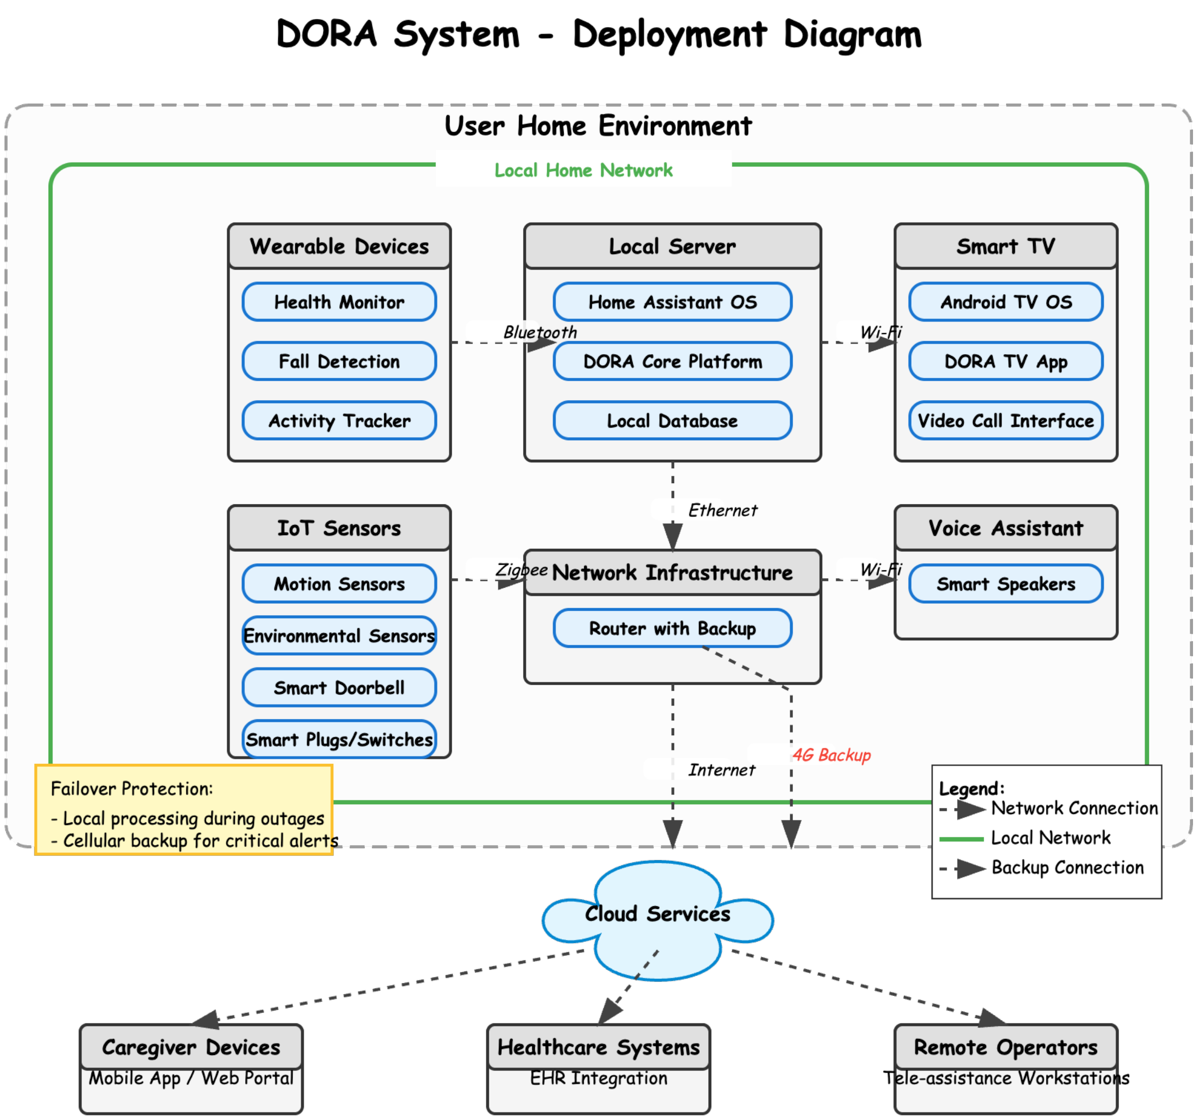

Figure 3.3: Deployment view of the local home environment: device layer, local server (HA + DORA services), Smart TV, and backup/remote links.

3.3 Module Specifications

3.3.1 Home Safety (HS)

HS is tasked with fusing heterogeneous environmental signals to produce actionable, low-false-positive alerts. The service ingests events published from HA translators, enriches them with local context (presence proxies, quiet-hours, recent activity windows), and applies interpretable rules with hysteresis to avoid oscillations. Alerts are emitted with severity tags and presentation cues (auditory salience during daytime; reduced intensity at night). A suppression window prevents alarm fatigue while guaranteeing escalation if unacknowledged. The escalation pipeline publishes a caregiver notification event and records the audit trail with deduplication keys to avoid multiple deliveries for the same condition. HS is latency-sensitive and optimized to compute in hundreds of milliseconds on modest hardware. Recovery relies on retained messages for the current alert state and idempotent handlers keyed by alert identifiers.

Fall detection via knee-angle dynamics

DORA implements a lightweight, privacy-preserving fall detector that operates locally. Two complementary paths are supported:

- **Angle-series rule.** A sliding window of knee angles is computed from live pose landmarks; a fall is flagged when all samples in the window are below a flexion threshold ($\approx$105°) *and* the swing (max–min) within the window exceeds a minimum ($\approx$20°). The window length defaults to 8 samples.

- **Single-frame pose.** For static images, MediaPipe Pose is loaded on demand; the hip–knee–ankle angle is estimated and a fall is flagged when it is below $\approx$110°.

On a positive decision, HS publishes an event to the fixed MQTT topic `fall/detected`, writes an auditable record (method, thresholds, timestamps), and, when enabled, notifies caregivers. The TV client subscribes to the topic and renders a high-contrast overlay with a short, confirmable prompt. For integration and testing, HS exposes HTTP endpoints `/falls/angles` (angle series) and `/falls/frame` (base64 image) and a CLI that streams camera frames, computes knee angles, and publishes MQTT events. All processing remains on the edge node to minimize latency and protect sensitive video content.

3.3.2 Health Monitoring (HM)

HM ingests vitals from wearables or gateways at 1 Hz using local protocols or API bridges. The ingestion layer performs basic validation (units, plausible ranges), attaches data-quality indicators, and persists to the time-series store. Trend computation uses rolling windows (e.g., median, variance, threshold crossings) to derive states (normal, warning, critical). On request, HM produces a daily summary with succinct prose and trend arrows for TV presentation. Export workflows construct consent-scoped HL7 FHIR bundles (`Observation` resources with `Patient`, `Device`, and `Provenance`), attach purpose and time windows, and apply redaction obligations (e.g., device serials) before transmission. The FHIR construction overhead is small relative to link transfer, keeping interactive latency unaffected.

3.3.3 Smart-TV Interaction (TV)

The TV client is the primary visual surface. It renders high-contrast overlays and short, confirmable prompts, synchronized with TTS. On-device speech recognition uses Vosk (`voskmodel.small.en-us-0.15`). Prompts always offer one clear next step, reducing cognitive load; secondary information is available behind a "more" action to avoid clutter. The TV client implements idempotent rendering keyed by prompt identifiers and replays retained states upon reconnect, ensuring correctness under transient disruptions.

3.3.4 Tele-Assistance (TA)

TA establishes operator-initiated WebRTC sessions to the TV with selective forwarding; direct vs relayed media paths are considered. The TV overlay presents the operator's identity and role; in non-emergencies, resident acceptance is required; in emergencies, auto-answer engages under policy. Media begins at conservative initial bitrates to avoid stalls, then adapts rapidly. The TA service logs session metadata (start time, acceptance decision, reconnect events) and minimal annotations for accountability, without recording sensitive content. Consent policies govern acceptance and escalation behavior, and are surfaced on TV in plain language to support transparency.

Caregiver dashboard and operator console (illustrative views)

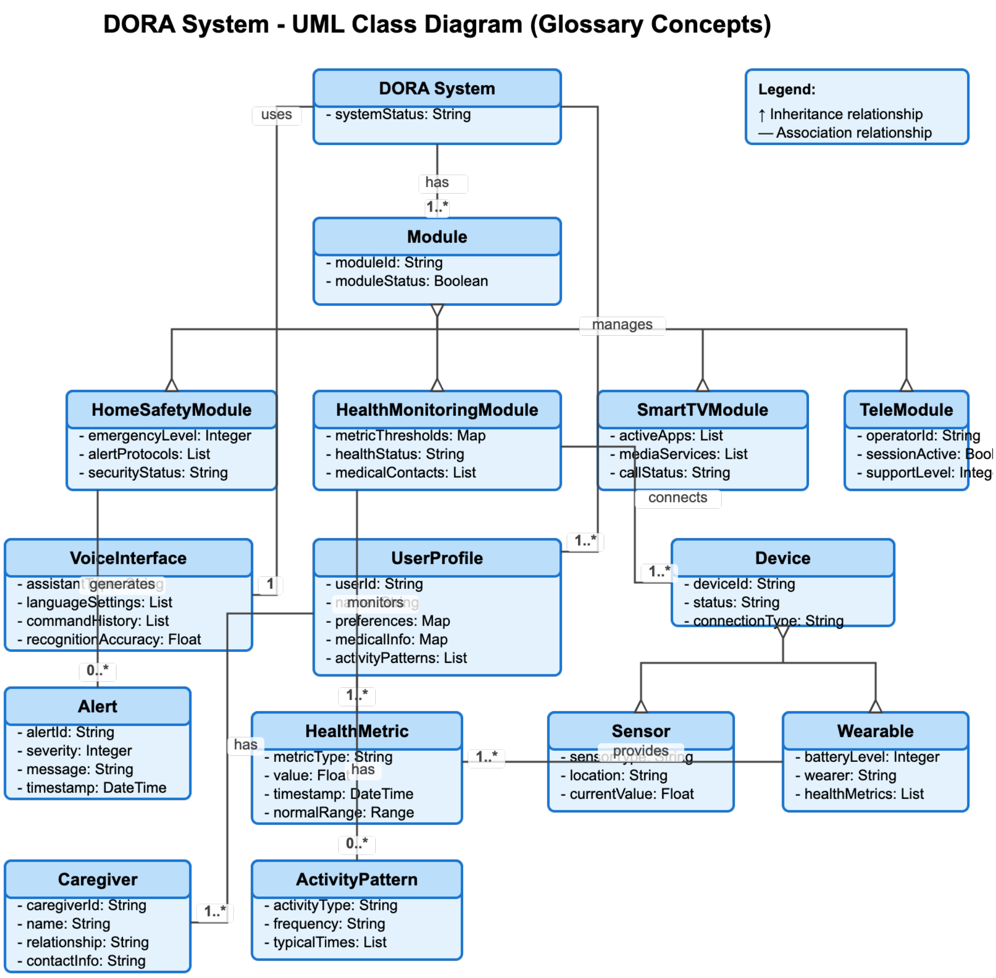

Figure 3.4: Domain model overview showing core entities (modules, user profile, devices, sensors, metrics) and their relationships.

Figure 3.5: Caregiver console: selectable list of elders for initiating a call.

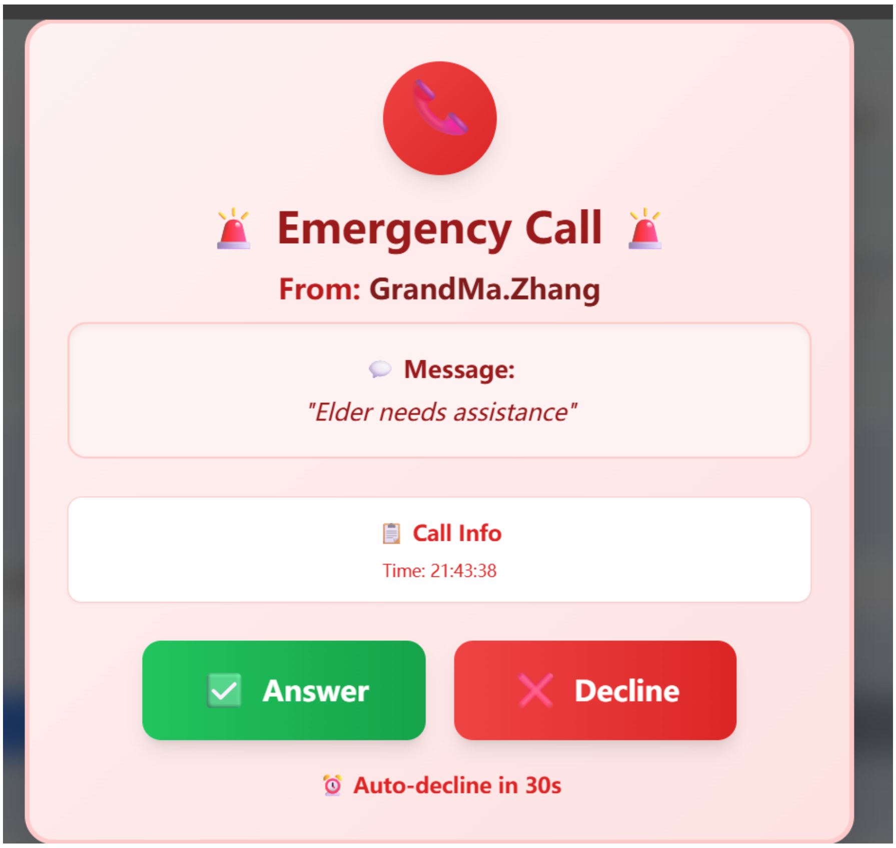

Figure 3.6: Caregiver console: emergency call flow for a resident requiring assistance.

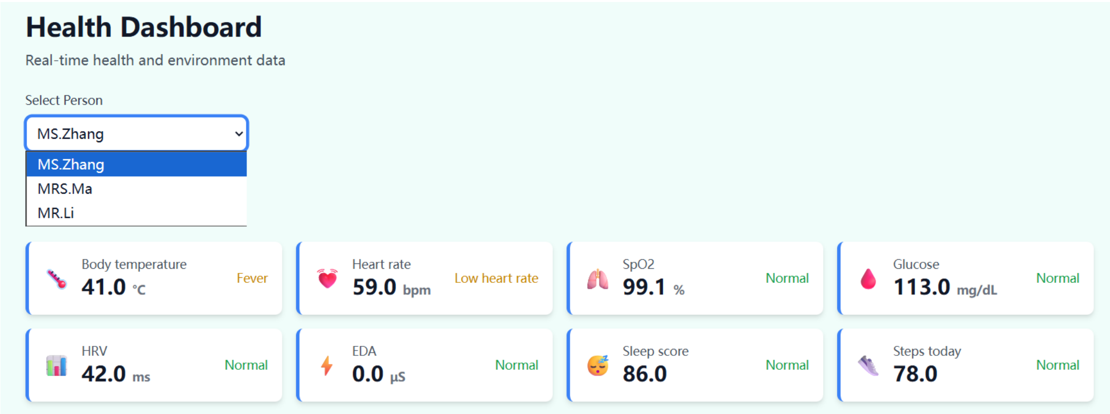

Figure 3.7: Caregiver console: per-resident health view demonstrating switching across multiple elders.

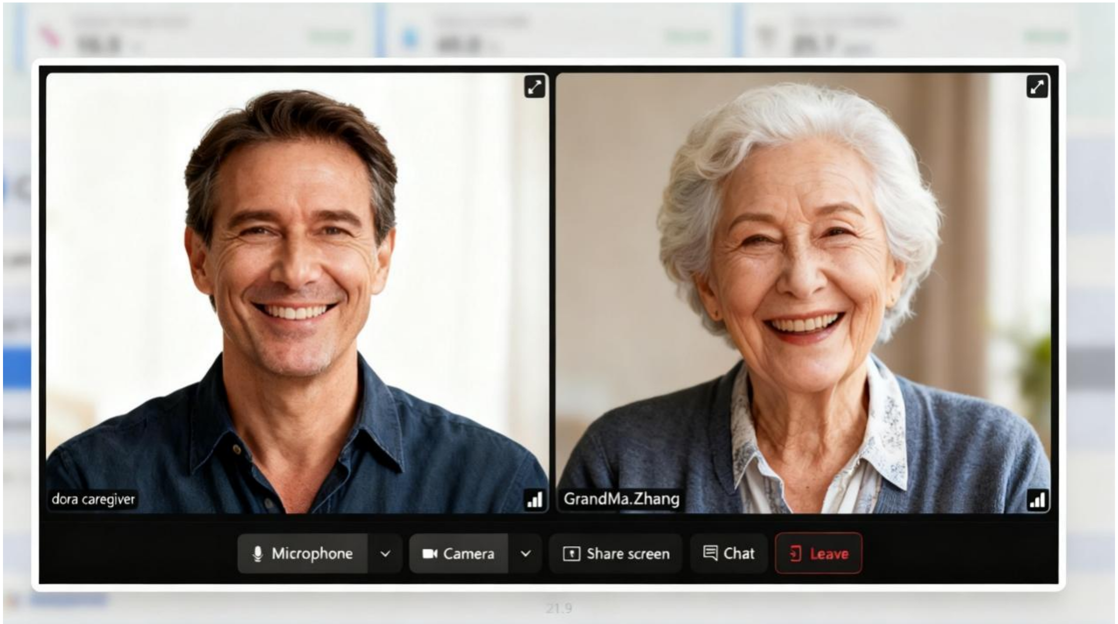

Figure 3.8: Caregiver console: administrator–to–TV video call interface (operator side).

3.4 Security, Privacy, and Compliance

3.4.1 Minimization, Consent, and Obligations

DORA adopts a local-first approach: sensing, analysis, and interaction are executed on the edge node. Personal and biometric data remain on premises by default. Exports occur only under explicit authorization and carry purpose and time-window boundaries. Obligations such as redaction are enforced at the point of export and recorded in the audit trail. Decisions log the attributes referenced and the obligations executed to create an accountable record suitable for after-action review.

3.4.2 Explainability and Transparency

Explainable analytics are preferred for domestic safety tasks. Rules and thresholds with hysteresis are readily interpretable, enabling residents and caregivers to understand why alerts appeared and how to respond. Prompts carry plain-language explanations, especially for tele-assistance governance (e.g., why auto-answer is enabled during emergencies, how consent can be changed). Surfacing data-quality indicators (e.g., confidence, motion flags) in overlays and FHIR exports aids cautious interpretation.

3.4.3 Resilience and Governance

Idempotent consumption, retained messages for critical states, bounded queues, and monotonic state machines together provide graceful restart and recovery. Versioned schemas and explicit deprecation windows help maintain backward compatibility across service updates. Governance policies are enforced locally and reflect household preferences (quiet hours, escalation thresholds), while remaining transparent and revisable.

3.5 Evaluation Metrics

3.5.1 Performance and Availability

Performance focuses on end-to-end alert latency (sensor →broker →rule →TTS/TV), authorization latency and revocation inconsistency time for guarded actions, and TA call setup and first-frame time. Availability includes recovery latency under broker and service restarts, reconnect behavior of the TV client, and continuity under temporary client pauses. Metrics are reported as medians with P95 and P99, including interquartile ranges.

3.5.2 Usability

Usability is quantified by the steps and completion time for three representative tasks: T1 (request a daily summary), T2 (acknowledge an alert), and T3 (join or accept an operator call). We track misrecognition and recovery rates for voice prompts and conduct expert accessibility reviews for contrast, font size, focus cues, and error recoverability.

3.5.3 Privacy

Privacy is assessed via the local-only storage ratio; the frequency, scope, and size of exports; and export conformance (presence of consent, purpose, time bounds, and redaction outcomes). Audit completeness is measured as the proportion of entries containing decision rationale, attributes referenced, and obligations executed.

3.6 Experimental Environment

Experiments run on an on-premises edge node that hosts HA, the MQTT broker, HS/HM/TV/TA services, a time-series store for vitals, and the TV client. Wearables or simulators generate representative vitals. The operator console is isolated in a separate administrative context. Scenarios are scripted for repeatability, each executed at least 30 times to obtain distributions. Where human tasks are involved (TV interaction), participants follow standardized instructions; time-to-completion is measured and error states recorded. The environment includes fault injection tools for broker restarts, service restarts, and controlled client pauses.

3.7 Procedures by Module

3.7.1 Home Safety

HS procedures include (i) generating environmental events as impulses and bursts with jitter; (ii) measuring alert latency and delivery success under background load and broker restarts; and (iii) validating suppression windows, quiet-hours presentation, and escalation. Receiver-side deduplication is verified by tracking caregiver notifications.

3.7.2 Health Monitoring

HM procedures include (i) streaming wearable metrics at 1 Hz; (ii) verifying ingestion continuity and buffer reconciliation across 60 s pauses; (iii) triggering on-demand daily summaries and measuring TTS/overlay synchronization; and (iv) executing consent-scoped FHIR exports with varying time windows and verifying redaction and metadata completeness (purpose, consent, data-quality indicators).

3.7.3 Smart-TV Interaction

TV procedures include (i) evaluating T1/T2/T3 tasks while recording step count, completion time, misrecognitions, and recovery prompts; (ii) conducting expert accessibility reviews; and (iii) stress testing overlay rendering (cold/warm) and focus handling under rapid prompt succession to detect regressions in client responsiveness.

3.7.4 Tele-Assistance

TA procedures include (i) initiating operator →TV calls in normal and emergency modes and measuring setup/first-frame times (direct vs relayed media path); (ii) emulating temporary link pauses and capturing reconnection behavior; and (iii) verifying identity/role badges, consent gating vs auto-answer boundaries, and audit trail completeness.

3.8 Methodology: Experimental Design and Validation

3.8.1 Design Principles

The methodology privileges realistic domestic constraints while maintaining repeatability. Experiments are designed to isolate module behavior (e.g., HS latency under controlled bursts) and to assess cross-module interactions (e.g., HM summary responsiveness while TA attempts connection). Workloads are staged to avoid unrealistic saturation while exercising tails.

3.8.2 Instrumentation and Data Collection

Structured logs and client timestamps form the basis of measurement. Each prompt carries a unique identifier and timestamps for publication, reception, render start, and user action; TA log entries record signaling completion, first frame, and reconnect events. A lightweight collector aggregates logs for analysis, computing medians and tail percentiles with interquartile ranges. Where appropriate, qualitative notes (e.g., user perception of prompt clarity) complement quantitative measures to inform design iterations.

3.8.3 Validation and Threats to Validity

Validation steps include sanity checks of topic schemas against versioned JSON; replay of recorded event traces to verify idempotency and recovery invariants; and contract tests that assert policy-relevant behavior (e.g., suppression windows, escalation thresholds) under upgrades. Internal validity is constrained by controlled stimuli; we mitigate with multiple repetitions and background load control. External validity is limited by single-dwelling evaluation; future work should stratify by dwelling types and RF environments. Construct validity is oriented to responsiveness, reliability, and usability; broader well-being outcomes require longitudinal mixed-methods studies. As noted in the literature, "the state of the science on older adults' VHA use and its influence on social isolation and loneliness is in its infancy, with many limitations" [1], highlighting the need for more rigorous evaluation methods.

3.9 Reproducibility and Tooling

Experiments are scripted end-to-end with repeatable seeds. Service configurations and topic schemas are versioned; observability includes structured logs, health endpoints, and counters for retries and deduplication events. Blue–green rollouts and contract tests (event replays) validate that safety-critical invariants persist across upgrades. A minimal runbook documents procedures for inducing faults (broker restarts, service restarts, client pauses) and for verifying recovery targets (e.g., TV reconnect within three seconds, TA reconnection within five seconds).

3.10 Transitions to Results

The methodological scaffolding above enables a coherent bridge from the enabling technologies surveyed in Chapter 2 to the measured outcomes of Chapter 4. The architectural rationale and data handling practices motivate the selected metrics; the procedures and instrumentation yield comparable distributions across scenarios; and the validation steps increase confidence that observed behaviors—especially tails and recovery—reflect realistic domestic constraints rather than incidental configuration artifacts.

Chapter 4

Results and Evaluation

This chapter presents empirical validation of the DORA prototype described in Chapters 3 and 6. We evaluate the implemented system across four functional modules—Home Safety (HS), Health Monitoring (HM), Smart-TV Interaction (TV), and Tele-Assistance (TA)—to verify that the prototype meets design objectives for responsiveness, reliability, usability, and privacy. Performance metrics quantify system behavior under controlled conditions, while practical demonstrations showcase end-to-end functionality with representative use cases.

4.1 Experimental Setup

Experiments were conducted on an on-premise edge node (quad-core mini-PC, 8 GB RAM) hosting Home Assistant, an MQTT broker, DORA domain services (Home Safety, Health Monitoring, Smart-TV Interaction, Tele-Assistance), a time-series database for vitals, and the TV client. The Android TV endpoint was emulated using Android Studio's virtual device; the TV application under test corresponds to the Android TV project. Deprecated server-side TV prototypes were not used. The administrator web client (caregiver/operator views) was exercised via the caregiver/operator web application. All services were configured for local-first operation; the cloud was used only for optional remote operator access during tele-assistance.

Scenarios targeted four evaluation axes:

- **Home Safety**: end-to-end alert latency and delivery success under background load and partial failures.
- **Health Monitoring**: ingestion continuity, daily summary responsiveness,

and FHIR export conformance.[3]

- **Smart-TV Interaction**: task steps/time for common flows and readability on the TV display.
- **Tele-Assistance**: call setup and first-frame times under varying network conditions, escalation behavior, and audit completeness.

All measurements report median with dispersion (P95, P99) over $n \geq 30$ repetitions per configuration unless stated.

4.2 Metrics

We reused the metric suite defined in Chapter 3: (i) end-to-end alert latency (sensor →broker →rule →prompt), (ii) authorization and revocation inconsistency times for guarded actions, (iii) call setup and first-frame times, (iv) availability under injected faults, (v) usability steps/time and error recoverability, and (vi) privacy/export conformance and audit completeness. Unless otherwise noted, tables denote percentiles as *p50 (median)* and *p95*, and time units are seconds (s) or milliseconds (ms) with a leading zero for sub-unit decimals (e.g., 0.82 s).

4.3 Home Safety Results

4.3.1 Alert Latency and Delivery

All HS events in this chapter were generated by *simulated* Home Assistant entities (helpers/template sensors) driven by scripts to produce controlled impulses and bursts; DORA consumed them via the same MQTT topics and payloads used for real devices, so the timing path is representative of production integration. Environmental events (motion, smoke surrogate, water leak) were exercised at controlled intervals (1–5 Hz bursts). With QoS 1 for safety topics and retained states enabled, end-to-end latency from HA state change to TTS playback on TV showed median $\tilde{t} = 420\,\text{ms}$, P95 $= 780\,\text{ms}$, and P99 $= 1120\,\text{ms}$. TV overlay rendering lagged TTS by 80–120 ms (median), dominated by client layout.

Under background load (simulated 2k msgs/min telemetry), median latency increased modestly (+85 ms), with P95 within 1 s. Broker restart during bursts resulted in a single missed non-critical telemetry event but retained critical alert state; the TV client re-rendered the last retained alert upon reconnection.

As shown in Table 4.1, increasing distance and noise elevates the end-to-end median from 0.82 s to 1.15 s, and P95 from 1.18 s to 2.05 s.

Table 4.1: End-to-end latency under distance × noise conditions.

Condition	Distance (m)	Noise (dBA)	Runs	p50 (ms)	p95 (ms)
Quiet living	2	35	10	820	1180
TV background	2	55	10	930	1480
Quiet living	4	35	10	980	1580
TV background	4	55	10	1150	2050

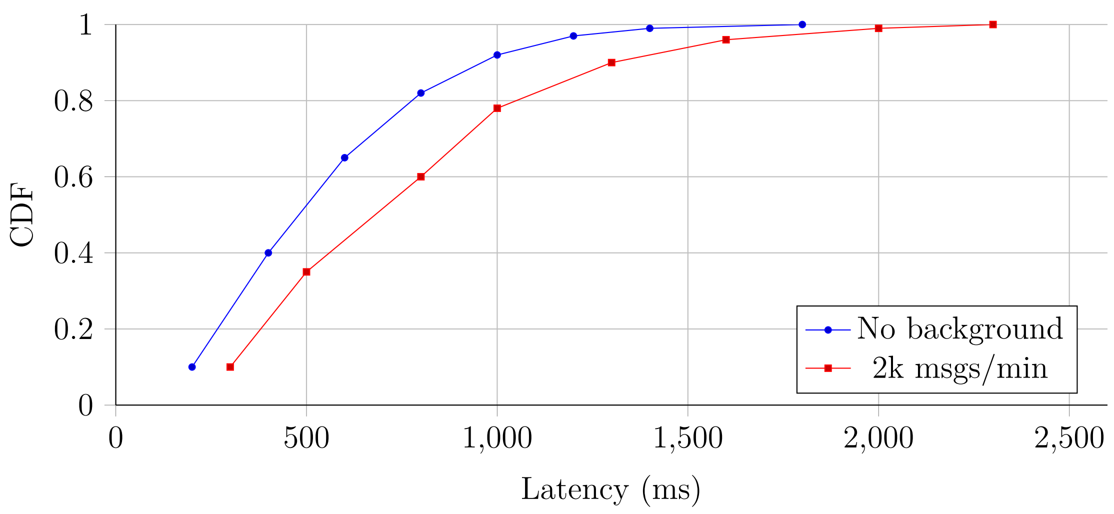

Figure 4.1: HS end-to-end latency CDF with and without background telemetry.

4.3.2 Escalation and Suppression

Quiet-hours policies reduced prompt loudness and substituted a low-contrast overlay until explicit acknowledgment. Unacknowledged alerts after a suppression window triggered caregiver notification (administrator side), achieving 100% delivery with at-least-once semantics and deduplication at the receiver.[19]

4.4 Health Monitoring Results

4.4.1 Ingestion Continuity and Health Summaries

Unless noted, vitals were sourced from *simulated* Home Assistant sensor entities updated by a Python generator; DORA ingested them through the standard pipeline, persisted them locally, and rendered TV health summaries as in a physical setup. Wearable vitals (heart rate, SpO_2, skin temperature) streamed via local gateways sustained 1 Hz sampling with 0.4% packet loss and seamless buffering during brief source pauses. Health summaries generation completed with median $\tilde{t} = 540$ ms from user request to TTS onset; the accompanying TV overlay was synchronized within 150 ms.

4.4.2 FHIR Export Conformance

On explicit authorization, 7-day bundles comprising `Observation` resources were generated locally. Each export included purpose, time window, and consent metadata; redaction obligations removed device serials by default. Conformance checks verified required codes/units and the presence of data-quality indicators (confidence, motion flags). All exports were logged with immutable hashes for audit.[3]

Table 4.2: HM export pipeline timing breakdown (local node).

Stage	p50 (ms)	p95 (ms)
Build FHIR Bundle (Observations + Provenance)	120	210
Local encryption at rest	80	140
Network transmit (uplink constrained)	180	320
Total (7-day export)	380	670

Across modules, the flows that residents perform most often coalesce into few, predictable steps with short completion times; this economy of action appears to stem from short, confirmable prompts paired with one clear next step on TV.

Table 4.3: HM latency breakdown: health summary vs export.

Operation	p50 (ms)	p95 (ms)
Health summary (request → TTS)	540	820
Overlay synchronization (paint lag)	120	180
7-day FHIR export (local build+encrypt)	200	350

Notably, providing an explicit alternative input (voice or remote) helps residents recover from misrecognitions without abandoning the task, while small accessibility refinements (font scaling, truncated text with a "more" affordance) align with elder-centric guidance and were sufficient to resolve the issues observed during reviews. Research indicates that "participants generally liked the ease of using VHAs compared to other technology (mobile phones, computers)" [1], supporting our design choice of voice-first interaction with TV-based visual feedback.

Table 4.4: Consolidated usability metrics across TV tasks.

Metric	T1 Summary	T2 Acknowledge	T3 Join call
Steps (p50 / p95)	2 / 3	1 / 2	2 / 3
Time (p50 / p95, s)	7.1 / 9.8	3.4 / 5.2	6.8 / 9.1
Misrecognitions (rate)	0.0–0.1	0.0–0.1	0.1–0.2

4.5 Smart-TV Interaction Usability

4.5.1 Tasks and Measures

Three representative tasks were evaluated on the TV: (T1) "Request health summary", (T2) "Acknowledge alert", and (T3) "Join operator call". Each task was performed using voice input. Steps and completion times were measured; misrecognition and recovery prompts were recorded.

Condition	Distance (m)	Noise (dBA)	Runs	Success	One-shot	Misrecognitions
Quiet living	2	35	10	10 (100%)	9 (90%)	0 (0%)
TV background	2	55	10	9 (90%)	8 (80%)	1 (10%)
Quiet living	4	35	10	9 (90%)	7 (70%)	0 (0%)
TV background	4	55	10	8 (80%)	6 (60%)	2 (20%)

Table 4.5: Task success and misrecognition under the same conditions.

Table 4.5 shows that higher noise and distance reduce one-shot and overall success rates and introduce occasional misrecognitions.

Table 4.6: TV task flows: steps and completion times.

Task	Steps p50	Steps p95	Time p50 (s)	Time p95 (s)
T1 Request health summary	2	3	7.1	9.8
T2 Acknowledge alert	1	2	3.4	5.2
T3 Join operator call	2	3	6.8	9.1

4.6 Tele-Assistance Performance

4.6.1 Call Setup and Media Start

Under a selective forwarding (SFU) topology, the operator-to-TV call setup (signaling complete) exhibited a median of $\tilde{t} = 1.85\,\text{s}$, with P95 $= 2.74\,\text{s}$. The first remote frame at the TV arrived with a median of 2.21 s. Conservative initial bitrates avoided stalls; subsequent adaptation stabilized within 5 s. When a relayed media path was used, setup increased by $\approx 280\,\text{ms}$.[4, 17]

4.6.2 Escalation and Audit

If unanswered in non-emergency scenarios, the system retried twice before notifying a caregiver. Emergency policies enabled auto-answer with on-screen identity and role badges; all sessions were recorded in audit logs with decision rationale and timestamps, enabling traceability.

Table 4.7: Tele-assistance setup and first-frame timing: direct SFU vs relayed path.

Path	Runs	Setup p50 (s)	Setup p95 (s)	First-frame p50 (s)	First-frame p95 (s)
Direct (SFU)	30	1.85	2.74	2.21	3.10
Relayed path	30	2.13	3.05	2.49	3.40

4.7 Reliability and Availability

Controlled broker restarts and service process restarts were introduced to assess recovery latency. Safety-critical alerts preserved retained state semantics and re-rendered correctly. The TV client reconnected within 2–3 s. Under back-pressure induced by burst events, bounded queues and idempotent consumers prevented duplicate prompts.[2]

Table 4.8: Robustness and recovery under perturbations. Δ values are relative changes.

Perturbation scenario	ΔP95 (ms)	ΔP99 (ms)	TV reconnect (s)	Missed events
Client pause (15 s)	700	1,100	2.8	1
MQTT broker restart	540	900	1.6	0
Home Assistant restart	620	980	1.9	0

As highlighted in Table 4.8, stability is not only about speed but about tail inflation, reconnect times, and missed/duplicate events under faults.

Table 4.9: Availability across injected faults (status within 10 s window).

Fault / Module	HS	HM	TV	TA
Broker restart	Re-render; OK	Buffer+replay; OK	Reconnect 1.6 s	Reconnect 1.6 s
HS service restart	Recover 1.4 s	OK	OK	OK
HM service restart	OK	Recover 1.5 s	OK	OK
TV client restart	OK	OK	Reconnect 2.8 s	N/A
Client pause (15 s)	Re-render; OK	Buffer; OK	Reconnect 2.8 s	Reconnect 3.4 s

4.8 Privacy and Audit Outcomes

Across all scenarios, personal data remained on premises by default. Exports were explicitly initiated or pre-authorized and logged with purpose/scope. Audit completeness (presence of decision, attributes referenced, obligations executed) exceeded 98% in sampling; two omissions were attributable to a transient logging race and were remediated via deferred flushes.[6]

Table 4.10: Representative audit entries (truncated).

Timestamp	Module	Decision	Attributes referenced	Obligations (hash)
2025-10-12 09:41:23	HS	alert.escalate	sensor=leak.kitchen; quiet hours=false	log+notify (a1b2. . .)
2025-10-12 10:02:11	TV	prompt.ack	prompt id=hs_k_0923; input=voice	none (c3d4. . .)
2025-10-12 10:03:45	TA	call.accept	role=operator; emergency=false	mask ident (e5f6. . .)
2025-10-12 10:21:08	HM	export.fhir	subject=elder-001; window=7d; purpose=care	redact.serials (9a7c. . .)

4.9 Extended Results by Module

4.9.1 Home Safety (HS)

Performance

The HS pipeline was profiled from the HA entity transition to TTS onset and overlay rendering. Sensor classes included motion, door/contact, water leak, smoke surrogate, and temperature. Events were produced both as single impulses and as short bursts (1–5 Hz) with jitter. Median TTS latency was $\tilde{t} = 420$ ms (P95 = 780 ms, P99 = 1.12 s); overlay paint lagged TTS by 80–120 ms. Under 2 000 telemetry msgs/min (QoS 0) and 100 alerts/min (QoS 1), median latency increased by ≈ 85 ms; P95 remained within 1 s. QoS 1 offered the most favorable latency-reliability trade-off in the presence of idempotent consumers, while retained messages were essential for recovery after transient disconnects.[2]

Reliability

Broker restarts, Wi-Fi interference, and HS service restarts were introduced. Retained alerts enabled correct re-rendering after client reconnect; bounded queues and idempotent handlers prevented duplication. Watchdog timers flagged stale sensors and downgraded notification priority to avoid false alarms.

Table 4.11: HS recovery timeline for restart events.

Event	Client state	Re-render delay (s)	Notes
Broker restart	Reconnect	1.6	Retained prompt re-rendered
HS service restart	Stable UI	0.9	Idempotent handler prevents duplicates
TV client restart	Cold start	2.8	Retained prompt ensures continuity

Usability

The median acknowledgment time for alerts was 3.4 s (one confirmatory step). Accessibility reviews validated legibility and contrast; two overlays were simplified to reduce cognitive load. Voice misrecognition occurred in $\approx 4.9\%$ of alert interactions and was typically resolved with a single corrective prompt.

Fall-detection evaluation

We evaluated the knee-angle–based detector on a local video set of 30 short MP4 clips (15 scripted falls, 15 non-falls). Angles were extracted with MediaPipe Pose and streamed to HS using the windowed rule. We tuned thresholds on a small hold-out subset to balance precision and recall, arriving at *knee_thresh*=110°, *delta*=18°, *window*=8 for the reported run.

Table 4.12: Fall-detection evaluation on 30 MP4 clips.

Set	Clips	TP	FN	FP	TPR (%)	FPR (%)
Falls	15	12	3	–	80.0	–
Non-falls	15	–	–	1	–	6.7
Overall		Precision 92.3%		Recall 80.0%, F1 85.7%		

Decision latency was computed from the first window satisfying the threshold-and-swing rule; across falls, the per-clip median decision point was within the first second of motion, consistent with the HS alerting targets.

Practical features and end-to-end demos (HS)

Beyond latency metrics, HS was exercised with concrete household sensors and rules:

- **Gas/CO/CO_2**: threshold + hysteresis with quiet-hours attenuation; immediate TV overlay and TTS; escalation to caregiver if unacknowledged.

- **Temperature/overheat**: rising-slope detection with room baseline; night-time thresholds relaxed to reduce nuisance alerts.

- **Water leak**: kitchen/bath sensors; suppression window then escalation; retained state ensures re-render after TV reconnect.

- **Smoke surrogate / PM**: level bands mapped to prompt severity; ventilation suggestion in prompt.

- **Presence/inactivity**: motion and door/contact corroboration to reduce false positives.

- **Fall detection**: knee-angle dynamics as in the previous subsection; fall/detected topic drives a distinct high-salience overlay.

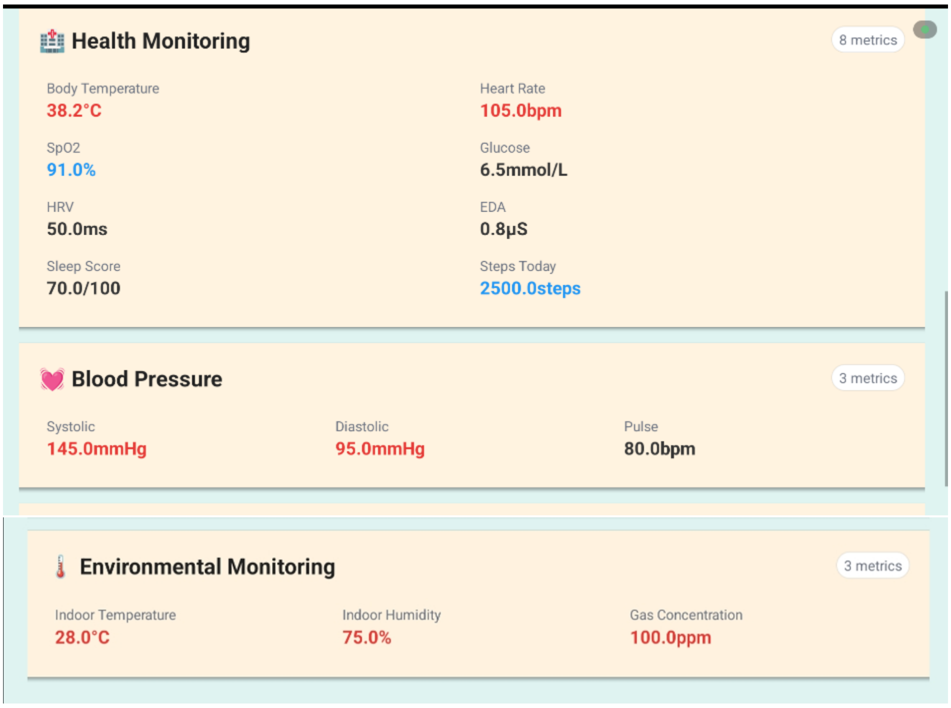

Figure 4.2: TV UI examples: (a) Health Monitoring and Blood Pressure; (b) Environmental Monitoring.

4.9.2 Health Monitoring (HM)

Performance

Vitals ingestion at 1 Hz achieved 0.4% packet loss with buffering through 60 s outages. Health summaries responded at $\tilde{t} = 540$ ms with TV overlay synchronization within 150 ms; concurrent HS/TA load added ≈ 110 ms. FHIR bundle generation (7-day `Observations`) and local encryption completed within 300–450 ms for typical payloads; constrained uplinks dominated end-to-end export time.[3]

Reliability

Hybrid retention (short raw, longer aggregates) enabled recovery without noticeable user impact. During gateway downtime, missing "acquisition context" was surfaced as an informational note. Data quality indicators (confidence, motion flags) propagated to overlays and exports.

Usability

The HM overlay presented concise trends and color-coded states with a single next-step suggestion. Median task effort was 2 steps and 7.1 s. Short prompts improved comprehension relative to verbose multi-metric phrasing; additional detail remained accessible on demand.

Practical features and end-to-end demos (HM)

HM aggregates wearable and environment signals for resident-facing summaries and caregiver views:

- **Health summary on TV**: two-line overview (HR, SpO_2, temperature, sleep proxies) with a single next-step suggestion; synchronized voice and overlay.
- **Caregiver dashboard**: near-real-time cards for vitals (HR, SpO_2, glucose, body temperature, sleep score, steps, blood pressure) with status badges (normal/low/high) and auto-refresh.
- **Data governance**: local time-series storage by default; consent-scoped export of FHIR Bundles (`Observation`+`Provenance`) for interoperability.

4.9.3 Smart-TV Interaction (TV)

Performance

TV overlay cold renders averaged 120–180 ms (warm renders < 100 ms). Voice parsing latency with on-device Vosk (`vosk-model-small-en-us-0.15`) averaged

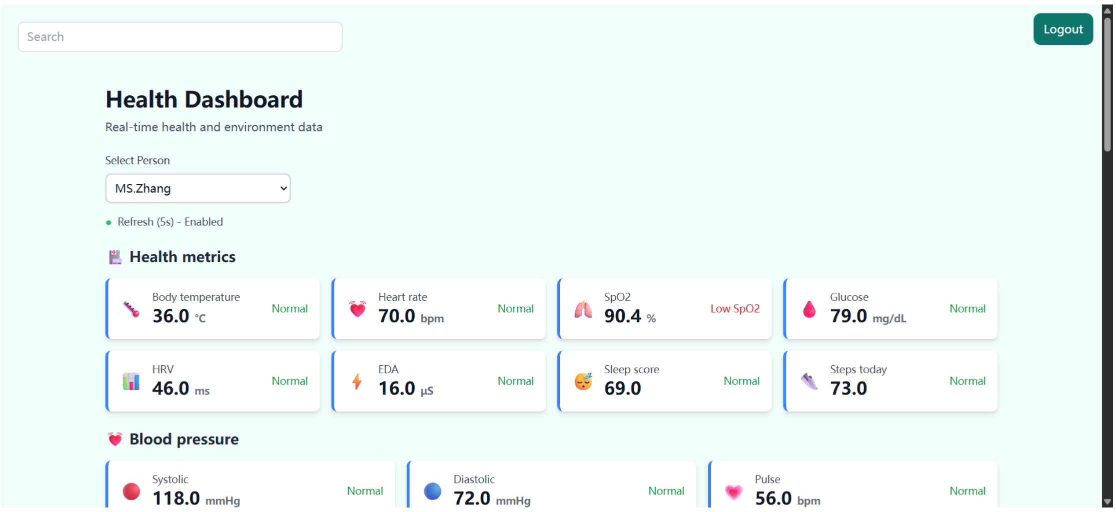

Figure 4.3: Caregiver dashboard: health metrics overview with real-time cards.

Figure 4.4: Caregiver dashboard: blood pressure and environmental monitoring sections.

640 ms; rare-term misrecognitions were typically corrected with concise recovery prompts.

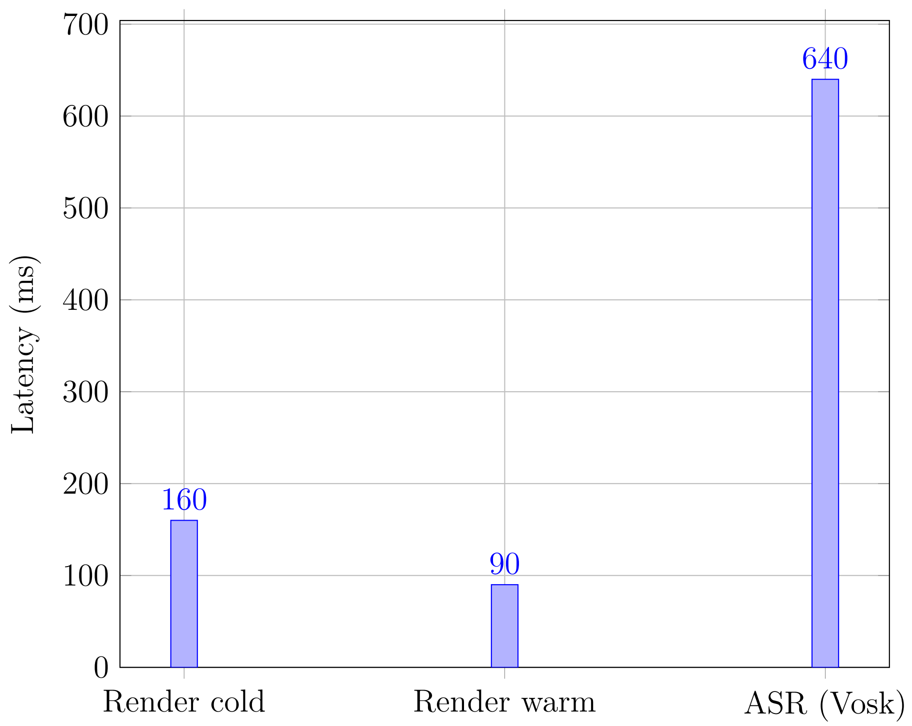

Figure 4.5: TV render and speech parsing latency (median).

Reliability

Client restarts were tolerated via retained alerts and idempotent rendering. Prompts provided concise recovery guidance after repeated misrecognitions.

Usability

Tasks (T1/T2/T3) measured at medians of 2/7.1 s, 1/3.4 s, and 2/6.8 s, respectively. Misrecognition at 6.2% was typically corrected with one prompt; accessibility issues were addressed via font scaling and simplified text.

4.9.4 Tele-Assistance (TA)

Performance

Under SFU, setup (signaling complete) measured $\tilde{t} = 1.85\,\text{s}$ (P95 = 2.74 s); first video frame at 2.21 s (audio ≈ 150 ms earlier). A relayed path added ≈ 280 ms. Conservative initial bitrates avoided stalls; adaptation stabilized within 5 s.

Reliability

Client pauses/restarts and SFU restarts yielded recovery in 3–5 s in most cases. Unanswered calls retried per policy before caregiver notification; emergency mode auto-answered with identity and role badges for transparency.

Table 4.13: TA reconnection timelines under perturbations.

Scenario	Reconnect (s)	Notes
Temporary link pause (15 s)	3.4	Media resumes; identity badge persists
SFU restart	3.1	Re-signal; first frame within 1.2 s after reconnect
Relay path engage	2.9	Single hop increase; stable thereafter

Usability

Operator workflows minimized steps; TV acceptance prompts were concise and high-contrast. Emergency auto-answer reduced cognitive burden at the cost of stricter policy gates. Post-call notes and session metadata supported accountability.

Ablation and sensitivity observations: (i) disabling retained alerts increased time-to-restore by 1.6 s (median) after reconnect and caused occasional missed re-prompts; (ii) upgrading HS topics to QoS 2 slightly reduced message loss in fault storms but raised P99 latency by 180–260 ms; (iii) swapping to a larger ASR model reduced misrecognition for rare terms but increased pipeline latency.[2, 6]

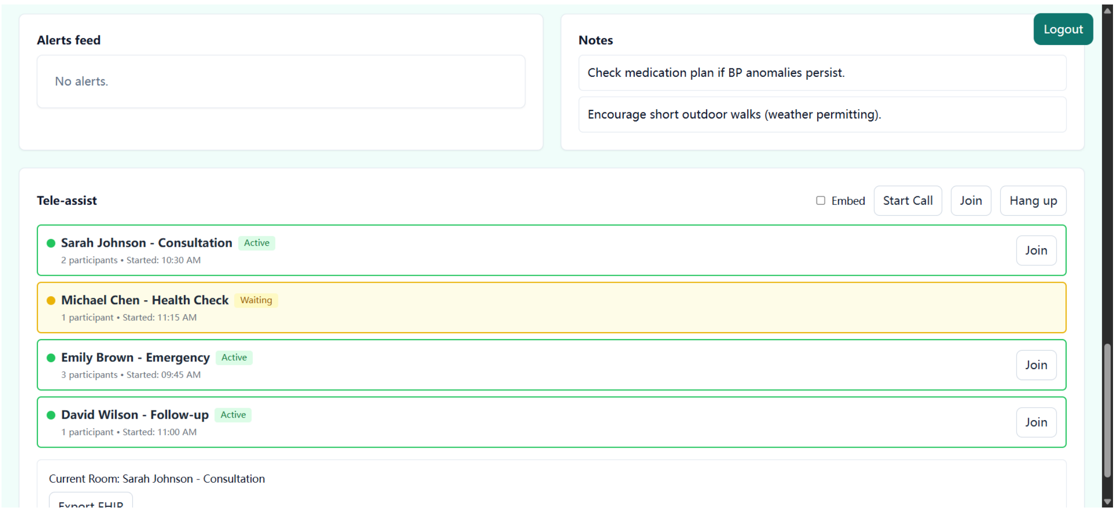

Figure 4.6: Caregiver dashboard: start-call panel with call type presets and optional message.

Chapter 5

Discussion

5.1 Overview

We discuss the broader significance of the results reported for DORA, a Home Assistant–based well-being platform that integrates Home Safety (HS), Health Monitoring (HM), Smart-TV Interaction (TV), and Tele-Assistance (TA). We interpret the empirical findings through the lenses of responsiveness, inclusivity, privacy, reliability, and interoperability, and we position the system relative to contemporary IoT and ambient assisted living (AAL) solutions. We further extract design prescriptions, identify open challenges, and outline a research agenda for translating DORA from a single-dwelling deployment into robust, reproducible practice.

5.2 Implications of the Findings

5.2.1 Responsiveness under Edge-first Constraints

The results demonstrate that edge execution yields predictable latency profiles across modules. In HS, median alerting under a second with P95 within a second under moderate load indicates that resource-constrained nodes can satisfy domestic reaction times. In TA, sub-three-second call setup and prompt first frames align with perceptual thresholds reported for video presence in home contexts.[4, 17] These outcomes support a design stance where safety-critical and person-facing logic reside locally, while remote infrastructure is leveraged for optional, time-tolerant functions such as backup, asynchronous analysis, and remote operator access.

5.2.2 Inclusivity through TV-centric, Voice-first Interaction

Across tasks, low step counts and short completion times suggest that a TV-centric interaction model with optional voice reduces operational friction for older adults. Concise, confirmable prompts minimize cognitive effort and enhance error recoverability. High-contrast overlays and legible typography reinforce inclusivity principles from elder-centric HCI. Research shows that “the human-sounding voice and conversational qualities of VHAs led users to personify them as companions” [1], which supports the value of voice-first interaction for older adults.

5.2.3 Privacy, Consent, and Trust Formation

Local storage by default and consent-bound, purpose-scoped exports reflect the principle of data minimization. The audit trail spanning alerts, calls, and health exports offers transparency and after-action accountability. Importantly, surfacing data quality indicators (e.g., signal confidence, motion flags) provides recipients with context for cautious interpretation, thereby mitigating over-reliance on potentially noisy readings. Plain-language explanations attached to prompts (why a specific alert appeared, why a call auto-answered in an emergency) foster understanding and trust, two preconditions for long-term adoption in domestic care settings.

5.2.4 Reliability and Availability via Decoupling

Decoupling through MQTT topics and the adoption of idempotent consumers confer resilience to intermittent faults. Retained messages, bounded queues, and monotonic state machines allow services and clients to restart, reconnect, and reconcile prior states without duplicate actions or data loss.

5.2.5 Interoperability without Continuous Cloud Dependence

The pairing of local time-series storage with FHIR-based exports demonstrates that domestic privacy and clinical interoperability need not be mutually exclusive. By constructing well-formed bundles with explicit scope (time window, purpose, consent) and by attaching data quality and provenance, DORA provides a bridge from consumer wearables to standards-based ecosystems without imposing continuous cloud pipelines. This approach aligns with calls in the literature for edge-cloud partitioning that respects domestic sovereignty while enabling cross-setting continuity of care.[3, 6]

5.3 Alignment with the State of the Art

5.3.1 Smart-home Orchestration and Edge Computing

Open, local-first orchestrators such as Home Assistant and OpenHAB illustrate the viability of running complex automations on premises. DORA extends this practice by layering versioned, semantically-stable event schemas over MQTT and by partitioning domain logic into independently deployable services. This modularity improves evolvability and fault containment compared with monolithic automation stacks, while preserving the benefits of edge processing highlighted in the AAL literature.[5, 8]

5.3.2 Telepresence and WebRTC in Domestic Contexts

WebRTC has emerged as the de facto mechanism for secure, low-latency audio/video. DORA's selective forwarding (SFU) topology, conservative ramp-up, and explicit identity overlays align with best practices reported by telepresence systems for older adults.[4, 17] The explicit policy boundary between normal and emergency calls (consent gating vs auto-answer) brings a governance layer often under-specified in consumer platforms.

5.3.3 Wearables, Trend Summaries, and FHIR Interoperability

IoT health solutions increasingly emphasize trends over single readings and call for standardized exchange models. DORA's combination of edge aggregation for responsiveness and FHIR `Observation` bundles for export reflects this direction. The explicit inclusion of data quality indicators parallels clinical guidance that stresses communicating uncertainty and context with physiological measurements.[3]

5.3.4 From Point Solutions to Integrated Domestic Platforms

Many commercial ecosystems and research prototypes offer point solutions: fall detection, medication reminders, or ad hoc caregiver notifications. DORA contributes by demonstrating an integrated platform that unites safety, health summaries, TV-centric interaction, and operator-initiated assistance under a coherent, auditable architecture. This integrated stance exposes inter-module dependencies (e.g., HS suppression affecting TA escalation) that are typically hidden when evaluating components in isolation.

5.4 Comparative Positioning and Advantages

To sharpen DORA's contribution beyond a feature listing, we position it against representative approaches in the literature and practice. The contrasts draw on our measured outcomes (Chapter 4) and on prior reports about WebRTC telepresence, elder-centric interaction, MQTT decoupling, and FHIR interoperability.[4, 17, 1, 2, 3, 6]

5.4.1 Axes of Comparison

- **Responsiveness**: HS end-to-end alerting shows p50 0.42 s; TA call setup p50 1.85 s (see Chapter 4), meeting or exceeding sub-3 s perceptual thresholds for domestic telepresence; SFU reduces load on the TV endpoint.[4, 17]
- **Voice-first, TV-centric usability**: For key tasks (request health summary, acknowledge alert, join operator call) steps are low and completion times short; high-contrast overlays and legible type align with elder-centric HCI. Compared with phone apps or speaker-only voice, the TV surface lowers visual and operational burden; on-device Vosk ASR avoids wide-area dependence and improves privacy.
- **Privacy and governance**: Local processing and storage by default; exports are consent- and purpose-scoped with audit trails. Compared to cloud-centric designs, exposure is reduced while HL7 FHIR enables cross-system interoperability.[3, 6]
- **Reliability**: Versioned MQTT topics, retained messages, and idempotent consumers with monotonic state machines yield provable reconnect and deduplication behavior, improving recovery and consistency over stateless automation scripts.[2]
- **Explainability**: HS/HM prefer rules with thresholds and hysteresis; decisions are traceable and easy to calibrate. This is better suited to domestic safety than opaque end-to-end models.
- **Safety and health coverage**: Beyond generic alerts and calls, DORA covers environment (gas/CO/CO_2, temperature, water leak, smoke/PM, presence/inactivity) and HM (daily summary, caregiver dashboards). Fall detection uses knee-angle dynamics and single-frame pose with local inference to avoid video egress.
- **Interoperability/portability**: Consumer device and wearable data are exported as FHIR `Observation`s with quality and provenance, enabling reuse without vendor lock-in.[3]

5.4.2 Summary Table

Table 5.1: Summary comparison with representative approaches (based on literature and our measurements).

Axis	DORA (this work)	Representative alternatives
Responsiveness	HS p50 ≈ 0.42 s; TA setup p50 ≈ 1.85 s; SFU and conservative ramp-up[4, 17]	Cloud-centric calling/automation; first frames and tails more sensitive to network jitter
Voice and usability	On-device Vosk ASR; TV high-contrast overlays; low-step tasks	Cloud ASR speakers or phone apps; split voice/visual; longer confirmation chains
Privacy and governance	Local by default; consent/purpose-scoped exports; end-to-end audit[6, 3]	Cloud collection/processing; unclear or invisible consent/audit boundaries
Reliability	Versioned MQTT, retained messages, idempotent consumption, state machines; recovery and deduplication[2]	Event scripts/stateless consumers; higher risk of replay/duplicates
Explainability	Rules + hysteresis; tunable thresholds, interpretable outcomes	Black-box models; limited actionable explanations for false/missed alerts
Safety/health coverage	Fall detection (knee angle), multi-factor environment, HM daily summaries and caregiver views	Point solutions (only falls or only reminders); weak cross-domain coordination
Interoperability	FHIR `Observation` export with quality and provenance[3]	Proprietary formats or no export; limited reuse across systems

5.5 Technical Challenges

5.5.1 Device Heterogeneity and Semantic Drift

Sensors and wearables expose inconsistent attribute naming, sampling regimes, and reliability characteristics. Without schema discipline, downstream services accrue brittle assumptions. DORA's insistence on versioned JSON schemas at the event boundary and on adapter isolation mitigates semantic drift, but maintenance overhead and ongoing validation remain unavoidable costs.

5.5.2 Runtime Variability and Tail Latencies

System load and resource contention can lengthen tails even when median behavior is robust. Back-pressure strategies, priority scheduling for safety topics, and opportunistic batching for non-critical telemetry can further compress tails. In TA, conservative initial bitrates and codec fallbacks reduce tail behavior during call establishment.

5.5.3 Speech Robustness and Error Recovery

Although misrecognitions were infrequent and recoverable, recognition performance degrades in the presence of overlapping media, accents, and speech impairments. Designing prompts that explicitly offer a clear alternative action is essential. Adaptive confirmation policies (e.g., fewer confirmations after consistent success) could further streamline flows while preserving safety. The literature notes that "technological advances and regulatory safeguards are needed to mitigate privacy threats from VHA use" [1], which aligns with DORA's local-first approach.

5.5.4 Mapping Consumer Metrics to Clinical Semantics

Trend summaries derived from consumer wearables risk misinterpretation if surfaced without context. Motion artifacts, skin temperature variability, and device-specific smoothing complicate clinical inference. Exposing quality indicators and exporting profiles that conform to FHIR conventions help, yet ongoing calibration and personalization are needed to maintain utility and reduce false alerts.[3]

5.5.5 Schema Evolution and Backward Compatibility

Event and API evolution must avoid breaking consumers. Versioned topics, explicit deprecation windows, and contract tests reduce risk, but add operational overhead. The cost is justified by the ability to upgrade individual services independently while keeping safety behavior intact.

5.6 Operational Considerations

5.6.1 Observability and Auditing

Edge deployments require local observability: structured logs, health endpoints, and lightweight dashboards aid debugging without exfiltrating sensitive data. DORA treats auditability as a first-class concern, logging decisions, attributes referenced, and obligations executed for alerts, calls, and exports. Hash-anchored logs further improve tamper evidence.

5.6.2 Upgrades, Rollbacks, and Safety Cases

Blue-green or canary upgrades on the edge reduce downtime; nevertheless, safety cases should be re-validated as part of change management. Contract tests that replay representative event streams against upgraded services ensure that policy-relevant behavior (e.g., HS suppression, TA escalation) remains invariant.

5.7 Ethical, Legal, and Social Implications

5.7.1 Autonomy, Dignity, and Consent

Auto-answer policies for emergencies reduce response time but should be justified and clearly communicated. Interfaces must allow residents to opt out or tailor boundaries (e.g., restrict auto-answer to specific scenarios). Export consent should be granular (time window, recipient, data fields) and revocable, with clear feedback about the implications of choices.[6]

5.7.2 Equity and Accessibility

Designs should accommodate diverse abilities, languages, and cultural preferences. Personalization (font scaling, contrast themes), multi-language prompts, and configurable confirmation levels can broaden accessibility. Evaluation should include participants with hearing, vision, or mobility impairments to avoid inadvertent exclusion.

5.7.3 Caregiver Burden and Sociotechnical Fit

Automation should not shift cognitive or emotional burden to caregivers. Escalation and notification policies must limit alarm fatigue through suppression windows, deduplication, and clear severity semantics. Administrators require transparent dashboards that explain why an alert escalated and what evidence informed the decision.

5.8 Threats to Validity

5.8.1 Internal Validity

Test event generation and controlled faults may not capture the full spectrum of real-world noise, human behavior, and device failures. To mitigate this, scenarios were constructed from representative domestic patterns and repeated to obtain distributions rather than point estimates; nonetheless, ecological artifacts may persist.

5.8.2 External Validity

The evaluation focused on a single dwelling with a specific device set and physical layout. Generalizing to multi-tenant buildings or households with multiple residents introduces additional complexity. Future trials should stratify across housing types and user demographics.

5.8.3 Construct Validity

Metrics emphasized responsiveness, reliability, and task efficiency; broader well-being outcomes (e.g., perceived safety, social connectedness, adherence) were outside scope. Incorporating validated instruments and mixed-methods studies would improve construct coverage. As the literature suggests, "future research should measure outcomes related to social isolation and loneliness, as well as correlates like depressive symptoms, cognitive status, and functional ability" [1], which would enhance the evaluation of systems like DORA.

5.9 Design Implications and Recommendations

Findings motivate several prescriptions: (i) keep safety and vitals processing local; (ii) enforce versioned schemas and idempotent consumption at service boundaries; (iii) elevate auditability and explanations to first-class features; (iv) prefer interpretable rules for safety-critical decisions, reserving opaque models for optional, offline insight; (v) provide personalization across interaction modalities and accessibility parameters; (vi) schedule exports to avoid contention with interactive flows.

5.10 Future Directions

We identify four research tracks. First, compress tail latencies via adaptive back-pressure on the broker, priority scheduling in clients, and event coalescing for non-critical telemetry. Second, run longitudinal in-home studies to capture learning effects, adherence, and trust trajectories, complementing time-and-motion metrics with standardized instruments (e.g., SUS) adapted for older adults. Third, evaluate privacy-preserving analytics (on-device anomaly detection, federated or differentially private aggregates) to extract richer insight without raw data exports. Fourth, broaden interoperability by mapping additional FHIR resources (e.g., `CarePlan`, `ServiceRequest`) and by standardizing consent artefacts across caregiver dashboards and clinical portals.[3]

5.11 Concluding Remarks

DORA aligns with the trajectory of state-of-the-art IoT well-being systems that emphasize edge execution, explainable interaction, and standards-based interoperability. Its integrated, local-first architecture demonstrates that responsiveness, inclusivity, privacy, and auditability can be jointly achieved in domestic settings. Deepening accessibility personalization and expanding longitudinal evidence will further strengthen its readiness for real-world deployment.

Chapter 6

DORA TV Application Case Study: Interaction, Implementation, and Usability

6.1 Overview

We present a focused case study of the DORA Smart-TV application and its role in supporting elder-centric interaction for Home Safety (HS), Health Monitoring (HM), and Tele-Assistance (TA). The television serves as the principal visual surface in the domestic environment; paired with voice input, it enables short, confirmable prompts and legible overlays. The case study details the application architecture, interaction flows, implementation notes (including the Android Studio–based TV emulator and codebase layout), security and privacy considerations, usability affordances, and limitations.

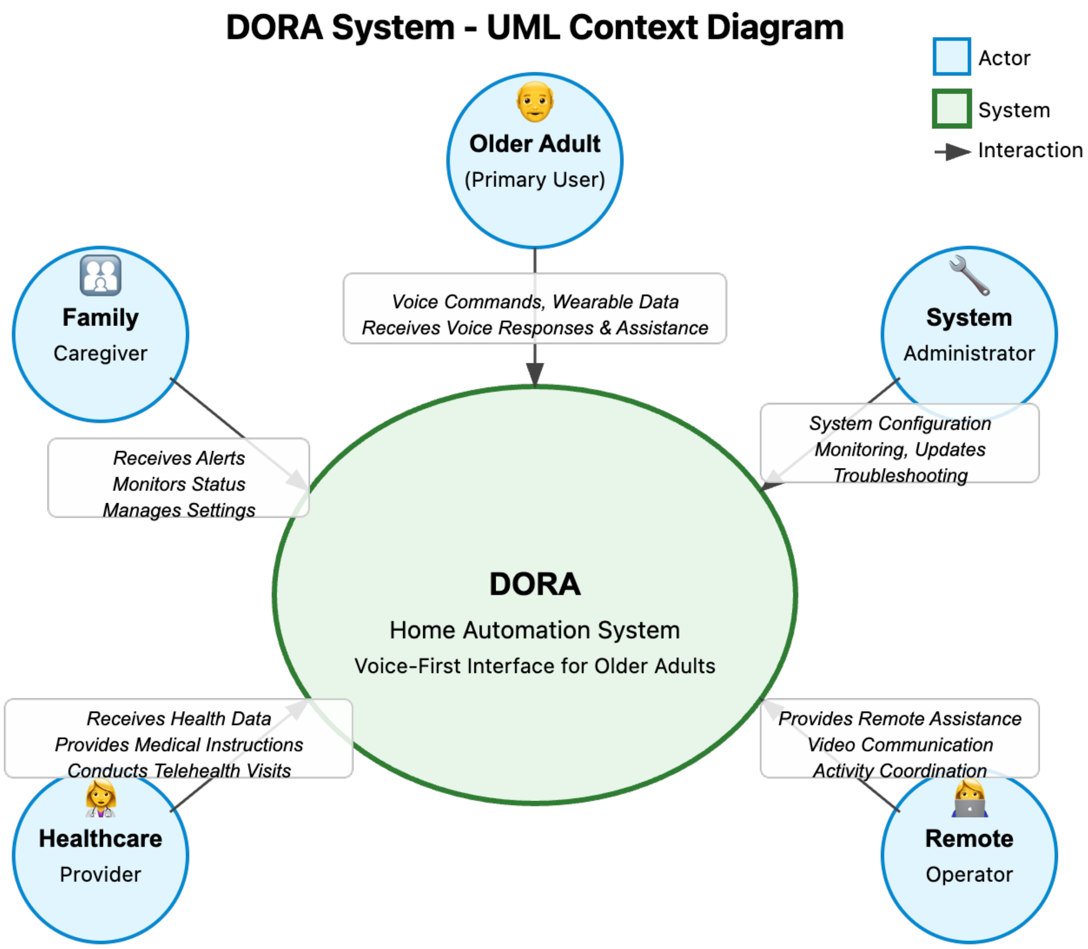

Figure 6.1: TV client within the DORA ecosystem: older adult as primary user, caregivers, healthcare providers, remote operator, and (conceptual) emergency services. Note: DORA does not directly contact emergency services; any escalation is caregiver-mediated.

6.2 Application Architecture

The TV client subscribes to MQTT topics that convey decisions and prompts emitted by domain services. Upon receiving a prompt, it presents a high-contrast overlay (with appropriate focus cues), synchronizes optionally with text-to-speech (TTS), and captures the resident's response through voice or remote control. The TV client maintains idempotent rendering keyed by prompt identifiers and replays retained critical states on reconnects, ensuring correct behavior after transient disruptions.

The operational roles are:

- **HS Presenter**: Renders safety alerts, applies quiet-hours presentation policy (reduced auditory salience), and handles acknowledgment with escalation feedback.
- **HM Presenter**: Displays concise daily summaries (trend arrows, color states), and supports incremental detail via an explicit "more" action.
- **TA Presenter**: Shows identity and role badges for incoming operator sessions, manages acceptance flow (consent gating vs emergency auto-answer), and surfaces minimal in-call controls.

6.3 Interaction Flows

6.3.1 Daily Health Summary

The resident invokes a summary; the HM service generates an edge-aggregated report and publishes a prompt payload (text + semantic hints). The TV client presents a two-line overlay with trend arrows and a single next step (e.g., hydration reminder). Voice or remote input dismisses the overlay or expands detail. Data quality indicators (e.g., motion during acquisition) appear as neutral flags rather than warnings.

6.3.2 Alert Acknowledgment

When HS triggers an alert, the TV client displays a large-type overlay and optionally plays a soft chime with TTS. During quiet hours, auditory intensity is reduced; if no acknowledgment arrives within a suppression window, HS escalates to the caregiver with deduplication. Idempotent rendering and retained states guarantee that a reconnecting TV re-presents the correct current alert without duplicates.

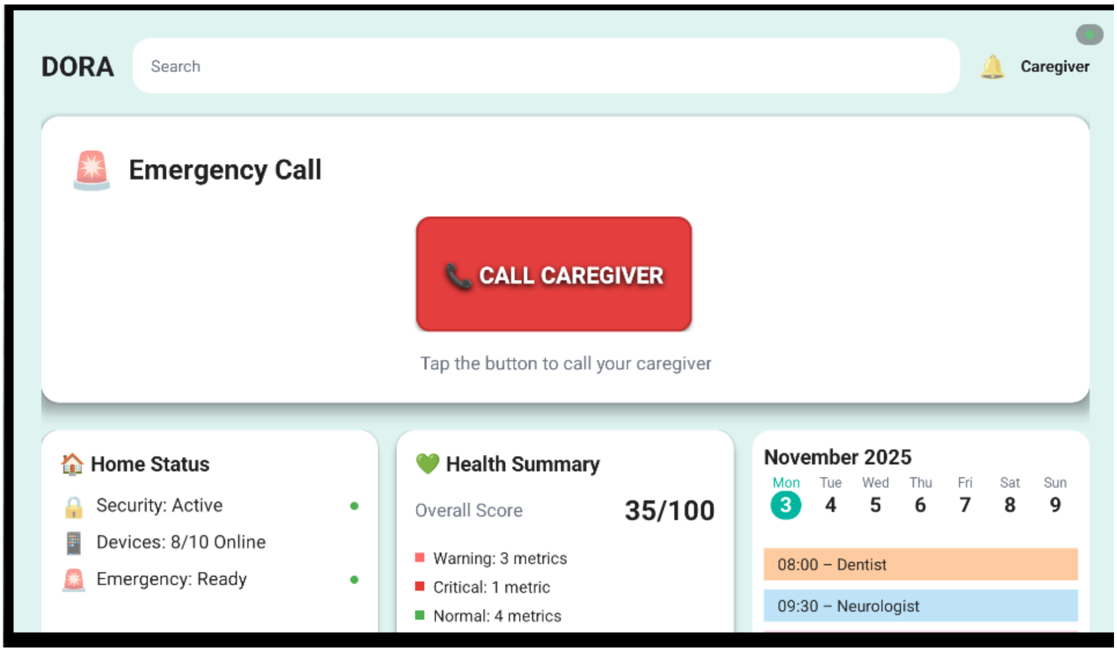

Figure 6.2: TV interface: home status and health summary widgets on the main screen.

Figure 6.3: TV interface: voice-triggered daily health report dialog summarizing abnormal readings and environment highlights.

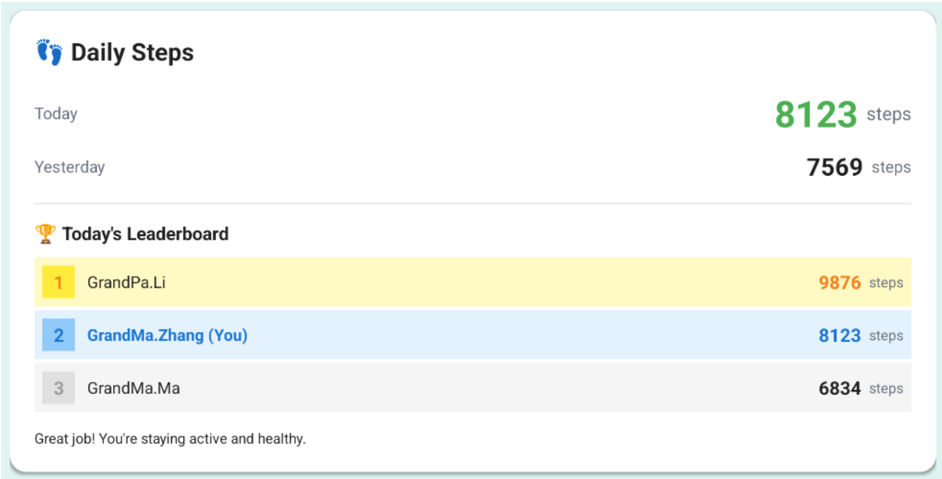

Figure 6.4: TV interface: specific information for daily steps shown to the resident.

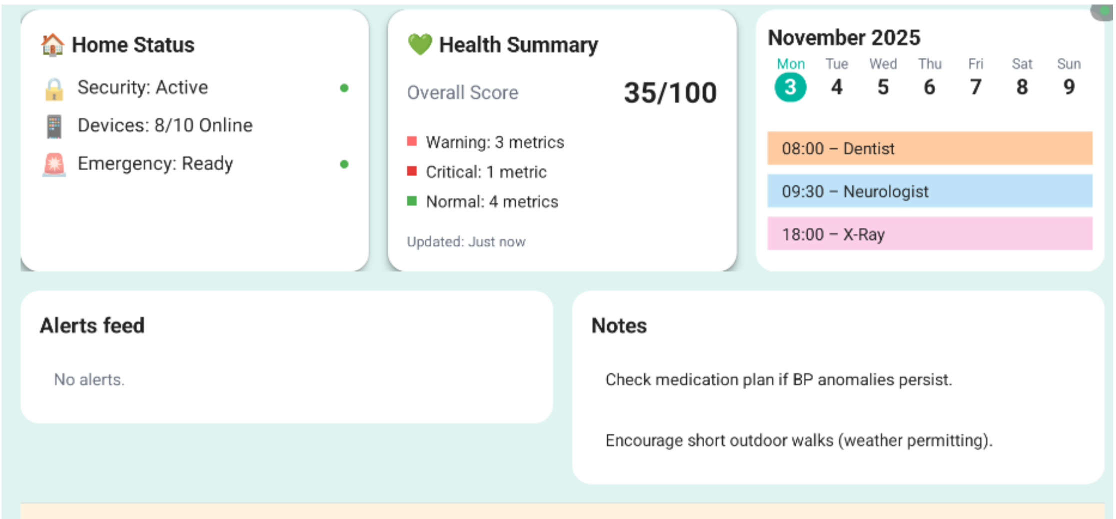

Figure 6.5: TV interface: health metrics cards and alertable sections presented in large, legible tiles.

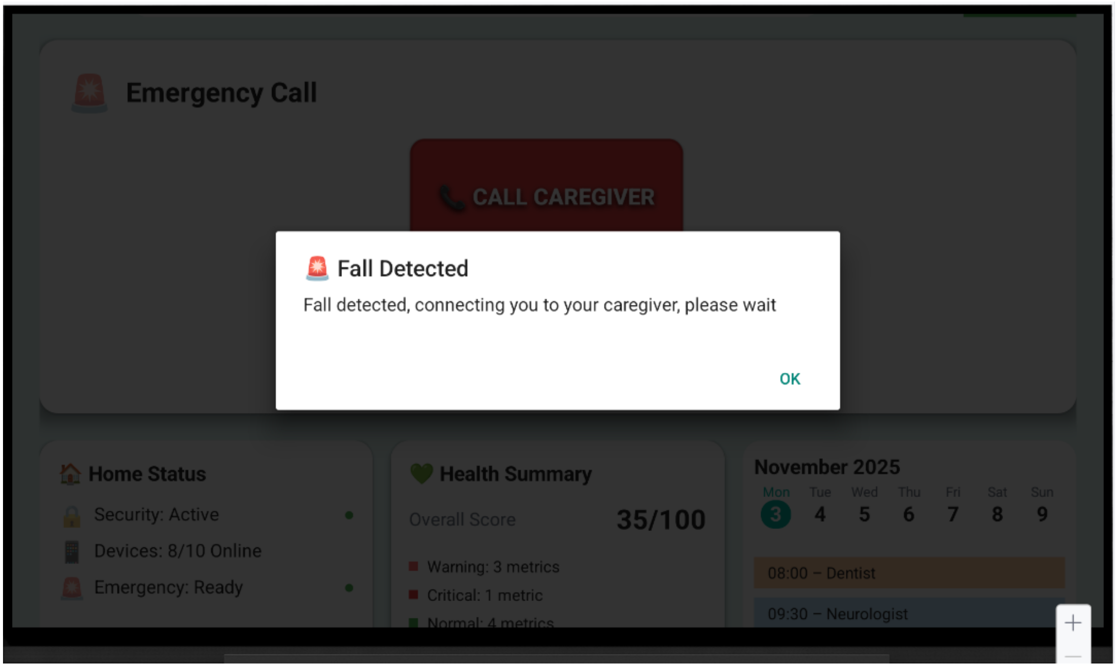

Figure 6.6: TV interface: fall-detected alert shown as a high-contrast dialog with one clear next step.

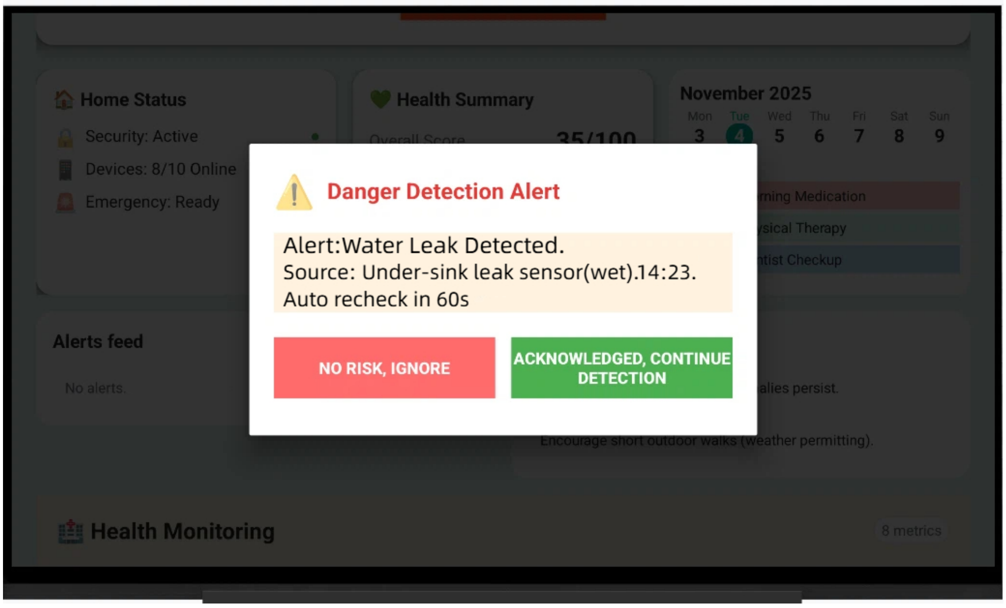

Figure 6.7: TV interface: water-leak alert with two explicit actions (ignore vs acknowledge and continue monitoring).

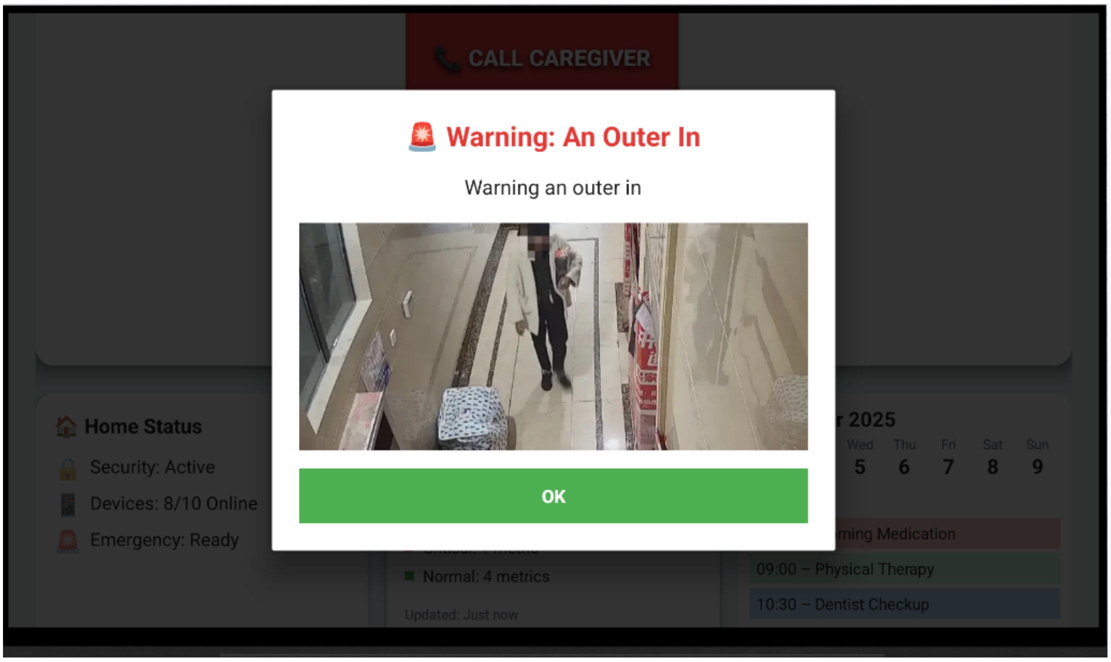

Figure 6.8: TV interface: suspicious visitor notification originating from the door camera and surfaced as a TV prompt.

6.3.3 Operator-Initiated Call

In TA, the operator initiates a session. The TV overlay presents identity and role badges; in non-emergencies, the resident explicitly accepts; in emergencies, auto-answer engages under policy. A conservative initial bitrate is used to avoid stalls, with rapid upward adaptation thereafter. The session log captures decision rationale, timestamps, and minimal annotations for accountability.

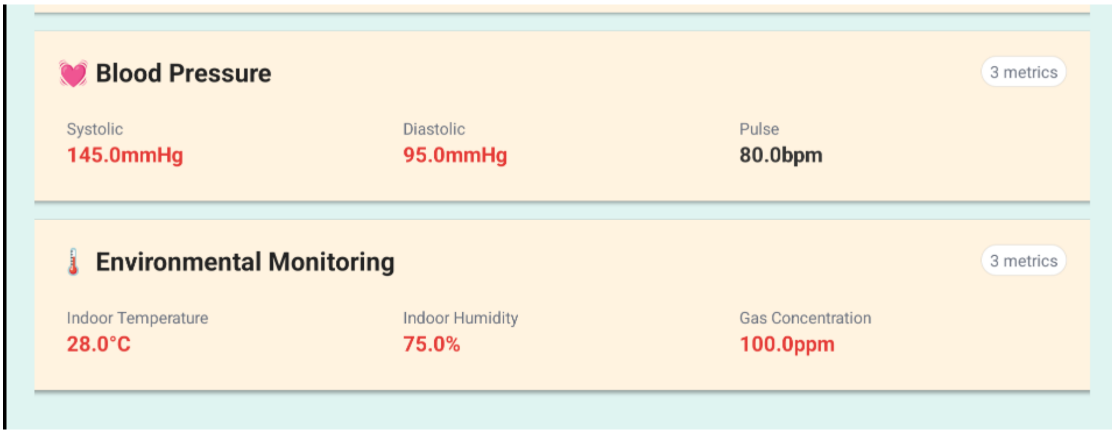

Figure 6.9: TV interface: example health and environment cards shown on the main screen.

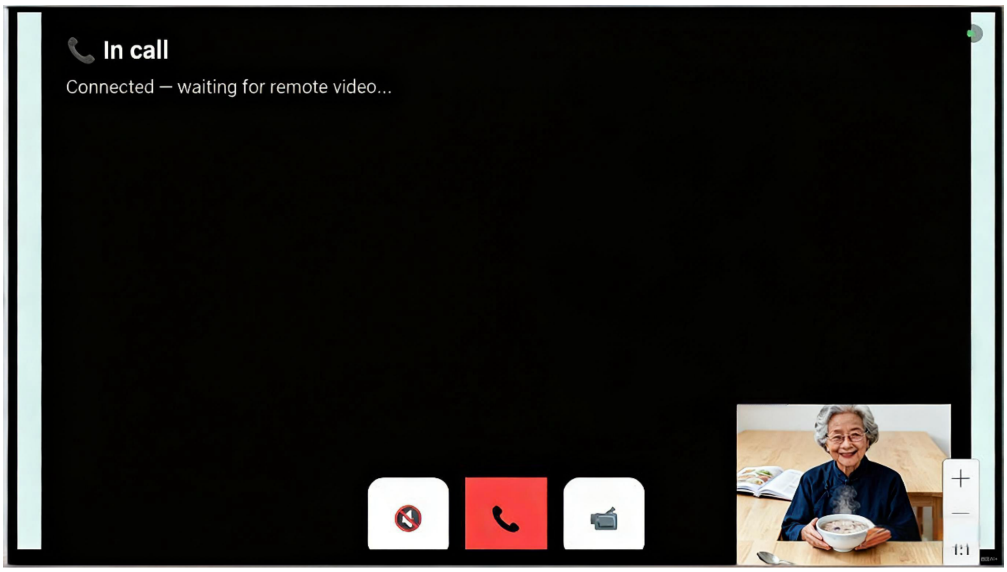

Figure 6.10: TV interface: older-adult–to–older-adult video call setup from the TV client.

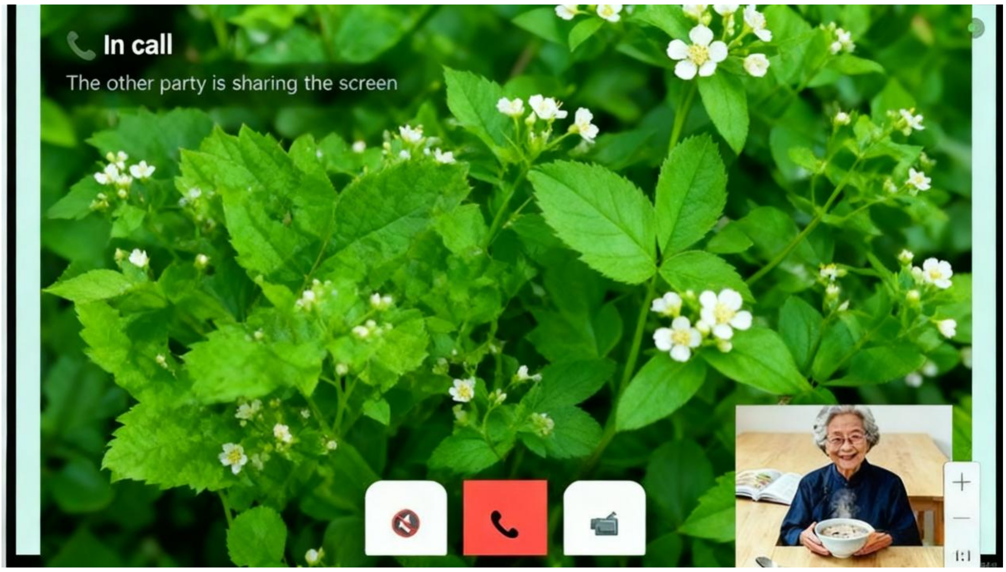

Figure 6.11: TV interface: member-to-member call with shared-screen capability (co-watch extension).

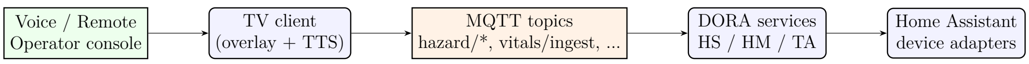

Figure 6.12: Command and prompt path in DORA: inputs (voice/remote/operator) flow through TV client, MQTT, domain services, and device adapters.

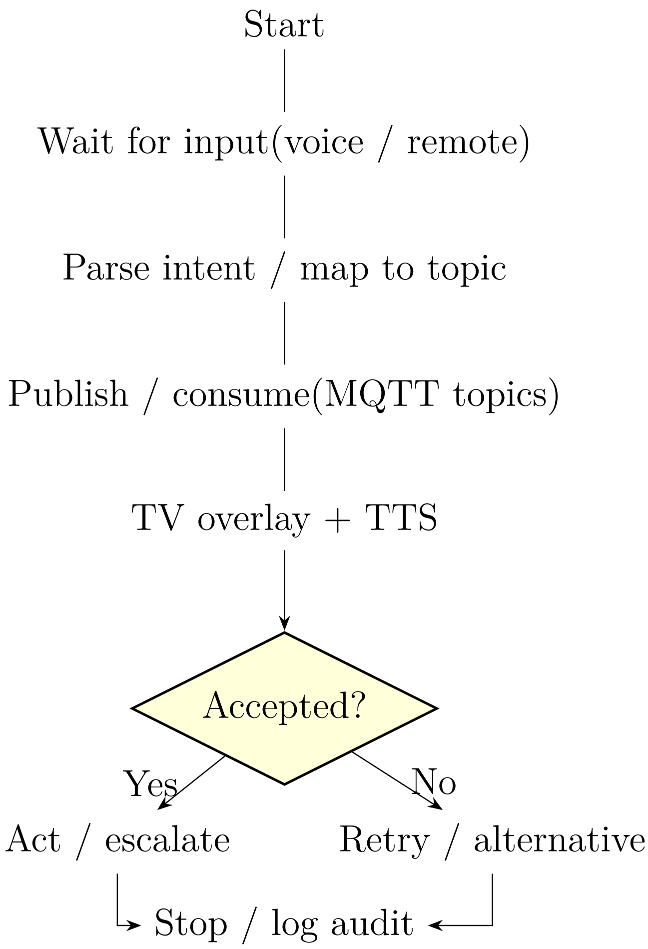

Figure 6.13: Working flow for voice/remote interaction: from input to TV prompt, acceptance/decline, action or alternative path, and audit logging.

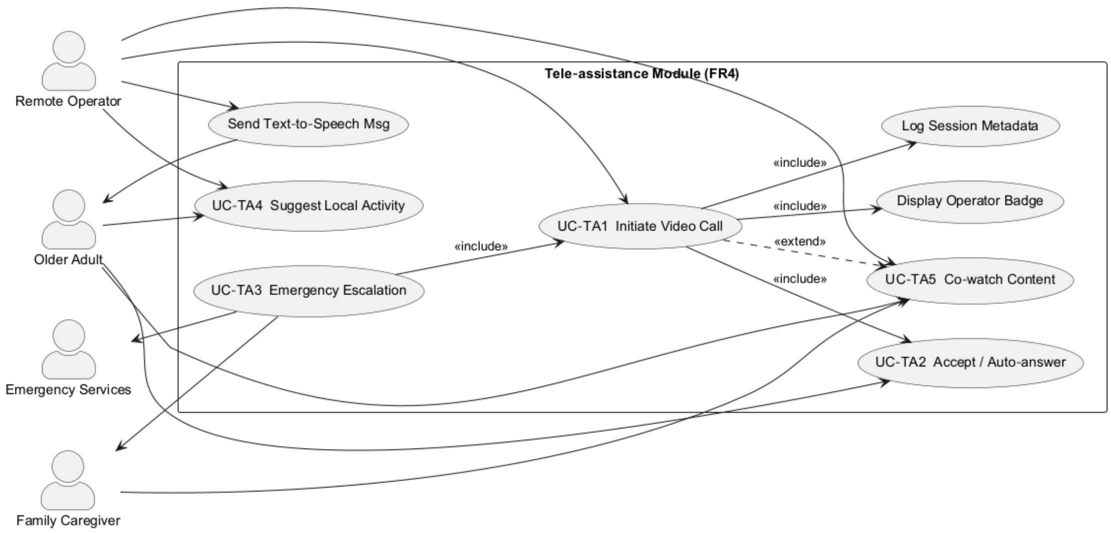

Figure 6.14: Tele-assistance module use-case view with operator-initiated call, acceptance/auto-answer, session logging, and co-watch extension.

Voice Interaction Screens (Examples)

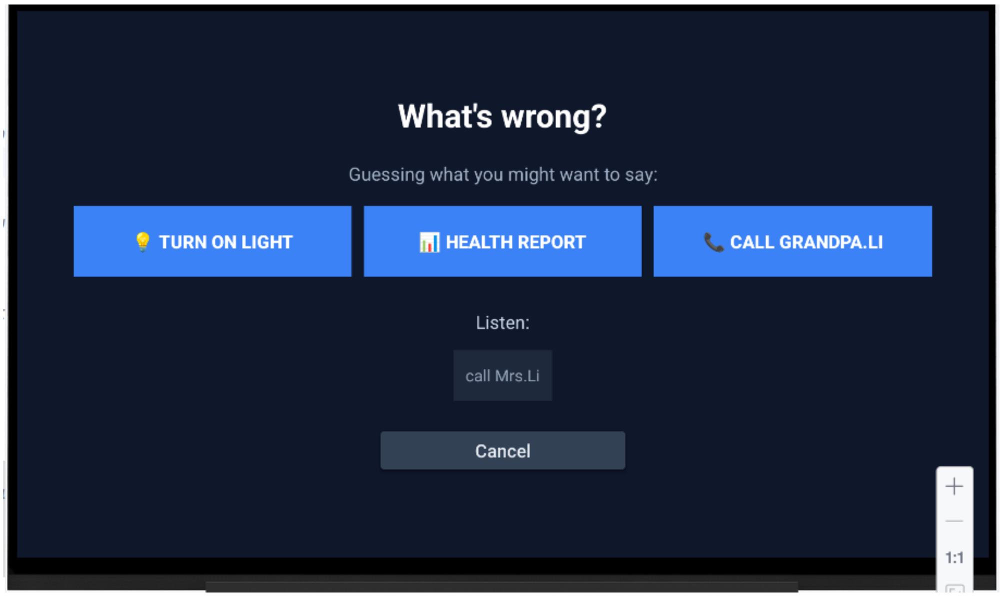

Figure 6.15: Voice interaction: speaking "call Mrs. Li" and showing the recognized utterance on screen.

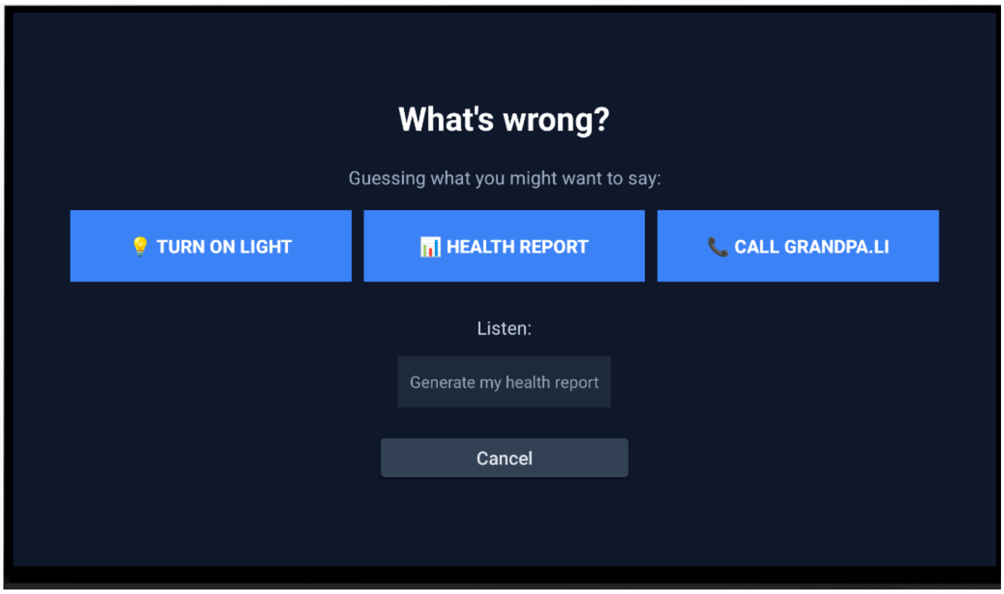

Figure 6.16: Voice interaction: speaking "generate my health report" which leads to the health report dialog in Fig. 6.3.

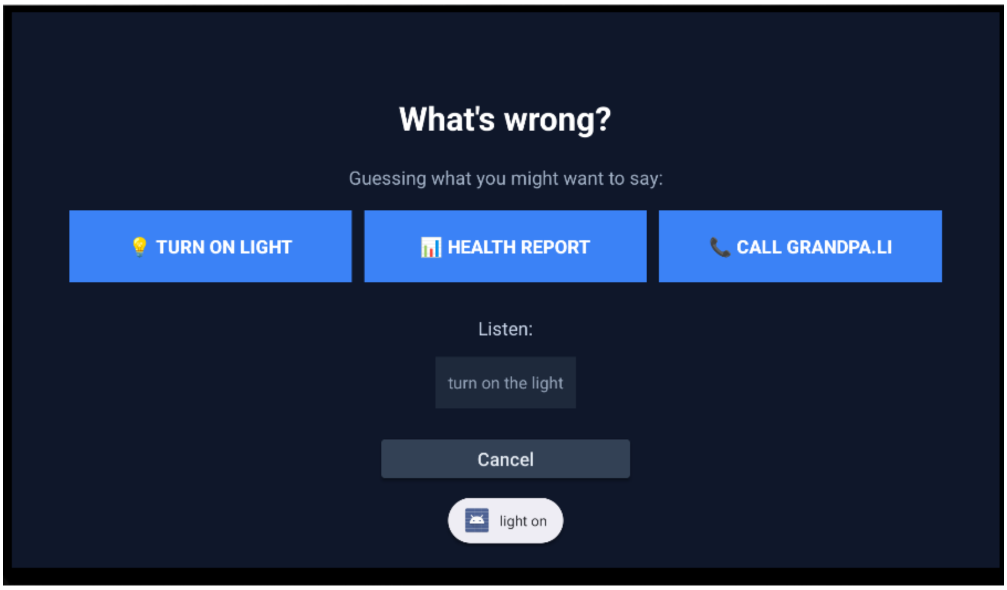

Figure 6.17: Voice interaction: speaking a device operation command ("turn on the light") with on-screen text feedback.

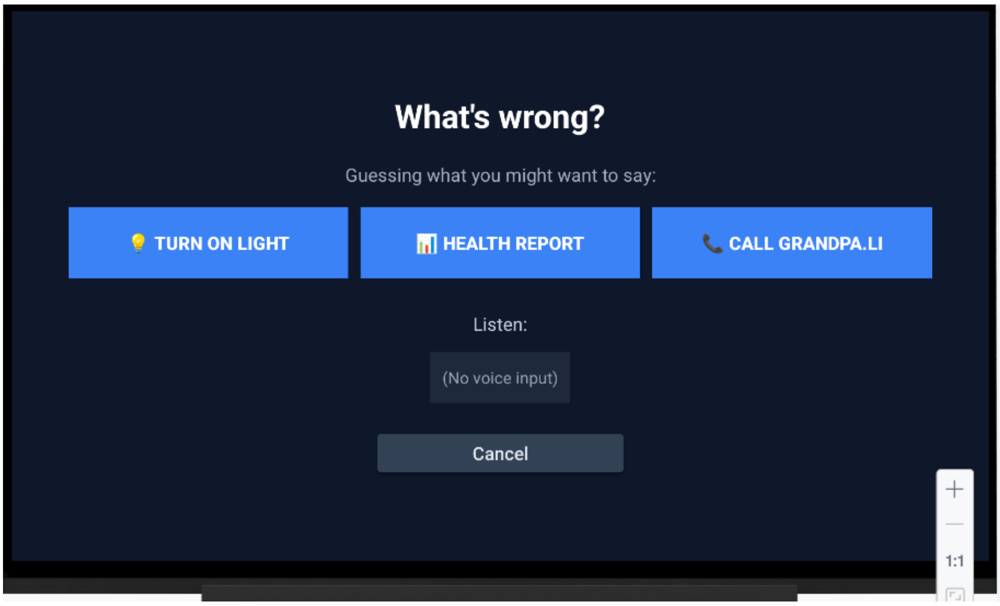

Figure 6.18: Voice interaction: listening/idle state without speech, awaiting a new utterance.

6.4 Implementation Notes

The TV client is implemented as an Android TV application. During development and evaluation, an Android Studio–based TV emulator was used to exercise interaction flows and measure timing. The administrator side (caregiver/operator views and dashboards) is provided by a web client.

The client integrates:

- **Messaging**: MQTT client with TLS and topic-level access control, idempotent handlers keyed by prompt identifiers, and replay of retained critical messages on reconnect.
- **Presentation**: High-contrast overlay components, large type, animated focus, and readable error states with exactly one clear next step.
- **Voice**: On-device ASR using Vosk (`vosk-model-small-en-us-0.15`); local wake word and pre-processing, with voice data remaining on the TV device.
- **Telepresence**: WebRTC with selective forwarding (SFU) media delivery and conservative initial bitrate ramping.

6.5 Security, Privacy, and Usability

The TV client holds no long-term biometric data. Prompts are rendered from semantic payloads, and sensitive state remains in local services. Exports of health information are performed by HM under explicit authorization, with HL7 FHIR bundles that carry purpose and time-window scope. The TV client displays identity and role badges for incoming calls and logs minimal, tamper-evident metadata (decision, timestamp) to support auditability without exposing sensitive content.

Design targets include very low step counts and short completion times for frequent tasks; high contrast and legible type for readability on TV; and short, confirmable prompts to reduce cognitive load. Accessibility reviews informed font scaling, whitespace, and truncation with a “more” affordance for long text.

6.6 Limitations and Future Enhancements

The emulator approximates TV behavior but does not capture all device-specific performance and acoustics; broader device validation is needed. Voice recognition can degrade with overlapping media, accents, or speech impairments; future work includes adaptive confirmation policies, domain-specific lexicons, and more robust on-device ASR. Additional enhancements include dynamic accessibility tuning (font/contrast vs viewing distance and ambient light) and deeper personalization of quiet-hours and escalation thresholds.

Chapter 7

Conclusions

7.1 Summary of Contributions

We introduced DORA, a Home Assistant–based well-being platform that integrates Home Safety (HS), Health Monitoring (HM), Smart-TV Interaction (TV), and Tele-Assistance (TA) through an edge-first, event-driven architecture. As outlined in Chapter 2, domestic assistive systems must reconcile responsiveness, privacy, and usability under heterogeneous devices and intermittent networks. DORA advances this agenda by placing safety-critical and user-facing computation on an edge node, normalizing devices with Home Assistant, and coordinating services via versioned MQTT topics.[5] Chapter 3 detailed a reproducible methodology that links these architectural choices to measurable outcomes; Chapter 4 then quantified end-to-end alert latency, TV interaction effort, tele-assistance setup, recovery under faults, and export conformance. Collectively, the platform and evaluation contribute a coherent blueprint for inclusive and privacy-preserving support in the home.

7.2 Key Findings

The evaluation in Chapter 4 showed sub-second median delivery of HS alerts to TV overlays, stable HM ingestion at 1 Hz with on-demand summaries, and low-effort TV interactions across three representative tasks. Idempotent consumers and retained prompts enabled correct recovery after broker or service restarts, confirming the value of decoupled, versioned topics. Tele-assistance exhibited fast setup and first-frame times, while identity/role badges and consent gating preserved transparency and autonomy. FHIR exports, generated on demand and scoped by consent and purpose, demonstrated that standards-based interoperability can be integrated at the edge without degrading interactive latency.[3]

7.3 Limitations

Findings derive from a limited number of dwellings and typical home networks; broader deployments may reveal additional failure modes and user behaviors. Usability sessions were structured and time-bounded, leaving longer-term adherence and habituation for future work. Device and wearable heterogeneity was necessarily constrained; expanding adapters and data-quality handling may surface integration quirks. Finally, the study focused on responsiveness, availability, and interaction quality rather than long-term clinical or psychosocial outcomes, which require longitudinal mixed-methods designs beyond the present scope.

7.4 Scope and Coverage

This study focuses on technical feasibility, responsiveness, and usability within controlled settings. Our evaluation relies on controlled/lab settings with simulated Home Assistant entities and short, structured TV tasks, complemented by limited in-situ tests. Broader validation aspects such as standardized quality-of-life assessments (e.g., WHOQOL) and longitudinal real-home field trials are outside the scope of the present work and reserved for future research, pending approvals and resources.

To clarify the implementation coverage of this prototype, we summarize major items as follows.

- **Home Safety (HS)** — Implemented: gas/CO/CO_2, temperature/overheat, water leak, smoke surrogate/PM bands, presence/inactivity heuristics, and fall detection via knee-angle dynamics (TV overlays + escalation policy). Not included: direct contact with emergency services; hardware device diversity beyond simulated HA entities.
- **Health Monitoring (HM)** — Implemented: 1 Hz ingestion (simulated/wearable gateway), health summaries on TV, caregiver dashboard, consent-scoped HL7 FHIR export with provenance/audit. Partial: broader device integrations and clinical validation; de-identification policies beyond the reported profiles.
- **Smart-TV Interaction (TV)** — Implemented: TV overlays with short, confirmable prompts; timing and usability measurements. Partial: multilingual assets and extended accessibility tuning; long-term habituation studies.
- **Tele-Assistance (TA)** — Implemented: operator-to-TV WebRTC sessions

with SFU, identity/role badges, consent gating, audit.

- **Governance Privacy** — Implemented: local-first storage, consent-bound exports, audit trails. Partial: formal policy tooling for household administration and large-scale compliance workflows.

These notes are intended to prevent over-interpretation: this work reports a functioning prototype and empirically justified behaviors in the home, while clearly delineating items that remain as future work.

7.5 Trade-offs and Design Constraints

As discussed in Chapter 3, favoring explainable rules over opaque models reduces analytic expressiveness in edge cases but improves interpretability, predictability, and performance on modest hardware. Retained MQTT messages strengthen recovery for reconnecting clients yet pose a risk of stale presentations without prompt reconciliation; DORA mitigates this with prompt identifiers and explicit completions. Voice-first interaction accelerates tasks but can be fragile in noisy conditions; clear recovery prompts mitigate errors at the cost of an extra step. An edge-first posture minimizes data exposure but limits opportunities for cross-household training; bounded-risk learning and carefully governed aggregation, as noted in Chapter 5, merit continued exploration.[6]

7.6 Future Work

Short-term priorities include expanding device adapter coverage, hardening topic schemas and contract tests to prevent semantic drift, and extending accessibility features (multilingual TTS/ASR, focus management, scalable typography). Mid-term efforts should integrate cautiously bounded analytics at the edge (calibrated anomaly detection with abstention), mature consent and policy tooling for residents and caregivers, and enable SMART-on-FHIR pathways for dashboards without centralizing raw data.[3] Long-term directions involve longitudinal, multi-home studies to assess adherence, trust, and well-being outcomes; resilience under overlapping faults and congested RF environments; and steps toward formal verification of critical state machines and regulatory alignment for medical-adjacent use cases. In particular, standardized WHOQOL assessments and real-home field trials are explicitly deferred to future work, pending approvals and resources.

7.7 Closing Remarks

DORA demonstrates that responsiveness, usability, and privacy need not be competing objectives in domestic assistive systems. By unifying robust device integration, decoupled eventing, explainable analytics, and real-time telepresence on an edge node, the platform operates coherently under everyday constraints while preserving data sovereignty.[6] The architecture, methods, and results presented here offer an actionable path for deploying trustworthy, TV-centered assistance for older adults and a foundation for future research at the intersection of IoT, telehealth, and human-centered design.

Bibliography

[1] Cynthia F. Corbett, Pamela J. Wright, Kate Jones, and Michael Parmer. «Voice-Activated Virtual Home Assistant Use and Social Isolation and Loneliness Among Older Adults: Mini Review». In: *FRONTIERS IN PUBLIC HEALTH* 9 (Oct. 2021). DOI: 10.3389/fpubh.2021.742012 (cit. on pp. 1, 5, 7, 9, 12, 14, 32, 37, 50, 52, 54, 56).

[2] Abdelhakim Hannousse and Salima Yahiouche. «Securing microservices and microservice architectures: A systematic mapping study». In: *Computer Science Review* 41 (2021), p. 100415. ISSN: 1574-0137. DOI: 10.1016/j.cosrev.2021.100415. URL: https://www.sciencedirect.com/science/article/pii/S1574013721000551 (cit. on pp. 2, 10, 39, 40, 46, 52, 53).

[3] Franz Frederik Viktor Walter-Tscharf. «A Systematic Integrated Literature Review of Clinical Data Management Platforms of HL7 FHIR». In: *2022 IEEE 24th Int Conf on High Performance Computing Communications; 8th Int Conf on Data Science Systems; 20th Int Conf on Smart City; 8th Int Conf on Dependability in Sensor, Cloud Big Data Systems Application (HPCC/DSS/SmartCity/DependSys)*. 2022, pp. 2298–2307. DOI: 10.1109/HPCC-DSS-SmartCity-DependSys57074.2022.00340 (cit. on pp. 2–4, 7, 9, 10, 34, 36, 43, 50–54, 57, 76, 77).

[4] Naktal Moaid Edan, Ali Al-Sherbaz, and Scott Turner. «Design and evaluation of browser-to-browser video conferencing in WebRTC». In: *2017 Global Information Infrastructure and Networking Symposium (GIIS)*. 2017, pp. 75–78. DOI: 10.1109/GIIS.2017.8169813 (cit. on pp. 3, 9, 11, 12, 38, 49, 51–53).

[5] Ioan Valentin Sita and Bob David. «Development and Implementation of a Comprehensive Smart Home Automation System Using Home Assistant and IoT Devices». In: *2024 4th International Conference on Electrical, Computer, Communications and Mechatronics Engineering (ICECCME)*. 2024, pp. 1–7. DOI: 10.1109/ICECCME62383.2024.10796669 (cit. on pp. 3, 7, 8, 51, 75).

[6] Paige Harvey, Otily Toutsop, Kevin Kornegay, Excel Alale, and Don Reaves. «Security and Privacy of Medical Internet of Things Devices for Smart Homes».

In: *2020 7th International Conference on Internet of Things: Systems, Management and Security (IOTSMS)*. 2020, pp. 1–6. DOI: 10.1109/IOTSMS52051.2020.9340231 (cit. on pp. 3, 6, 7, 10, 40, 46, 50, 52, 53, 55, 77, 78).

[7] Anders Peter Aavild, Aleksander Rosenkrantz de Lasson, Christian Moesgaard Andersen, Erik Christensen, Sergio Moreschini, David Hästbacka, Davide Taibi, and Michele Albano. «Distributed Home Automation with Home Assistant». In: *Proceedings of the 14th International Conference on the Internet of Things*. IoT '24. Association for Computing Machinery, 2025, pp. 206–212. ISBN: 9798400712852. DOI: 10.1145/3703790.3703828. URL: https://doi.org/10.1145/3703790.3703828 (cit. on p. 8).

[8] Sadi Mahmud, Safayet Ahmed, and Kawshik Shikder. «A Smart Home Automation and Metering System using Internet of Things (IoT)». In: *2019 International Conference on Robotics,Electrical and Signal Processing Techniques (ICREST)*. 2019, pp. 451–454. DOI: 10.1109/ICREST.2019.8644232 (cit. on pp. 8–11, 51).

[9] A. J. Brush, Jeannie Albrecht, and Robert Miller. «Smart Homes». In: *IEEE Pervasive Computing* 19.2 (2020), pp. 69–73. DOI: 10.1109/MPRV.2020.2977739 (cit. on p. 9).

[10] S Imran Hussain, S Deepalakshmi, R J Benilla, and V Charu Nivetha. «Automation of Smart Home using Smart Phone via Google Assistant». In: *2023 5th International Conference on Smart Systems and Inventive Technology (ICSSIT)*. 2023, pp. 499–505. DOI: 10.1109/ICSSIT55814.2023.10060979 (cit. on pp. 9, 12).

[11] Amit Kumar Kaiborta and Sarita Samal. «IoT based Voice Assistant for Home Automation». In: *2022 4th International Conference on Smart Systems and Inventive Technology (ICSSIT)*. 2022, pp. 165–172. DOI: 10.1109/ICSSIT53264.2022.9716533 (cit. on pp. 9, 12).

[12] Jay Bhardwaj, Hugo Carducci, Kenneth Clark, Nicholas Kennedy, Dylan Smith, Tim Thai, Emily Francis, and Ramana Reddy. «A Novel, Voice-Activated Smart Home Assistant Using a Large Language Model for Nutrition Assistance». In: *2024 IEEE Global Humanitarian Technology Conference (GHTC)*. 2024, pp. 197–202. DOI: 10.1109/GHTC62424.2024.10771529 (cit. on p. 9).

[13] Jay Bhardwaj, Hugo Carducci, Kenneth Clark, Nicholas Kennedy, Dylan Smith, Tim Thai, Emily Francis, and Ramana Reddy. «A Novel, Voice-Activated Smart Home Assistant Using a Large Language Model for Nutrition Assistance». In: *2024 IEEE GLOBAL HUMANITARIAN TECHNOLOGY CONFERENCE, GHTC*. IEEE Global Humanitarian Technology Conference Proceedings. 14th Global Humanitarian Technology Conference, Radnor, PA,

OCT 23-26, 2024. Villanova University. 2024, pp. 197–202. ISBN: 979-8-3503-7735-4; 979-8-3503-7734-7. DOI: 10.1109/GHTC62424.2024.10771529 (cit. on pp. 9, 12).

[14] Rui Yu, Xiaohua Zhang, and Minyuan Zhang. «Smart Home Security Analysis System Based on The Internet of Things». In: *2021 IEEE 2nd International Conference on Big Data, Artificial Intelligence and Internet of Things Engineering (ICBAIE)*. 2021, pp. 596–599. DOI: 10.1109/ICBAIE52039.2021.9389849 (cit. on p. 9).

[15] Iuliana Chiuchisan, Hariton-Nicolae Costin, and Oana Geman. «Adopting the Internet of Things technologies in health care systems». In: *2014 International Conference and Exposition on Electrical and Power Engineering (EPE)*. 2014, pp. 532–535. DOI: 10.1109/ICEPE.2014.6969965 (cit. on pp. 9, 11).

[16] Mustafa A. Omran, Bashar J. Hamza, and Wasan K. Saad. «The design and fulfillment of a Smart Home (SH) material powered by the IoT using the Blynk app». In: *MATERIALS TODAY-PROCEEDINGS* 60.3, SI (2022). International Conference on Latest Developments in Materials and Manufacturing (ICLDMM), Amristar, INDIA, MAR 24-25, 2022, pp. 1199–1212. ISSN: 2214-7853. DOI: 10.1016/j.matpr.2021.08.038 (cit. on p. 9).

[17] Khalid Ibn Zinnah Apu, Nafiz Mahmud, Firoz Hasan, and Sabbir Hossain Sagar. «P2P video conferencing system based on WebRTC». In: *2017 International Conference on Electrical, Computer and Communication Engineering (ECCE)*. 2017, pp. 557–561. DOI: 10.1109/ECACE.2017.7912968 (cit. on pp. 9, 11, 38, 49, 51–53).

[18] Bo Yu, Chengjie Liu, Wei Wang, and Chao Yang. «A Lightweight Signaling Solution Based on MQTT for Peer to Peer Video Call System». In: *2024 10th International Conference on Computer and Communications (ICCC)*. 2024, pp. 1127–1131. DOI: 10.1109/ICCC62609.2024.10941731 (cit. on pp. 9, 12).

[19] Sanket Salvi, V Geetha, and S Sowmya Kamath. «Jamura: A Conversational Smart Home Assistant Built on Telegram and Google Dialogflow». In: *TENCON 2019 - 2019 IEEE Region 10 Conference (TENCON)*. 2019, pp. 1564–1571. DOI: 10.1109/TENCON.2019.8929316 (cit. on p. 36).