



TECHNICAL DOCUMENTATION

Security & Architecture Overview

The ARC conversational AI kiosk platform — system architecture, the two points at which ARC integrates with an organisation's systems, and the security controls applied at each boundary.

Document reference	ARC-SEC-ARCH-001
Version	1.1
Date of issue	September 2026
Issued by	Amity Robotics
Audience	Organisations evaluating ARC and their security teams

CONTENTS

1	Purpose and scope	3
2	System overview	3
2.1	Component architecture	4
3	Integration model	6
3.1	Integration point 1 — kiosk to customer systems	6
3.2	Integration point 2 — ARC cloud to customer systems	7
4	Network architecture and connectivity	7
5	Security controls	9
5.1	Controls at the kiosk	9
5.2	Controls at the cloud boundary	9
5.3	Transport and encryption	10
5.4	Tenant isolation	10
5.5	Administrative access	10
5.6	Logging and auditability	10
6	Hosting and deployment	11
7	Customer responsibilities	11
8	Security contact	11
A	Appendix A — control summary	12

REVISION HISTORY

VERSION	DATE	SUMMARY
1.0	28 August 2026	First issue for customer architecture and security review.
1.1	1 September 2026	Added the component architecture view (section 2.1) and the network diagram with connection matrix (section 4), for review by customer security and architecture teams.

1 · PURPOSE AND SCOPE

This document describes the architecture of the ARC conversational AI kiosk platform and the security controls that apply to it. It is written to support architecture and security review by an organisation's technical and security teams.

The following are covered:

- The components of the platform and how they relate to one another (section 2).
- The two integration points between ARC and customer systems, and what traverses each (section 3).
- Network connectivity and the direction of every connection (section 4).
- Authentication, transport security, tenant isolation and logging (section 5).
- Hosting arrangements and the responsibilities held by each party (sections 6 and 7).

Configuration options are identified where they exist. Where a control is selectable, the customer's own standard determines which option is applied, and this is agreed in writing before go-live.

2 · SYSTEM OVERVIEW

An ARC deployment consists of three parties. The **ARC kiosk** is the physical device installed on customer premises; it runs the conversational interface and an embedded browser. The **ARC cloud** is the managed platform that performs speech processing, language understanding, knowledge search and answer generation, and holds configuration and analytics. **Customer systems** are the web properties, service APIs and knowledge sources that the customer chooses to expose.




Figure 1 — Platform components and the two integration points.

2.1 Component architecture

Figure 2 opens the three parties of Figure 1 into their components and trust zones. Every request from a kiosk or a console user enters the ARC cloud through a single service edge, which authenticates it before any internal service is reached. Conversation processing is transient: the speech and language services hold data for the active turn only, and anything persisted is written to the tenant-partitioned stores shown in the data plane.




Figure 2 — Component architecture and trust zones. Numbered paths are the two integration points of Figure 1; solid lines stay inside ARC, dashed lines cross into customer-controlled systems.

Single entry, single exit. All inbound traffic to the ARC cloud passes through the service edge and is authenticated there; all traffic toward customer systems leaves through the egress path. No component is reachable from the internet directly, and no component opens a connection into the customer network beyond the endpoints the customer has published.

The components shown in Figure 2, their function, and the data each one handles:

COMPONENT	FUNCTION	DATA HANDLED
Kiosk application	Runs the conversational interface and the embedded browser panel; captures visitor speech and plays synthesised replies.	Audio in transit; its own tenant-scoped device credential. Browser sessions cleared between visitors.
Service edge	Terminates TLS 1.2+ and authenticates every kiosk and console request before any internal service is reached.	Requests in transit; session tokens.
Conversation orchestrator	Coordinates a visitor turn end-to-end across speech recognition, knowledge search, answer generation and speech synthesis.	Transcript and answer for the active conversation.
Speech services	Convert visitor speech to text, and synthesise the platform's spoken replies.	Audio and transcript for the active turn.
Answer generation	Generates the reply from the knowledge retrieved for the tenant. The language model region is selected per organisation (section 6).	Transcript and retrieved knowledge for the active turn.
Knowledge search	Retrieves the tenant's synchronised knowledge relevant to the visitor's question.	The tenant's own knowledge base only.
Knowledge synchronisation	Pulls customer sources on the customer's schedule; read-only, mapped field by field (section 3.2).	Only the fields the customer has mapped.
Configuration and knowledge stores	Hold each organisation's configuration and synchronised knowledge.	Encrypted at rest (AES-256), partitioned per tenant.
Interaction records	Per-message records of what was asked and answered, for audit and reporting (section 5.6).	Stored under opaque session identifiers.
Managed secret store	Holds integration credentials — mutual-TLS material and tokens — with encryption and access logging.	Customer-issued credentials; never in configuration files or code.
Management console	Customer-facing administration: content review, configuration and reporting. Access controls in section 5.5.	Configuration and reporting scoped to the signed-in organisation.
Egress gateway	Optional dedicated NAT gateway giving cloud-to-customer calls a static IPv4 source for firewall allow-listing (section 5.2).	Traffic in transit only.

ARC integrates with customer systems at exactly two points. No other channel exists between the platform and the customer estate. In both cases the connection is initiated by ARC or by the kiosk; the customer is never required to open an inbound path into ARC.

3.1 Integration point 1 — kiosk to customer systems

Where a visitor's request can only be answered by a customer system in the moment — a booking, a case status, a live queue position — the kiosk contacts the customer system directly over its own network connection. **This traffic does not pass through the ARC cloud.**

Mode A — presenting a customer web page

The kiosk opens a customer-owned URL in an embedded browser panel, and the visitor interacts with the customer's own page on the kiosk screen.

- The embedded browser will open only domains that are explicitly allow-listed for the deployment.
- The browser session is cleared between visitors; cookies and local storage are discarded on reset.
- Page content is under customer control and can be changed without any release on the ARC side.

Mode B — calling a customer API

The kiosk issues an HTTPS request to an endpoint published by the customer and presents the result to the visitor.

- Requests are stateless and request/response only. ARC holds no persistent session into the customer network.
- Authentication follows the customer's standard; see section 5.2 for the supported options.
- Credentials are issued per deployment and can be rotated or revoked by the customer unilaterally.

Network boundary. The kiosk operates on a public mobile network and therefore has no route to any customer internal system unless the customer publishes it. That boundary is enforced by the customer's own perimeter, not by ARC configuration.

3.2 Integration point 2 — ARC cloud to customer systems

The ARC cloud reads from customer sources in two modes. In both, ARC initiates the connection outbound to an endpoint the customer publishes.

Mode A — scheduled knowledge synchronisation

Information the kiosk should hold in advance — directories, service catalogues, opening hours, published policies — is synchronised from customer sources on a schedule set by the customer.

- **ARC pulls on a timer** — nightly, hourly or on demand. No inbound path to ARC is required.
- **Read-only.** The synchronisation cannot create, modify or delete anything in customer systems.
- **Field-level mapping.** The customer selects which fields are transferred; unmapped fields are never requested.
- **Review before publication.** Synchronised content is presented in the management console for customer review.

Mode B — live retrieval during a conversation

Where information is only valid at the moment of asking, the ARC cloud queries a customer API during the conversation and incorporates the result into the answer.

- The request carries the visitor's question and the search terms derived from it. The customer defines both the request and the response structure.
- The call runs in parallel with the platform's own knowledge search, under an independent timeout (10 seconds by default, tunable per endpoint).
- **Fails safe.** If the endpoint does not respond in time, the kiosk answers from information it already holds rather than speculating.
- Authentication follows the options in section 5.2.

4 · NETWORK ARCHITECTURE AND CONNECTIVITY

- **The kiosk uses a dedicated cellular connection** (4G/5G) provisioned for the deployment. It does not join the customer's corporate network or wireless infrastructure.
- **All connections are outbound.** No service on the kiosk listens for inbound traffic, and no inbound path into the customer network is required by ARC.
- **No lateral path exists.** Because the kiosk is not attached to the customer LAN, a compromised kiosk yields no foothold on customer internal infrastructure.
- **Transport is HTTPS/TLS in every direction**, including kiosk to cloud, cloud to customer, and kiosk to customer.

Figure 3 places every connection on the network, and the connection matrix that follows states its protocol, direction and initiator. This is the complete set: a connection that does not appear in the matrix does not exist.




Figure 3 — Network zones and every connection in an ARC deployment. Arrows point in the direction the connection is initiated. Dashed lines terminate at the customer's own perimeter.

REF	FROM → TO	PROTOCOL	INITIATED BY	PURPOSE AND CONTROLS
C1	Kiosk → ARC service edge	HTTPS / WSS, TLS 1.2+, port 443, certificate-pinned	Kiosk	Conversation traffic, configuration and telemetry. Tenant-scoped device credential (section 5.1).
C2	Kiosk → customer pages / APIs	HTTPS, TLS 1.2+, port 443	Kiosk	Integration point 1: allow-listed web pages and live look-ups, direct from the kiosk (section 3.1).
C3	ARC cloud → customer endpoints	HTTPS, TLS 1.2+, port 443; static IPv4 source optional	ARC cloud	Integration point 2: knowledge synchronisation and live retrieval. Mutual TLS or secret key / token (section 5.2).
C4	Administrator browser → management console	HTTPS, TLS 1.2+, port 443	Customer staff	Console access with per-user accounts, role-based access control and MFA (section 5.5).
—	Inbound to the kiosk	None. No service on the kiosk listens for inbound connections.		
—	Inbound into the customer network required by ARC	None. Every connection toward the customer terminates at endpoints the customer has published.		

Controls are described by the boundary at which they apply. At the kiosk, the objective is that a device standing in a public space can be trusted. At the cloud boundary, the objective is that every call between the platform and customer systems is verifiably authentic in both directions.

5.1 Controls at the kiosk

Certificate pinning

The kiosk trusts only the certificate pinned into it. A proxy presenting an otherwise-valid certificate is rejected, so kiosk traffic cannot be intercepted on the mobile network it uses.

Tenant-based authentication

Each kiosk authenticates as the customer's tenant and is scoped to that tenant's data only. Credentials are held per device, so an individual kiosk can be revoked without affecting the remainder of the fleet.

Network separation

The kiosk's dedicated cellular connection keeps it off the customer network entirely, as described in section 4.

5.2 Controls at the cloud boundary

Two authentication options are available for calls between the ARC cloud and customer systems. The customer's own standard determines which is used; both may be combined with the network control below.

OPTION	DESCRIPTION
Mutual TLS	Certificates are presented and verified on both sides. The customer verifies ARC, and ARC verifies the customer, so no other party can present itself as the platform.
Secret key / token exchange	A shared secret or bearer token, held in a managed secret store, never written into configuration files or source code, and rotated on the customer's schedule.

Source IP allow-listing

The ARC cloud is built on serverless compute (AWS Lambda), so outbound calls do not originate from a fixed address by default. Where a customer firewall requires a source IP allow-list, ARC provisions a dedicated NAT gateway with a static IPv4 address reserved for that customer's integration.

Revocation is unilateral. Because ARC authenticates using credentials the customer issues, the customer can terminate either integration at any time by revoking those credentials, without coordination with Amity Robotics.

5.3 Transport and encryption

- **In transit.** All traffic between kiosks, the ARC cloud and storage is encrypted with TLS 1.2 or higher.
- **At rest.** Customer data held by the platform is encrypted using AES-256.
- **Secret material** is held in a managed secret store with encryption and access logging, and is not written into configuration files or source code.

5.4 Tenant isolation

Each customer's data is logically isolated within the platform. Access controls and data partitioning are applied such that data belonging to one organisation cannot be read by another. Kiosk credentials are themselves tenant-scoped, as described in section 5.1, so a device can only ever address its own organisation's data.

5.5 Administrative access

- **Authentication** to the management console is provided through AWS Cognito, with per-user accounts and automatic session timeout.
- **Multi-factor authentication is available** and can be enabled for customer administrator accounts.
- **Role-based access control** governs permissions per user role, including feature-level permissions.
- **Organisation separation** is enforced at the console layer as well as in storage.

5.6 Logging and auditability

- **Integration activity is logged.** Each call to a customer endpoint is recorded with its timestamp, target and outcome, and is available to the customer for review.
- **Synchronisation runs are logged** with timestamps and record counts, so the customer can confirm what was transferred and when.
- **Administrative activity in the console is tracked** for security review.
- **Infrastructure-level activity** is captured by the underlying cloud provider's audit logging.

6 · HOSTING AND DEPLOYMENT

The ARC cloud runs on Amazon Web Services. The hosting region is agreed for each deployment. Where a customer requires the platform to operate within a specific jurisdiction for production, ARC can be deployed to a region in that country; the architecture and the controls described in this document are unchanged by the choice of region.

Platform infrastructure is defined as code, so deployment into an additional region is a configuration exercise rather than a redesign.

7 · CUSTOMER RESPONSIBILITIES

The following are held by the customer, and ARC depends on them:

AREA	CUSTOMER RESPONSIBILITY
Endpoint exposure	Publishing any web page or API that ARC is to reach, and applying the customer's own perimeter controls to it.
Authentication method	Selecting mutual TLS or secret key / token exchange, and issuing the corresponding certificates or credentials.
Credential lifecycle	Rotation and revocation of credentials issued to ARC, on the customer's own schedule.
Source data accuracy	The correctness of information exposed for synchronisation or live retrieval.
Content approval	Reviewing and approving synchronised knowledge before it is used in visitor-facing answers.
Console administration	Managing the customer's own console user accounts, roles and multi-factor enrolment.

8 · SECURITY CONTACT

Amity Robotics welcomes procurement security reviews, security questionnaires and vulnerability disclosure. Deployments requiring stricter or additional controls can be discussed directly with the security team.

Security team	security@amityrobotics.com
Public trust information	Trust Center and architecture overview, published on the Amity Robotics website

APPENDIX A · CONTROL SUMMARY

A consolidated view of the controls described in this document, for use as a review checklist.

AREA	CONTROL	APPLIES AT
Transport	TLS 1.2 or higher	Every connection
	Certificate pinning	Kiosk
Authentication	Tenant-based device authentication	Kiosk
	Mutual TLS	Cloud to customer
	Secret key / token exchange	Cloud to customer; kiosk to customer
Network	Outbound connections only	Kiosk and cloud
	Dedicated cellular; no LAN attachment	Kiosk
	Static egress IPv4 via dedicated NAT gateway	Cloud to customer
Data separation	Encryption at rest, AES-256	Cloud
	Logical tenant isolation and data partitioning	Cloud
Access	Per-user console accounts (AWS Cognito)	Management console
	Role-based and feature-level permissions	Management console
	Multi-factor authentication, available	Management console
Assurance	Integration call logging	Kiosk and cloud
	Synchronisation run logging	Cloud
	Customer approval of synchronised content	Management console
Resilience	Independent timeout on live retrieval	Cloud to customer
	Fail-safe answering when an endpoint is unavailable	Kiosk

End of document. Questions arising from this review may be directed to security@amityrobotics.com.