VERSION 1.0 · SEPTEMBER 2026

Voho Security Whitepaper 2026

Voho puts an AI layer on top of systems you already run. It answers calls, reads documents and completes bounded actions inside your core systems. This document explains where the data goes, what the agents are permitted to do, what they are structurally prevented from doing, and what evidence is produced for every interaction.

For security, risk and compliance teams evaluating Voho for regulated operations in the Kingdom of Saudi Arabia and the Gulf.

DEPLOYMENT	TRAINING ON YOUR DATA	CONTROL SYSTEMS
Your servers, a Saudi region, or Voho cloud	Never, without a separate written agreement	Read-only, in every deployment

01

Eight statements this document has to defend.

Everything that follows is detail behind these. Each one is a property of how the platform works, demonstrable in a technical session rather than asserted in a slide.

01

It reads your systems. It does not become one.

Voho is not a data lake and not a system of record. It queries the systems that already hold your data, uses what it needs for that interaction, and writes results back into the system that owns them. There is no requirement to copy an archive, a customer base or a historian into Voho to begin.

02

Identity before disclosure, on every channel.

An agent has nothing about an account, a patient, a case or a tenancy until the caller passes your own verification rules. The bar is set per action: a balance enquiry and a transfer are not the same risk and are not permitted by the same check.

03

Permissions are inherited, never granted.

An agent acts with the entitlements of the person it is serving. Nobody gains access to a document, a record or a screen by asking an agent instead of opening the system themselves, and a contractor sees only what their own scope allows.

04

Anything consequential waits for a person.

Money movement, clinical decisions, eligibility outcomes, disconnections, discretionary exceptions and spend above a configured limit are drafted with their evidence attached and held for a named human. The approval is part of the record, not a step outside it.

05

Operational technology is read-only. Always.

SCADA, DCS, EMS, BMS and safety systems are read. Voho issues no setpoint, no switching instruction, no override and no control action of any kind. Its output at that boundary is a work order, a notification or an escalation for a person.

06

Your data does not train our models.

Call audio, transcripts, documents and records are used to serve the request in front of them. They are not used to train, fine-tune or evaluate models for Voho or any third party. Where a customer wants their own data used to improve their own deployment, that is a separate, written, revocable agreement.

07

Every interaction leaves evidence.

08

Nothing goes live on assertion.

Recording, transcript, identity result, every tool call with its system and payload, model and version used, approvals, refusals and the final outcome, on one record, exportable to your own SIEM.

Agents are evaluated against your own historical calls, cases and events before deployment, including the ones they should refuse or escalate, and monitored against the same measures afterwards.

02

Where the arrows point.

The architecture question that decides a security review is not what the system can do, but which direction data and instructions may cross each boundary.

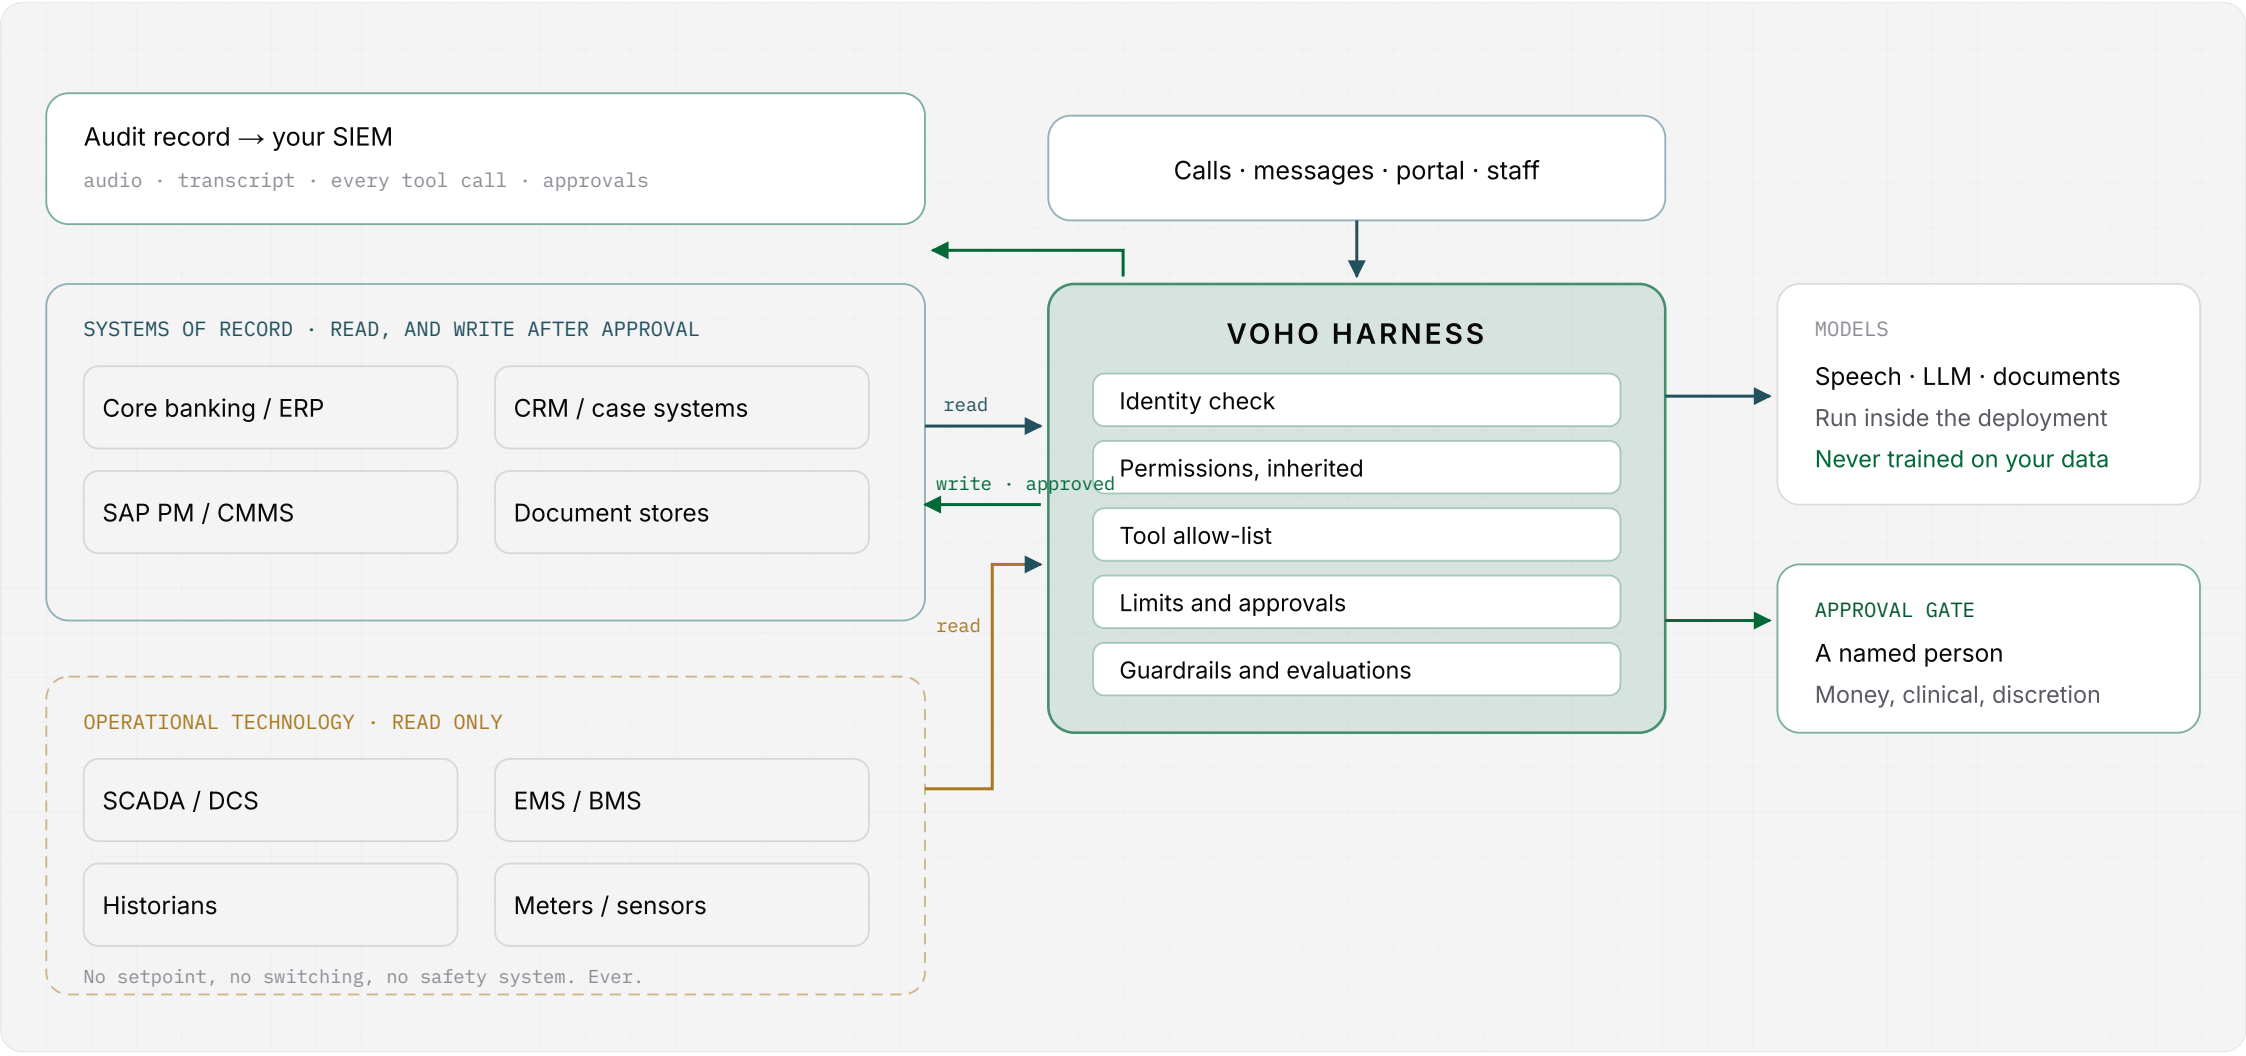

Read-only paths are one-way by construction, not by policy.

03

Three deployment models, and what changes between them.

Residency is usually the first question and the easiest to answer. What matters more is what each model changes about access, egress and update control.

Your own infrastructure

INCLUDING AIR-GAPPED SITES

WHERE MODELS RUN
On your servers, in your data centre

WHERE AUDIO AND TRANSCRIPTS LIVE
Your storage, never leaves your network

OUTBOUND INTERNET
Not required

WHO CAN REACH THE ENVIRONMENT
Your directory, your network controls

UPDATES
Delivered as versioned artefacts, applied on your change process

TYPICAL BUYER
Ministries, OT-connected operations, banks with no-egress policy

Saudi cloud region

IN-KINGDOM, SINGLE TENANT

WHERE MODELS RUN
In-Kingdom region, isolated per customer

WHERE AUDIO AND TRANSCRIPTS LIVE
In-Kingdom storage under your retention policy

OUTBOUND INTERNET
Egress restricted to the integrations you approve

WHO CAN REACH THE ENVIRONMENT
Named Voho engineers under your access process

UPDATES
Managed by Voho on an agreed window

TYPICAL BUYER
Enterprises wanting residency without running the stack

Voho cloud

FASTEST TO START

WHERE MODELS RUN
Voho-managed infrastructure

WHERE AUDIO AND TRANSCRIPTS LIVE
Voho-managed storage under your retention policy

OUTBOUND INTERNET
Standard managed egress

WHO CAN REACH THE ENVIRONMENT
Voho operations team under least privilege

UPDATES
Continuous

TYPICAL BUYER
Pilots, smaller deployments, partners reselling under white-label

04

What data exists, why, and for how long.

Every class of data the platform touches, with the reason it exists. Retention is configuration you own; the defaults are the shortest period your regulator permits.

DATA	WHY IT EXISTS	WHERE IT LIVES	RETENTION	WHO CAN READ IT
Call audio	Recognition during the call; recording where you require one for quality or regulation	Your storage or your region, by deployment	Your retention policy. Default is the shortest period your regulator permits	Named roles you nominate; every access logged
Transcripts	Understanding the request, handover context, quality assurance, audit	Alongside the audio, same boundary	Configurable independently of audio, commonly longer	Role-based; PII redaction available on the stored copy
Tool-call payloads	The reads and writes an agent made in your systems	The audit store inside your boundary	Matched to your audit requirements	Security and audit roles; exportable to your SIEM
Documents	Extraction, classification, answering with citations	Read in place where possible; derived fields stored with the case	Source stays in your document system; derived data follows the case	Inherited from the document system’s own permissions
Operational and network data	Context for an answer: trends, alarms, meter reads, shipment position	Read from the source system, not copied into a second store	Not retained beyond the interaction unless you ask for it	Read-only service accounts, scoped to named tags or objects
Model prompts and outputs	The reasoning behind a recommendation, kept so it can be reviewed	Audit store inside your boundary	Your retention policy	Security, audit and the evaluation team

05

What an agent may never do.

A capability list describes what a system does on a good day. This describes what it will not do on a bad one, which is the list a risk committee is actually looking for.

Operational technology

OIL AND GAS, UTILITIES, FACILITIES, MANUFACTURING

NEVER · Write a setpoint, operate a breaker, override a schedule, touch a safety instrumented system, or issue any control action

INSTEAD · Raise a work order, notify a person, escalate to the control room

Clinical

HEALTHCARE

NEVER · Diagnose, advise, interpret a result, comment on medication, or decide whether a symptom can wait

INSTEAD · Follow your written triage protocol and connect a clinician; red flags stop the booking flow

Money

BANKING, UTILITIES, TELECOM, LOGISTICS CASH ON DELIVERY

NEVER · Move funds, issue credit, waive a fee or adjust a balance outside configured limits

INSTEAD · Draft the action with its evidence and hold it for a named approver

Entitlement and eligibility

GOVERNMENT, HEALTHCARE, INSURANCE

NEVER · Decide eligibility, grant an exemption, or exercise discretion reserved for an officer

INSTEAD · Prepare the case with the rule cited and route it to the officer who decides

Identity and credentials

EVERY DEPLOYMENT

NEVER · Reset, read out, store or request a password, OTP or credential

INSTEAD · Take the request and route it to the process that owns it

Supply and service continuity

UTILITIES, TELECOM

NEVER · Disconnect a customer, bar a line, or suspend a service automatically

INSTEAD · Prepare the action for a person, with vulnerable-customer checks applied first

06

Model governance.

The model is the least controllable part of the system, so the controls sit around it rather than inside it.

Routing, not one model for everything

Speech recognition, synthesis, language models, document models and your own scorecards are separate components chosen per task. Each is versioned, and the version that produced an output is recorded with it.

Grounded answers, cited sources

Answers about an account, an asset or a policy come from a tool call into your systems or a retrieval over your documents, with the record or the paragraph attached. Where there is no source, the agent says so and hands over rather than improvising.

Guardrails at the boundary, not in the prompt alone

What an agent may call, read, write and spend is enforced by the harness around the model. A prompt is not a security control, and is not treated as one.

Evaluation before go-live

Agents are scored against your own history: calls that already happened, cases already closed, incidents already resolved. The measures that matter most are refusal and escalation accuracy, not average handling time.

Continuous evaluation after it

The same suite runs against live traffic samples. Regression on a refusal case is treated as a production incident, not a tuning exercise.

Change control

Prompt, tool, model and policy changes are versioned and released through your change process in on-premise deployments, and on an agreed window in managed ones.

07

Platform and operational security.

How the software is run, reached and observed, in the deployment model you choose.

Access to your environment

Least privilege, named individuals, multi-factor authentication, and access granted through your own process in on-premise and in-Kingdom deployments. Voho engineers hold no standing access to customer environments.

Segregation

Single-tenant by default for enterprise deployments. Where a partner resells under white-label, each of their clients is an isolated workspace with its own agents, data and administrators.

Encryption

In transit with modern TLS on every hop, including telephony media where the carrier supports it. At rest with platform-managed keys, or your own key management where the deployment provides it.

Telephony

Voho connects to the Cisco, Avaya or SIP estate you already run, on your existing numbers and queues. No new numbers are required, and call media can be kept inside your network in on-premise deployments.

Network placement

Deployed in the zone your architecture requires, with OT read paths through the historian or data diode arrangement your OT security team specifies rather than direct control-network access.

Logging and monitoring

Platform and application logs, plus the full interaction record, streamed to your SIEM. Voho does not require log data to leave your environment for the platform to function.

08

Assurance and regulatory position.

Voho is a young company and does not yet hold the attestations a larger vendor would put on this page. Rather than imply otherwise, here is the exact position. The controls described in this document are real and can be demonstrated; the certificates that independently confirm them are in progress, and we will say so until the day they are not.

SOC 2 Type II	Not yet held. Platform built to the Trust Services criteria; readiness work in progress
ISO/IEC 27001	Not yet held. Information security management practices aligned; certification planned
GDPR	Aligned. Data minimisation, retention control, deletion path and processor terms in the agreement
HIPAA	Aligned. PHI handling controls and clinical boundaries as described in this document
EU AI Act	Aligned. Disclosure, human oversight, logging and evaluation practices built in
Saudi PDPL (SDAIA)	Deployment designed to support your obligations; residency, retention and DPIA support provided
NCA Essential Cybersecurity Controls	Control mapping provided for in-Kingdom deployments on request
SAMA Cyber Security Framework	Control mapping provided for financial-sector engagements on request
Independent penetration test	Scope, tester and date of the most recent test supplied under NDA where one has been performed
Sub-processors	None in air-gapped deployments. Managed deployments: list supplied per environment

Voho does not claim an attestation it does not hold. Where a row above is not yet complete, the honest answer is given in the technical session rather than implied here.

09

Sector control summary.

The same architecture, with the boundaries each regulator and each risk committee cares about most.

Banking and financial services

- Authentication bar set per action; nothing readable before it passes
- Transaction limits, with anything above them drafted for an approver
- Card, dispute and beneficiary changes treated as high-risk actions
- Every attempt, successful or refused, on the audit trail

Government

- Decisions, exemptions and eligibility remain with officers
- Procedural answers cite the circular and clause they came from
- In-Kingdom or on-premise deployment, including no-egress entities
- Retention and PII redaction configured to the entity’s own policy

Healthcare

- No diagnosis, advice, medication guidance or result interpretation
- Red-flag symptoms stop the flow and reach a clinician immediately
- Identification before any patient matter is discussed
- Escalation accuracy is the gating measure at go-live

Oil, gas and utilities

- Read-only to SCADA, DCS, EMS and historians; no control actions
- Actions land in SAP PM, CMMS or work management after approval
- Vulnerable and critical customers handled before bulk notification
- Restoration and safety information sourced from the incident record

Telecom

- Higher verification bar for SIM swap, port-out and account changes
- Spend limits on anything chargeable, with approval above them
- Network alarms read; no configuration change to network elements
- Fraud-sensitive actions flagged to the fraud team as they happen

Facilities and logistics

- Read-only to building management systems
- Cost limits on callouts and work orders raised out of hours
- Cash-on-delivery amounts read from the order, never from the call
- Address and recipient changes gated behind verification

10

The questions security teams ask.

Answered the way they would be answered on a call, which is the same way they are answered here.

Where does our data physically sit?

Where your policy requires. On your own servers, in an in-Kingdom region, or in Voho cloud. In air-gapped deployments nothing leaves the building, and the platform is designed so that no function depends on egress.

Is our data used to train your models?

No. Customer audio, transcripts, documents and records are used to serve the interaction in front of them and nothing else. Improvement using a customer's own data, for that customer's own deployment, requires a separate written agreement that can be withdrawn.

What stops an agent doing something catastrophic?

Three things, in order. It can only call the tools it has been given. Those tools enforce limits and permissions outside the model. And anything consequential is drafted and held for a named person. The prompt is not the control.

Can it touch our control systems?

No. Operational technology is read-only in every deployment. There is no code path from an agent to a setpoint, a breaker or a safety system, and the read path itself is placed where your OT security team specifies.

What does a caller hear about being on an AI?

What you configure, and in every deployment the agent tells the truth the moment it is asked. Handover to a person is always available rather than a trap door.

How do we audit a decision months later?

Every interaction has one record: audio, transcript, identity result, each tool call with its payload, the model and version, approvals and refusals, and the outcome. It can be exported to your SIEM as it happens.

What happens when the model is wrong?

Grounded answers carry their source, so an engineer or officer can check them in seconds. Refusal and escalation cases are tested before go-live and monitored after it, and a regression on one is handled as a production incident.

Who at Voho can see our data?

In on-premise and in-Kingdom deployments, nobody by default: access is granted through your process, to named individuals, for a defined period, and is logged. There is no standing access.

What do you need from us to start a security review?

The deployment model you require, your retention and residency policy, the systems to be connected with their owners, and your change and incident processes. A control mapping to NCA ECC or the SAMA framework is

provided where relevant.

Reviewed every six months, or on any material change to the platform. Questions that are not answered here should be brought to a technical session with our engineers, where the platform can be shown doing what this document describes.

VOHO SECURITY WHITEPAPER 2026 · VERSION 1.0 · SEPTEMBER 2026 · VOHO.AI