



PLATFORM · CONVERSATIONAL IVR

Smart IVR

Voice is the front door, not a phone tree.

A carrier-grade, standards-based IVR platform that resolves what it can, understands what it can't, and hands off with full context to a human agent — or a video call — on the same suite.

VoiceXML 2.x

W3C standard

Hundreds

of calls per server

6+

ASR/TTS/NLU vendors

Cloud / on-prem

deployment



THE AUTOMATED VOICE PLATFORM

A carrier-grade IVR that speaks any engine you choose.

Standards-based

VoiceXML 2.0/2.1, MRCPv2 and SIP — open, not proprietary.

Vendor-neutral speech

HTTP and MRCP connectors plug in the ASR, TTS and NLU vendor you already trust.

Deploy anywhere

Bare metal, VMware, KVM, or AWS, Azure and Google Cloud.



Everyone promises AI voice. Most still make you press 1.

Conversational AI is the pitch. In practice, buyers are asked to choose between a rigid legacy tree, a forced move to someone else's cloud, or a rebuild inside a brand-new bot-builder canvas.



Press 1 for menus, forever

Legacy DTMF trees stay rigid because rebuilding them is treated as a separate, risky project.



Cloud-only, take it or leave it

Nuance's on-prem Conversational IVR is being retired by Microsoft — sustaining support ends June 2026.



Bot-builder lock-in

Modernising often means rebuilding the whole IVR inside a new proprietary canvas, from zero.



Smart IVR starts from telephony, not a chatbot canvas.

A standards-based VoiceXML engine with open speech-engine connectors, and — as of August 2026 — an AI Agent Gateway that sits alongside VoiceXML rather than replacing it, so existing applications keep running.

A carrier-grade IVR that speaks any engine you choose

Smart IVR is a software-based, 100% packaged conversational IVR platform, VoiceXML 2.X compliant with extended tags, built on an open-standards core. It runs on your own servers, on a virtual machine, or in the cloud.



Standards-based

VoiceXML 2.0 and 2.1 (W3C),
MRCPv2 (RFC 6787) and SIP (RFC
3261 / 5552) — open, not
proprietary.



Carrier-grade

Multi-tenant architecture built to
process hundreds of
simultaneous calls per server.



Vendor-neutral speech

HTTP and MRCP connectors plug
in the ASR, TTS and NLU vendor
you already trust — not a walled
garden.



Deploy anywhere

Bare metal, VMware, Proxmox
and KVM, or AWS, Azure, Google
Cloud and OVH — your choice.

VoiceXML 2.0/2.1

W3C compliant

MRCPv2

RFC 6787 speech control

SIP

RFC 3261 / 5552

Per port

subscription or licence

What it means for voice to be a real front door

Not a bot bolted in front of the phone system. The telephony layer and the conversation layer, designed to work as one.



Owens telephony end to end

SIP trunking, DID numbering, PBX/CCaaS interconnection, barge-in and transfer — natively, not bolted on.



Hands off with context

Call variables and conversation context carry through to Web ACD's queue and agent desktop at escalation.



Coexists with new AI

The AI Agent Gateway lets a connected AI own conversation while Smart IVR keeps owning telephony — no rebuild.



Vendor-neutral speech

Any MRCP- or HTTP-compatible ASR, TTS or NLU engine plugs in — Nuance, Verbio, Lumenvox, Google, Amazon, Azure.



Escalates to a real video agent

A call that needs a human can route into the same suite's Web ACD queue — voice, or face to face over video.



One telephony layer, one vendor

IVR, ACD and video all run on the same open-standards core, not stitched together from separate SKUs.

Entry points: any carrier or VoIP system over standard SIP trunking — no proprietary handset or app required.




Did you know you can take your company's calls to the next level with Natural Language and AI?

`<vxml>` version="2.x"

Interactive Powers

From touch-tone menus to real conversation, one engine

Directed-dialog and natural-language mode run on the same VoiceXML engine, so a flow can start as a menu and graduate into a conversation without changing platform.



Directed-dialog menus

DTMF and grammar-based prompts, with standard VoiceXML flow control throughout.



Mixed-initiative conversation

Natural-language mode via the native `<nlu>` element lets callers say what they want.



Context carried across turns

The `<data>` element and NLU context grammars keep state through a whole conversation.



Barge-in enabled

Callers can interrupt a prompt at any point — standard VoiceXML behaviour, honoured throughout.



Speech and DTMF verification

Identity or transaction-confirmation steps ahead of a sensitive action, in either input mode.



Recording & transcript capture

NLU sessions can record and retrieve caller speech for QA, training and dispute resolution.

Bring your own ASR, TTS and NLU — or use ours



ASR

Speech recognition, MRCP-compliant

Nuance

Loquendo

Lumenvox

Verbio

Corporate site also lists: Deepgram · Voci · Azure SR



TTS

Text-to-speech, HTTP or MRCP

Flite (bundled)

Mbrola

Cepstral

Loquendo

Acapela

Verbio

Voxygen

Nuance

VoiceRSS

v2.x adds cloud APIs: Google · Amazon · Azure



NLU / Dialog

Native <nlu> element, since v2.2.0

Google GRPC Speech API

Google Dialogflow

AI Agent Gateway (2026) adds: Azure · AWS · Deepgram · Nuance · Verbio

From native NLU to a live voicebot, four steps

The native <nlu> VoiceXML element shipped in v2.2.0 (March 2020) with Google GRPC Speech API support, and was hardened through 2021–2022 with context grammars, pause handling and session cleanup. A documented runbook takes a Dialogflow agent from design to a live trunk.

1



Design the agent

Build a conversational agent in Dialogflow, covering intents, entities and responses.

2



Export & import

Package the agent as a ZIP and hand it to IntPowers to import into the runbook.

3



Grant access

Create a GCP service account with the Dialogflow API Client role and generate a JSON key.

4



Go live on the trunk

Send the Project ID and JSON credentials to IntPowers to wire the agent into a live SIP trunk.

Outbound calling and call control, built in

Smart IVR isn't only inbound. Outbound campaigns and system-level call control run on the same open-standards core.



Outbound Call Files

Native .call files dropped into the outgoing spool drive reminders, notifications and proactive outreach — no separate outbound dialer product required.



Direct call control, no CCXML

VoiceXML applications reach native call-control variables and commands directly — transfer, hang up or branch logic without CCXML.



SIP-standard everywhere

RFC 3261 SIP and the dedicated RFC 5552 "SIP Interface to VoiceXML" keep outbound and inbound calls on the same open protocol as the rest of the suite.

One console for scripts, logs, licences and health



VoiceXML code editor

Create, edit and upload .vxml, .php and .wav files in the browser, with syntax colouring.



CLI console access

SSH access to the command-line console for control and management — routine checks and deeper diagnostics alike.



System reports

CDR search, Call Compare, Monthly Traffic, Daily Load, and Telephony / VoiceXML logs and statistics.



Fail2Ban visibility

Brute-force protection status is surfaced directly in the admin panel, not buried in a log file.



IVR Watchdog

A scheduled job (every 30 minutes by default) restarts failed components automatically.



License management

Channel count, ASR/TTS enablement, licence type and expiry — one config block, one activation flow.

Encrypted media, hardened by default



Encrypted media

SRTCP over the RTP media path, added via the `res_srtp` module in v2.4.0.



Hardened install path

SELinux guidance and Fail2Ban brute-force protection are part of the documented install, not an afterthought.



Brute-force protection

Fail2Ban status is visible in the admin panel, so intrusion attempts don't go unnoticed.



License-gated capacity

Channel count, ASR/TTS enablement and expiry are enforced per licence, in one configuration block.



Self-healing watchdog

The IVR Watchdog restarts failed components on a schedule, keeping the service resilient without a human on call.



Deployment is a privacy control

On-prem and virtualized installs keep call audio and conversation data inside your own perimeter.

On-premises and virtualized installations keep call audio and conversation data inside your own infrastructure — for regulated voice data, often the deciding factor.

The same engine, wherever it has to run

Most conversational-IVR platforms sold today exist in exactly one place: the vendor's cloud. Smart IVR does not force that choice.



Cloud

AWS, Azure, Google Cloud or OVH — install the same package, or take the AWS-hosted and managed option.



On-premises

CentOS Linux on your own hardware, inside your own network perimeter.



Virtualized

VMware, Proxmox, KVM or Xen — fits the virtualization estate you already operate.



Hybrid

Run the IVR at the edge, close to your PBX, while the rest of the suite lives wherever it needs to.

Why it matters right now


Microsoft is retiring Nuance's on-premise Conversational IVR stack — enterprise licence sales stopped in August 2024, and sustaining support for on-prem installs ends in June 2026. Buyers who don't want to be forced onto a public-cloud-only or subscription-only stack need a standards-based landing spot. Smart IVR is built to be one.

Carrier-grade capacity, self-healing by design



Hundreds of calls per server

Multi-tenant, carrier-grade concurrency — exact capacity depends on hardware and codecs configured.



Multi-node architecture

Company positioning (Aug 2026) describes tens to hundreds of concurrent channels across nodes.



Modern engine core

The VoiceXML Engine has been built with C++11 since v2.2.0, for a leaner, faster runtime.



Self-healing watchdog

Failed components are restarted automatically on a schedule — no human has to notice first.



Built for AWS scale-out

Packaged for deployment across AWS as well as private infrastructure and other public clouds.



Standards-based, not black-box

VoiceXML, MRCP and SIP are open standards, which keeps capacity planning and vendor portability honest.

Exact port-per-server figures vary by hardware, codecs and generation — ask your account team for a sizing worksheet against your traffic profile.

A standards-based engine, continuously hardened



Native NLU arrives

v2.2.0 · MAR 2020

The VoiceXML <nlu> element ships with Google GRPC Speech API support — the engine's first native language tag.



Encrypted media & context grammars

v2.4.0 · MAR 2021

SRTP support (res_srtp) and <context> grammars for NLU recognition land in the same release.



PJSIP & voice capture

v2.5.0 · NOV 2021

PJSIP support added alongside recording and retrieval of NLU user voice, for QA and training.



AI Agent Gateway

AUG 2026 · Company positioning

A connected AI Agent can own conversation logic while Smart IVR keeps owning telephony — coexisting with existing VoiceXML apps, no rebuild.

A track record of hardening, not just features

Random garbage-collector crashes fixed by compiling Spidermonkey thread-safe (2019) · a VoiceXML Engine memory leak closed (2020) · NLU session-cleanup and mid-phrase pause handling improved (2021). Ask your account team for the latest roadmap beyond the versioned technical changelog.

Smart IVR is the front door to three more products

Each platform stands alone and speaks standard protocols — but bought together they behave as one system, with no integration project in between.



Smart IVR

AUTOMATION LAYER


Carrier-grade VoiceXML IVR with speech engines, handling the call before a human is needed.



Web ACD

DISTRIBUTION LAYER


Queues, routes and distributes voice and video to the right agent, and runs their desktop.



Video RTC

MEDIA GATEWAY

Hybrid SIP + WebRTC gateway for browsers, mobile devices and kiosks, built for high availability.



Video REC

RECORDING LAYER

HD video recording, storage and archive, with export for analytics and compliance.

Smart IVR also connects to any carrier or PBX over standard SIP trunking — which is how many organisations add automation to a voice estate they've already standardised on.

Cloud Services • hosting, monitoring and managed operation across all four platforms

Standards first, so it fits what you already run



Telephony

- Standard SIP trunking to any carrier or VoIP system
- Direct native call-control access from VoiceXML — no CCXML required
- Outbound calling via native Call Files
- Licence-gated channel capacity, per server



Standards

- VoiceXML 2.0 and 2.1, W3C-compliant with extended tags
- MRCPv2 (RFC 6787) for ASR/TTS resource control
- SIP (RFC 3261) and "SIP Interface to VoiceXML" (RFC 5552)
- ECMA-262 scripting inside VoiceXML applications



Speech & AI

- Vendor-neutral ASR / TTS / NLU connectors (HTTP or MRCP)
- Documented Dialogflow voicebot deployment runbook
- AI Agent Gateway coexistence model, alongside VoiceXML
- Native <nlu> element for in-engine conversation state

Because the platform speaks standard SIP and VoiceXML, Smart IVR can sit in front of an incumbent contact center and take on the automation layer first.


AI Agent Gateway: telephony stays, conversation evolves

Announced August 2026: Smart IVR splits the job in two, so modernising the conversation never means rebuilding the telephony underneath it.



Smart IVR owns telephony

SIP trunking, DID numbering, PBX/CCaaS interconnection, barge-in and transfer stay with the platform.



An AI Agent owns conversation

Business-logic and dialogue decisions are delegated to a connected AI Agent, wired in alongside VoiceXML.



No rip and replace

The gateway coexists with existing VoiceXML applications rather than requiring a rebuild.



Same multi-vendor speech layer

Google, Azure, AWS, Deepgram, Nuance and Verbio connect the same way they always have.



The voice of AI agents

That's the company's own framing, published 1 August 2026: a standards-based telephony layer that a modern conversational AI can plug into immediately, without forcing a rebuild of applications already in production. Certification and capacity claims tied to this announcement should be confirmed with your account team before quoting.

Five ways to resolve things before an agent has to



Self-service menus

Directed-dialog and natural-language flows that resolve routine requests without an agent.



Voicebot deployment

A documented Dialogflow runbook takes a conversational agent from design to a live trunk.



Outbound campaigns

Native Call Files drive reminders, notifications and proactive outreach at scale.



Identity & verification

Speech- and DTMF-based verification steps ahead of a sensitive transaction or transfer.



Escalation to a human

Any flow can hand off — with context — into the same suite's Web ACD queue.

Every one of them ends the same way if it needs to

A caller who starts on a self-service menu, a voicebot or an outbound campaign call is one hand-off away from a queued conversation with a named agent — voice or video — on the same suite, with the context of everything that happened before the transfer. The entry point changes; the operating model does not.

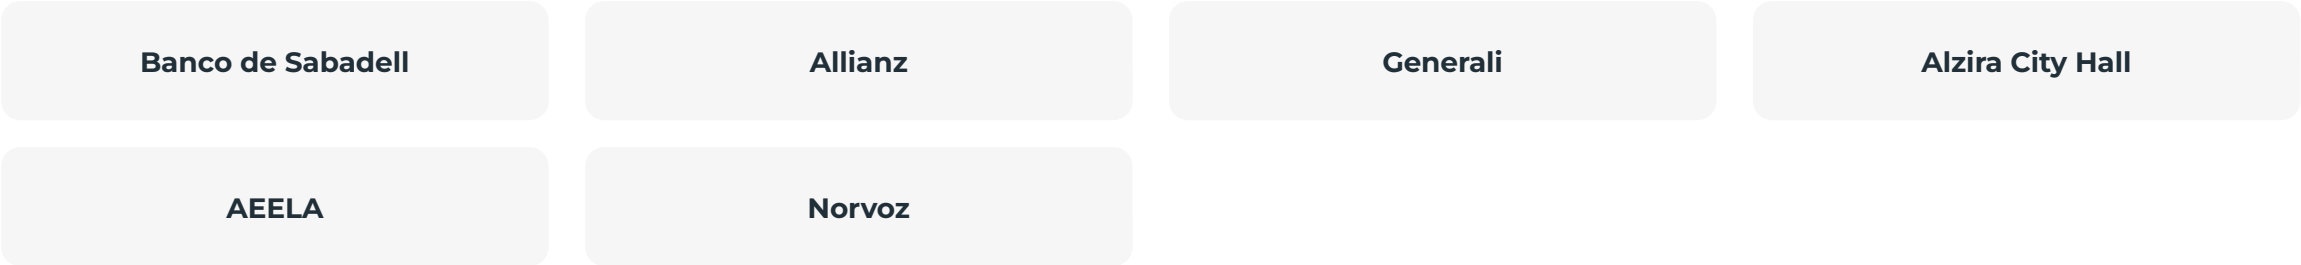
Where Smart IVR parts company with bot-builder platforms

	A typical bot-builder platform	Smart IVR
Speech engine choice	Proprietary or single-vendor bot builder	Vendor-neutral ASR/TTS/NLU via MRCP/HTTP connectors
Deployment	Cloud-only, increasingly enforced	Cloud, on-premises, virtualized or hybrid
AI / voice-bot access	Gated behind the top subscription tier	Bundled — no separate AI tier to unlock
Adoption path	Rebuild inside a new bot-builder canvas	AI Agent Gateway coexists with existing VoiceXML apps
Commercial model	Per agent, per month, plus per-session fees	Subscription or licences
Hand-off to human agents	Separate product or SKU to integrate	Same suite as Web ACD, Video RTC and Video REC

"A typical bot-builder platform" summarises patterns published on the capability and pricing pages of Five9 IVA Studio, Genesys Cloud CX, NICE CXone, Nuance/Microsoft, Twilio Voice and Vonage, reviewed August 2026. Specific vendor comparisons available on request.

Running where the call actually matters

Smart IVR automates phone conversations for banks, insurers, public agencies and service providers across Latin America and Europe — organisations for whom a call is a case file, not a ticket.



Named clients by industry (published case-study base)



Client counts by industry per [interactivepowers.com/en/success-cases](https://www.interactivepowers.com/en/success-cases). G2-rated (Smart IVR listing). Quantified case-study results and named references available under NDA through your account team.

Port-based subscription, with a licence available

Smart IVR is contracted by subscription, per port, on any deployment model — cloud, on-premises or hybrid. For the public sector and regulated environments that need to own the licence, we still offer a perpetual per-port purchase, with the same annual maintenance.



Per-port subscription

Extensions are licensed per server, quoted on release of the binary package — not a public per-seat meter.



Annual maintenance

Covers new versions, updates and support — the upgrade path is included, not re-sold.



Full-function trial

Complete functionality, capped at one port, for evaluation before you commit to capacity.



Quote-based sales motion

A conversation with sales rather than a self-serve checkout — in exchange for a predictable long-term cost.

What this means over a five-year horizon

Port-based subscription means you pay for simultaneous channel capacity, not a named licence tied to each agent. Cost tracks actual business usage instead of a headcount roster — with AI included, not upsold. And where your organization needs to own the licence, a perpetual per-port purchase is still available, with the same annual maintenance.

SMART IVR • CONVERSATIONAL IVR

Give your callers a real front door.

Let us show you a call resolved, escalated and handed off with context — on your own use case, with a full-function trial and no bot-builder rebuild required.

[Talk to Sales](#)

[Request a Demo](#)



Int**Powers**

interactivepowers.com

