

Ovum Decision Matrix: Selecting an Intelligent Virtual Assistant Solution, 2020–21

Publication Date: 17 Dec 2019 | Product code: INT002-000264

Michael Azoff



Summary

Catalyst

In the past half-decade, the intelligent virtual assistant (IVA) (Ovum uses the term interchangeably with digital assistant, automated virtual assistant, virtual digital assistant, intelligent chatbot, and intelligent agent) market has become transformed by the impact of artificial intelligence (AI). There is no leading player in the market that doesn't have some AI-related capability driving its solution. What is also evident is that the major high-tech players with AI interests see the IVA market as an opportunity to apply their technology, while incumbents with decades of experience in the field have embraced AI and used their domain expertise to stay ahead in the market. This Ovum Decision Matrix (ODM) focuses on the leading IVA solutions targeting the enterprise market.

Ovum view

The changing face of virtual assistants and the impact of AI

A key driver for adopting IVAs is the loss that many businesses are facing through poor customer service. The drive toward cost savings using contact center services with human agents that read prepared scripts and have limited freedom to act beyond this might be one reason for this. Another is first-generation virtual assistants that have serious limitations. Whatever the cause, end users (we use this term collectively for consumers, businesses, or employees initializing a conversation) vote with their feet toward the better experiences.

The legacy virtual assistants on the market are not intelligent, have single intents, and cannot conduct a natural conversation that requires retaining context and navigating intents. Unless specified, a conversation can be either text- or voice-based. The best IVAs today have a more natural, human-like conversational ability, can handle multiple intents, and can achieve customer satisfaction scores on a par with the best human agents. They can perform at this level because AI has made it possible to track conversational flow and to detect and manage multiple intents, so that the end user has as near as possible a natural human-like interaction.

There are, however, limits to this capability. IVAs can still struggle when confronted with very long and verbose sentences and unclear vocal utterances (but so can humans). Another feature of this market is that the state-of-the-art is improving year by year, so we expect IVAs to overcome some of their current limitations. This is on the back of a large research effort taking place in conversational AI in both academia and the commercial sector.

This begs the question of whether end users understand they are talking with a robot and how this changes their behavior. Organizations that deploy IVAs do make clear the person is interacting with a virtual assistant, but if a human-like image of an IVA is depicted, some end users can forget. This has both advantages and disadvantages. Speaking with the customers of the vendors assessed here, some of their end users have been found to be curt and less polite when interacting with IVAs and because the machine needs as much information as possible to ascertain the correct intent, this lack of verbosity can be a hindrance. On the other hand, too much verbosity can stress the natural language understanding (NLU) part of natural language processing (NLP) to the limit.

Despite the challenges, all the customers of this technology we spoke to said IVAs make a massive difference to customer engagements, and the organizations that have introduced this technology as a first in their industry sector end up being copied by their peers, which explains the speed with which IVA adoption is spreading.

The effectiveness of IVAs

Benchmarking the performance of IVAs by comparing the call metrics of different vendors and their customers is not possible. This is due to customers having different requirements in the given domain being assessed. For example, the customer will build the conversation flow models and decide at which point the IVA is to hand over to a human agent (called deflection). This point will vary case by case and the information knowledge base that the IVA uses to answer questions will vary from case to case. An IVA vendor customer example with 85% of all calls handled automatically is not therefore necessarily performing better than another of its customers with only 25% of calls handled automatically. Instead it's a question of what the customer wanted, so, for example, a bank would want certain call operations only handled by a human agent. But clearly, the more calls handled automatically, the greater the cost savings for the customer.

IVAs handle multiple calls simultaneously and can easily scale up or down depending on call density so are perfect in running campaigns and other irregular, burst-type activities. They are also perfect for fielding the bulk of calls that are not typically demanding, such as those to IT support for resetting passwords, allowing human agents to handle more demanding and satisfying work.

IVAs are also useful where a 24x7 call center operation is not yet in place but is desired, as was the case with one vendor customer we spoke to. Rather than go down the costly route of extending the hours of their existing human call center, they brought in IVAs. If the IVA hits a difficulty, the caller is requested to call back to the manned call center during opening hours. This resulted in 5,000 net new callers per month with a 72% automated call resolution rating (90% when including the request to call back).

We spoke to a major global IT services company that uses many different IVAs in its customer engagements. Its verdict on IVA technology is that on average it can fully and automatically handle some 20% to 25% of calls within the desired process without deflection to a human agent. There are examples where this could be as high as 90%, depending on the required task. For enterprises handling millions of contact center calls, even 75% to 80% deflection therefore amounts to a significant saving. The IT services company also said it preferred to work with a smaller and nimbler IVA vendor. No doubt the smaller IVA vendors are eager to please and perhaps a lesson for the larger vendors is that they need to think as a small player with their customer engagement approach.

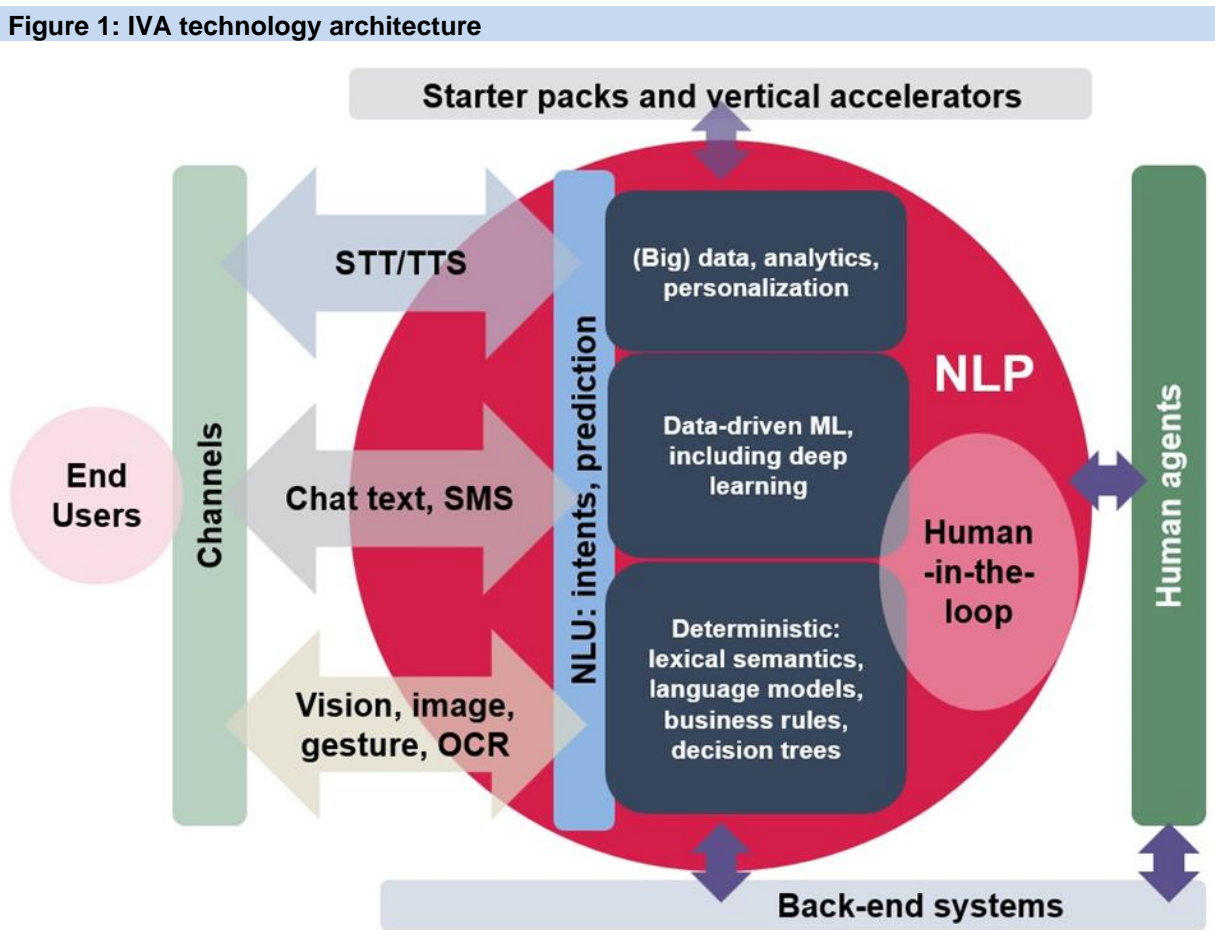
A good balance between machines and humans in call center task allocations is to have IVAs automatically handle high-volume, low-value calls, and deflect to human agents for low-volume, high-value calls.

Technology view

To understand how the market splits between players, it is helpful to look at Figure 1, which provides a high-level IVA technology architecture diagram. End users connect through various channels and are processed by NLP technology that includes the key component NLU. This is driven by various approaches that typically combines data-driven ML models with deterministic systems. The IVA will deflect to human agents as designed in the conversation journey or if it faces problems.

Some vendors add a core human element, represented as human-in-the-loop in Figure 1, as distinct from the human agents who are business application-oriented. The human-in-the-loop can perform clarification functions, almost as a cog in the machine, when the robot part of the IVA is facing difficulties, such as when the confidence level is below a threshold. The human helps and then passes control back to the robot or to a qualified human agent as necessary. Some vendors have created large scale human-in-the-loop operations to support their IVAs.

Each player in the market is approaching the IVA challenge in a different way, placing different emphasis on the various components in Figure 1, and bringing in to play varying degrees of AI in terms of where the AI is applied (there is potential throughout the whole stack).



Source: Ovum

Enterprise focus

This report focuses on enterprise solutions as opposed to the consumer market. Large enterprises with call centers servicing millions of end users around the globe, covering many languages, will need to have IVAs that integrate smoothly with back-end systems, and will want IVA vendors that offer an enterprise-grade level of customer support. Enterprises also need to be able scale IVAs in and out effectively. They need to have high levels of security and be compliant with data privacy regulations (and all the variations in the global jurisdictions they operate in). Not least, they need to work with IVA vendors that understand their business domain and can help them create workable solutions. While

there are a lot of open source tools in the market and many new AI-focused startups with segments of the IVA technology stack, Ovum's selection of vendors here all offer enterprise-focused IVA products.

Key findings

- The IVA allows end users to interact in a natural way rather than conform to the limitations of a legacy (non-intelligent) virtual assistant.
- For voice interactions, the leading IVAs are vastly better in capturing speech first time than legacy solutions which frequently request utterances to be repeated.
- In the early days of chatbots, when something went wrong, they had no graceful exit leaving end users in a continuous loop. Today's IVA technology has learned from this and has various strategies to avoid these scenarios.
- IVA system training is essential because each customer use case is unique, having its own business jargon, as well as having to contend with the natural language vagaries of normal conversation that may be specific to the local region in which it is active.
- Sentiment analysis is mostly of interest to detect negative end-user conversation, so while the art is still in a fledgling state, some vendors are better in bringing in additional context to detect emotion or sentiment.
- Vendors in the customer experience space have built into their IVA features reflecting the best human agent commercial acumen, such as for example, cross-sell and upsell, which makes these IVAs high performers from the start.
- Customers of IVA technology find end-user satisfaction increases and their human agents can focus on higher priority or more demanding queries. In addition, cost savings can be made.
- In hearing the vendors' customer use cases in researching this report it was typical to hear of million-dollar range cost savings by introducing IVAs (calculated as cost per call), bringing ROI within the first year.
- While established in markets including customer care, technical support, and banking, some vendors are also targeting a wider spectrum of use cases, such as vision applications to extend IVAs to product or parts recognition with a host of applications from marketing to maintenance.
- The market is roughly divided into pure-play IVA vendors (with many years' experience in the virtual assistant industry) and relative newcomers, high-tech giants leveraging their AI expertise (some through acquisitions) into the IVA domain. Rationalization has been going on for some time and will continue.

Market and solution analysis

Ovum Decision Matrix: Intelligent Virtual Assistant, 2020–21

The market splits into two types of player: dedicated IVA pure-plays and large technology players with investment in AI entering the IVA space. In this report, in the former group includes Artificial Solutions, Creative Virtual, IPsoft, Interactions, and Verint. The latter group includes IBM, Oracle, Salesforce, and SAP. The latter group is also a relative newcomer to IVAs whereas many in the former group

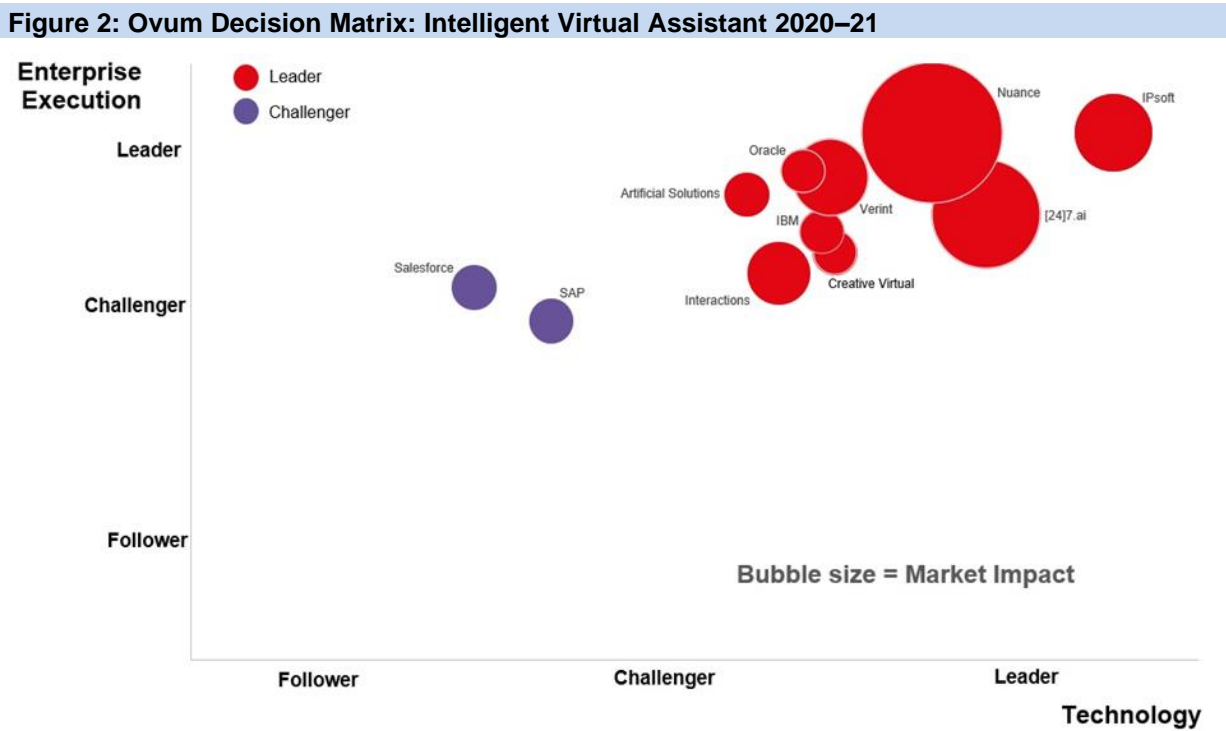
have a long history in the virtual assistant space with deterministic NLP models and have combined these with AI as the field has improved and made an impact in the space. The former group can also boast possession of many years of data in the field, and a lot of data is a must-have for training AI applications.

Pure-play vendor [24]7.ai, which launched in 2000, formed an alliance agreement with Microsoft in 2012 that involved ongoing technology transfer, mostly for voice technology, with Microsoft taking an equity stake in it, so the company sits in the middle of this dual grouping of the market.

While several vendors here have an everything-all-in IVA package, there are aspects to this technology that lend themselves to modularization, such as, for example, speech to text (STT) and text to speech (TTS). This enables IVA customers to opt for a best-of-breed approach, exploiting an IVA vendor for its NLU, NLP, and AI technology, and then perhaps using another solution for the TTS/STT function.


Automatic speech recognition (ASR) and natural language recognition (NLR) are terms that are also used for STT. Similarly, somewhat related to TTS is interactive voice response (IVR). Unrelated (and not covered here) is natural language generation (NLG) technology that takes structured data and converts it into natural language.

The vendors included in this report do not represent an exhaustive coverage of the market but a selection that covers the leading vendors. Figure 2 shows the results of our evaluation, the ODM on IVAs, and an expanded view is shown in Figure 3. We have ranked the vendors accordingly in Table 1 (see Appendix for further details about the evaluation process).



Source: Ovum

Figure 3: Expanded view of Ovum Decision Matrix: Intelligent Virtual Assistant 2020–21



Source: Ovum

Table 1: Ovum Decision Matrix: Intelligent Virtual Assistant 2020–21

Market leaders	Market challengers	Market followers
[24]7.ai	Salesforce	
Artificial Solutions	SAP	
Creative Virtual		
IBM		
Interactions		
IPsoft		
Nuance		
Oracle		
Verint		

Source: Ovum

Vendor solutions

Rankings

In this ODM we rated as leaders the pure-play vendors: Artificial Solutions, Creative Virtual, Interactions, IPsoft, Nuance, Verint, and [24]7.ai (partly owned by Microsoft). Interactions and [24]7.ai have significant investment in human-in-the-loop NLP augmentation. In the case of Interactions, these humans are trained to clarify utterances flagged as unclear by the IVA robot. This is their sole function and they are not otherwise involved in the conversation, passing control back to the robot, or, if they are also unclear about the utterance, to the human agents. [24]7.ai uses human-in-the-loop to handle chat and SMS when the robot is having difficulties with the query. The other vendors can also introduce human-in-the-loop, but these two vendors have large-scale operations in place to provide the manpower.

The high-tech vendors, IBM and Oracle, also rated as leaders, have seized the opportunity to enter the IVA market on the back of their AI expertise. IBM has for many years developed its Watson cognitive technology in-house and applied it to conversational AI. Its solution went live in 2018 with well-known brands as customers. Similarly, Oracle has for many years been quietly developing its AI technology in-house and only went public with its IVA solution in 2019, also with many well-known brands as customers.

We rated Salesforce and SAP as challengers, and we expect them to continue to improve. Both have relatively new solutions on the market, have a large customer base for their other core solutions that generate a huge amount of data, and are well placed to have an impact with their IVA technologies.

Technology

AI has had a huge impact on the virtual assistant market, and today no vendor can compete without moving to the IVA. Of course, there are still many legacy virtual assistant systems in place, and anyone encountering them often must endure a poor customer service as a result. The conversational AI startup space is highly active with venture capital investment, with many startups focusing on segments of the IVA stack rather than end-to-end provision. Some large enterprises with research resources have picked up open source tools and tried to build their own IVA from these pieces, but the net result typically falls short of desired performance. The missing factor is knowledge of the conversational journey: how to construct conversational flows that are relevant for the business as well as how to combine the right mix of data-driven ML with deterministic methods and models, augmented by human-in-the-loop when necessary. Vendors with many years' experience in the field have amassed a huge amount of data that they can use in starter packs and vertical accelerators. The vendors in this report all understand the complexities of an IVA and how to design and build conversation flows.

Some of the vendors have decided not to offer STT and TTS, including leaders Artificial Solutions, Creative Virtual, and Verint, and challengers Salesforce and SAP. These are specialized technologies and these vendors have deferred to a best-of-breed approach. The other vendors offer a complete end-to-end solution, and IVA customers must decide for themselves in proof-of-concept trials whether they need to mix and match or buy one end-to-end solution.

Enterprise execution

The leaders all have exceptional NLP capabilities. What might attract a customer to one versus another could be the ability of a vendor to reach that customer (some of the leaders have more global offices and partners than others). Having good relationships with SIs and specialist consultancies is also essential to be visible in the market. Some of the vendors have excelled in the breadth of business starters and vertical accelerators they offer.

Market impact

Our market impact metric, the ODM chart bubble size, is based on revenue earnings (in a few cases our own estimates). We have created a bubble size metric that is normalized to relative earnings, based on the largest player in the market, and used a floor cut-off to represent the lower end of the scale.

Vendors not in this report

There are two high-tech players in the IVA space who have not participated in this report but are also pressing at the heels of the others covered here: Amazon with Lex and Polly, and Google with Dialogflow. Both Amazon Lex and Google Dialogflow came up in conversations with IVA customers, but neither provides a full end-to-end IVA package. Missing is designing the conversation journey and managing the whole lifecycle of an IVA. The IVA field is continually evolving and improving, and these vendors and others from the VC space could emerge with an end-to-end IVA offering.

We have not examined IVA solutions pre-built for specific industries, but worth mentioning as a good example is Symphony RetailAI. It has developed a turnkey ready-built solution for the retail and consumer packaged goods (CPG) industry, comprising the IVA called Conversational Insights and Decision Engine (CINDE), and Eureka, a comprehensive data platform that feeds CINDE. The company is beginning to explore creating similar turnkey IVA solutions in other industries.

Vendor analysis

[24]7.ai (Ovum recommendation: Leader)

Product: [24]7 AIVA

Availability: Cloud SaaS or managed service

Introduction

[24]7.ai is a private company based in Silicon Valley and has R&D centers around the globe. It offers an IVA called [24]7 AIVA, an AI-powered virtual agent for voice, text, and visual interactions across customer channels, including web, mobile, messaging, advertisements, and phone calls. [24]7.ai includes contact center agents in the loop to improve AI models and enhance the customer experience. These agents can be the company's own agents or provided by [24]7.ai as an outsourced service. The company has been running since 2000 and has developed expertise in virtual assistants by building AI models trained on large datasets of human agent conversations. [24]7.ai holds more than 150 patents and patent applications, including several specifically related to artificial intelligence and customer experience. It has an understanding of communications between consumers and businesses, and its customer journey analytics provides insights into intents and journeys.

In 2012 [24]7.ai formed an alliance with Microsoft. This included Microsoft merging its interactive self-service assets (clients, people, and technologies) into [24]7.ai, as well as an R&D partnership, long-term IP licensing, and Microsoft taking an equity stake in [24]7.ai. AIVA's TTS and STT interface is based on Microsoft's deep neural networks research in this space.

[24]7.ai has also formed an agreement with IBM whereby clients can plug IBM Watson into [24]7.ai, so that the Watson NLP can be used within [24]7 AIVA if desired.

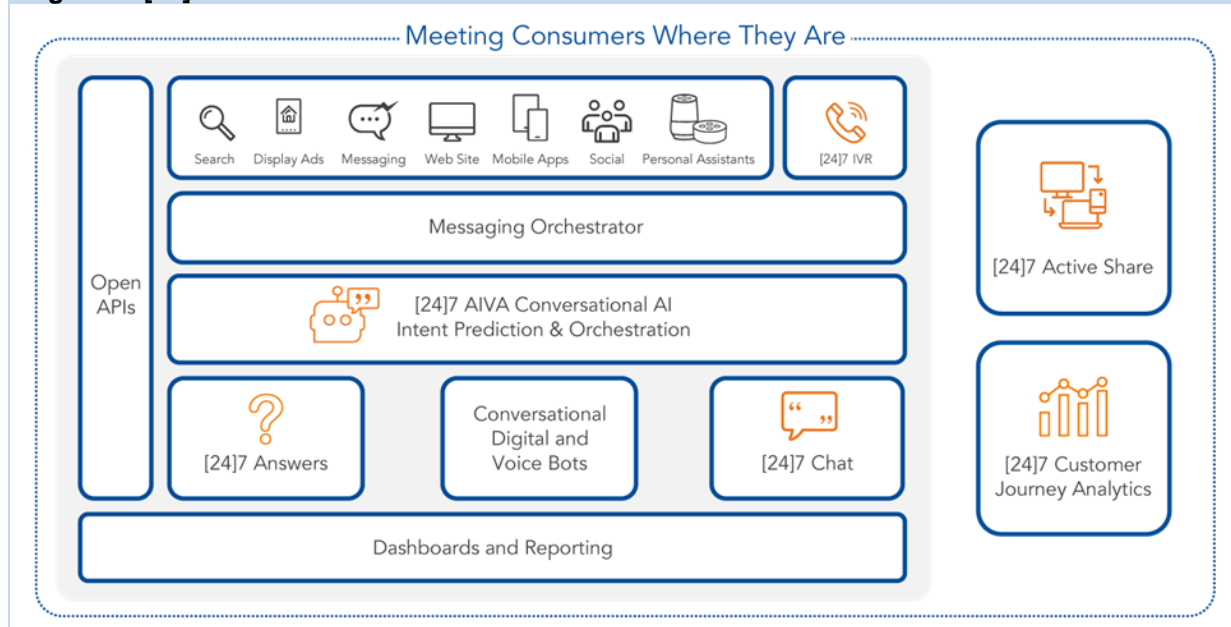
[24]7 AIVA blends AI with live agents, where the human agents can be provided by the client or [24]7.ai. The advantage of the [24]7.ai human agents is that they are specialized on digital channels (live chat and messaging) and trained to work with AIVA, reducing average handle time and boosting efficiency. According to [24]7.ai, typical call centers are focused on voice communications and lack the capabilities to manage chat and messaging communications. [24]7.ai provides the technology and know-how to run agent operations for synchronous and asynchronous digital channels, from agent recruiting to development to performance management. In addition, [24]7.ai has professional services that help a client to roll out AIVA.

Some of the advantages of AIVA include:

- Unified voice and digital: A single platform for connecting voice and digital channels.
- Bot-agent collaboration: The blend of AI and human agents (agents helping bots, bots helping agents) to improve customer experience.
- Deep intent prediction: AIVA goes beyond NLP to use all data signals (behavior, transaction, preferences) to anticipate and understand customer intent.
- Vertical domain expertise: [24]7.ai has predefined, extensible intent ontologies and flows for major verticals.
- Journey management: AIVA maintains context across time, devices, and channels, allowing customers to pause and continue conversations on their own terms.
- Development and management tools:[24]7.ai has design tools to take its clients through the development cycle when implementing a conversational solution.

According to [24]7.ai, the performance to expect from AIVA will depend on the tasks a client decides to automate. Informational requests can be fully automated and represent some 10% to 20% of typical call center workloads. The remainder are transactional requests that require execution of business rules and integration with back-end systems. [24]7.ai recommends a phased rollout that starts with the more common transactions (leading to 50% to 60% automation), followed by the less common transactions (increasing automation to 70% to 80%). Depending on the complexity of transactions and the availability of data, some clients can reach 90% of full automation.

Figure 4: [24]7 AIVA architecture



Source: [24]7.ai

Ovum strength and weakness assessment

Strengths

- **Strong intent analysis:** [24]7.ai uses the conversation source plus external signals, such as web history, CRM data, and the [24]7 big data platform, to understand a user's intent. The company says to understand true customer intent, an IVA must go beyond intent classification by NLP. NLP classification relies only on the words or expressions used in the customer's query. [24]7.ai adds intent prediction based on non-language data, including behavioral data, transactional data, user profile data, and environmental data to determine the true intent.
- **Outcome-based pricing:** [24]7.ai uses an outcome-based pricing model, paying only for results as agreed with the customer when specific performance targets are met or exceeded. Possible metrics include resolved chats and improved NPS and CSAT scores.
- **Offers STT and TTS:** [24]7.ai has its own STT/TTS solution through a privileged technology relationship with Microsoft, giving it direct access to its complete stack of speech recognition technology. This includes advanced deep neural networks algorithms, providing [24]7.ai with best-in-class speech recognition.

Weaknesses

- **Image and facial recognition:** In line with most of the vendors in this ODM, there is no image or facial recognition offered. However, for a leader in IVAs we believe it will be necessary in future. Passing documents and images between end users and IVAs will become normal, as it is with human agents where it is possible to send such materials by email or via cloud drive folders. There are also applications in fraud detection.
- **Safety-critical use cases:** [24]7.ai was one of four IVA vendors in this research not able to say its technology could be used in safety-critical scenarios where human life could be at risk. An example would be fielding end user calls to an emergency service.

Appendix

Vendor solution selection

Inclusion criteria

This ODM required vendors to offer an IVA platform targeting enterprises wanting to build IVA solutions, either for internal or external facing (B2C or B2B) applications. We invited the leading players to participate, able to demonstrate end-to-end conversational journey technology features. Our research has shown that AI has had a huge impact on improving the performance of virtual assistants and hence we focused on players that used AI to power their platforms.

Exclusion criteria

In general, the ODM is not designed to exhaustively cover all the players in a market but a representative set of the leading players. Ovum also invites smaller, possibly niche, vendors that have innovative solutions and are on a fast growth path. With this flexibility we consider each participant on its merits as a good fit to the ODM topic. This ODM does exclude vendor solutions that are complete out-of-the-box IVA solutions for specific use cases. The project also excludes solutions not targeting enterprises, such as consumer market virtual assistants.

Methodology

- Vendors complete a comprehensive capability questionnaire in a spreadsheet, covering the three dimensions of the ODM. The resultant matrix of responses is appropriately scored (not all questions score, for example some are purely informational, and some have higher weighting), and run through our ODM algorithm that scores and ranks the solutions.
- There is a series of comprehensive, structured briefings with all key vendors, including a demonstration where possible.
- We requested customer references and followed up with customer conversations on a selection.
- Supplemental information is obtained from vendor literature and websites and from the results of Ovum surveys.
- The report is peer reviewed by Ovum analysts.

Definition of the ODM

The ODM spans three assessment dimensions.

Technology

In this assessment dimension, Ovum analysts develop a series of features and functionality that provide differentiation between the leading solutions in the marketplace.

Market execution

In this dimension, Ovum analysts review the capability of the solution around key areas such as: enterprise fit, business vision and innovation, interoperability, language support, product roadmap, business starters, and partner ecosystem.

Market impact

The global market impact of a solution is assessed in this dimension. Market impact is a metric normalized to the market leader and is mainly based on the solution's global revenue.

Ovum ratings

- **Market leader.** This category represents the leading solutions that Ovum believes are worthy of a place on most technology selection short lists. The vendor has established a commanding market position with a product that is widely accepted as best of breed.
- **Market challenger.** The solutions in this category have a good market positioning, and the vendors are selling and marketing the product well. The products offer competitive functionality and good price-performance proposition and should be considered as part of the technology selection.
- **Market follower.** Solutions in this category are typically aimed at meeting the requirements of a narrower range of customers. As tier-1 offerings, they should be explored as part of the technology selection.

Author

Michael Azoff, Distinguished Analyst

michael.azoff@ovum.com

Ovum Consulting

We hope that this analysis will help you make informed and imaginative business decisions. If you have further requirements, Ovum's consulting team may be able to help you. For more information about Ovum's consulting capabilities, please contact us directly at consulting@ovum.com.

Copyright notice and disclaimer

The contents of this product are protected by international copyright laws, database rights and other intellectual property rights. The owner of these rights is Informa Telecoms and Media Limited, our affiliates or other third party licensors. All product and company names and logos contained within or appearing on this product are the trademarks, service marks or trading names of their respective owners, including Informa Telecoms and Media Limited. This product may not be copied, reproduced, distributed or transmitted in any form or by any means without the prior permission of Informa Telecoms and Media Limited.

Whilst reasonable efforts have been made to ensure that the information and content of this product was correct as at the date of first publication, neither Informa Telecoms and Media Limited nor any person engaged or employed by Informa Telecoms and Media Limited accepts any liability for any errors, omissions or other inaccuracies. Readers should independently verify any facts and figures as no liability can be accepted in this regard – readers assume full responsibility and risk accordingly for their use of such information and content.

Any views and/or opinions expressed in this product by individual authors or contributors are their personal views and/or opinions and do not necessarily reflect the views and/or opinions of Informa Telecoms and Media Limited.

CONTACT US

ovum.informa.com

askananalyst@ovum.com

INTERNATIONAL OFFICES

Beijing

Boston

Chicago

Dubai

Hong Kong

Hyderabad

Johannesburg

London

Melbourne

New York

Paris

San Francisco

Sao Paulo

Shanghai

Singapore

Sydney

Tokyo

