

MyWave

Virtual Agent

Technical

Recommendations

May

2018

myWave

Commercial in confidence
©MyWave 2018

Contents

Executive Summary	4
Introduction	6
Virtual Agent vs Chatbot	6
Process Recommendations	8
1. Define and design the VA experience	8
2. Explore and identify scenario(s)	9
3. Feasibility study and cost benefit analysis	9
4. Implementation	9
Natural Language Processing	11
Customer Knowledge Model	11
Customer Centric Knowledge Model	11
Benefits of Customer Centric Knowledge Model	12
Voice Application Considerations	12
Delivery Channels	12
Voice Assistant	12
Mobile App	13
Web Application	13
Conversation considerations	13
Target audience	14
Supported languages	14
Legal / regulatory requirements	14

MyWave **Virtual Agent Technical Recommendations**

Checklist	14
Recommendation	15
Handover to Live Agents	15
Live Agent channels	15
Live Chat	15
Phone call	16
Customer Tolerance	16
Understanding a User's request	16
Stuck in Conversation	16
Customer service centre system integration	16
UX Considerations	17
Checklist	17
EA Back End Integration	17
Customer Data Synchronisation	17
Customer Authentication	17
Back End API Gateway	17
Cost & Time	17
EA "Vault" System	18
Security and Privacy	18

Executive Summary

The MyWave Virtual Agent Technology Platform is a revolutionary tool that can quickly and easily be used to transform the current customer experience. The Virtual Agent can be used to automate and simplify current processes, when integrated with the right back-end systems, saving time and cost while greatly improving the customer experience and outcomes. When the Virtual Agent hands customers over to a live agent, the Virtual Agent provides context and continuity, so the customer never starts from zero again. The Virtual Agent can also be used to offload traffic from the Live Chat customer service channel, the most regular requests can be automated through the Virtual Agent and the exceptions can be handed to a Live Chat Agent or Call Centre. The Live Chat platform would be integrated into the Virtual Agent to provide a seamless hand-over and experience. Backend system integration may increase the implementation cost but the savings on operational costs and improved customer satisfaction create a strong justification and business case.

When creating the next use cases for EA's Virtual Agent we recommend an iterative and incremental implementation approach with continuous learning from customer feedback. We recommend starting with mobile and web channels first before experimenting the voice channels as they have greater flexibility and scope for complex use cases and conversations.

In terms of next steps, in order to deliver the VA within the recommended six-month timeframe and be the first Energy Company with a VA, there are some considerations that need to be taken into account. If the prioritisation is to integrate the VA with EA's Live Chat to offload volume from the call centre, then assessment of the current Live Chat platform and customization of the VA will likely need to be done in parallel to deliver the VA within six months.

MyWave Virtual Agent Technical Recommendations

The reason behind this is due to the ordering of the process steps in implementing a VA:

1. To determine prioritisation of which conversations to start with, a technical review of which EA systems are currently accessible to the VA will need to be undertaken
2. If any changes to those EA systems are required for the chosen Conversational Use Cases, getting the changes actioned may cause a delay and increase the time to go live
3. Branding, Personality and Tone of the VA will need to be defined and go through Marketing approvals. Legal will need to be consulted on Authentication and Terms and Conditions
4. The experience of the VA for the chosen Conversational Use Cases will be designed and customer tested

5. Integration with the required EA Back End Systems will need to take place while the Conversational Use Cases are being customised and the Natural Language trained.

6. The VA will then go through a controlled release to a small customer segment or internally to EA for further learnings

7. Integrate with EA's Live Chat and/or Call Centre systems

The integration to EA's Live Chat system would occur potentially in stage 5 or after stage 6. Stages 1 and 2 can take time, especially if any changes are required to EA's back-end systems and since much of the work can be performed prior to integrating to EA's Live Chat, it is far more efficient to perform the work in parallel.

Introduction

In this paper, we are going to provide recommendations and discuss the technical considerations of implementing a Virtual Agent for Customer Service in Energy Australia (EA) based off our experience with the MyWave Intelligent Agent Platform.

Virtual Agent vs Chatbot

Virtual Agents and Chatbots are frequently confused, and most sources use them interchangeably, leaving many with the impression that a Virtual Agent such as MyWave's Myia (My Intelligent Agent) is the same as a Chatbot. However, Virtual Agent technology is a newer more powerful implementation while Chatbots are a ten-year-old technology.

A Chatbot handles requests case by case and does not remember the customer. For example, a Chatbot in an airline will ask the customer for a booking reference and their surname to identify the booking involved in the conversation. If the same customer talks to the Chatbot later on the same booking, they need to provide the booking reference again.

A Virtual Agent is the customer's personal assistant and its role is to curate, guide and orchestrate generating an outcome. It remembers and organises information centred on the individual customer. For example, after the customer has logged into the airline's Virtual Agent, it will find the customer booking from their backend

system. The customer simply needs to select from the bookings displayed on the computer screen. The Virtual Agent can also provide other valuable services to the customer such as reminding the customer to exchange currency with an indicative exchange rate for his destination a few days before his departure or other travel related services with context to the booking. This is because the Virtual Agent has memory and context around what the customer is doing now and what they may want to do next. Another important differentiator is that Intelligent Agents have the ability to be not only reactive but proactive. They can take data insights or triggers such as time passing, or a requirement being met to generate a proactive conversation to further assist the customer.

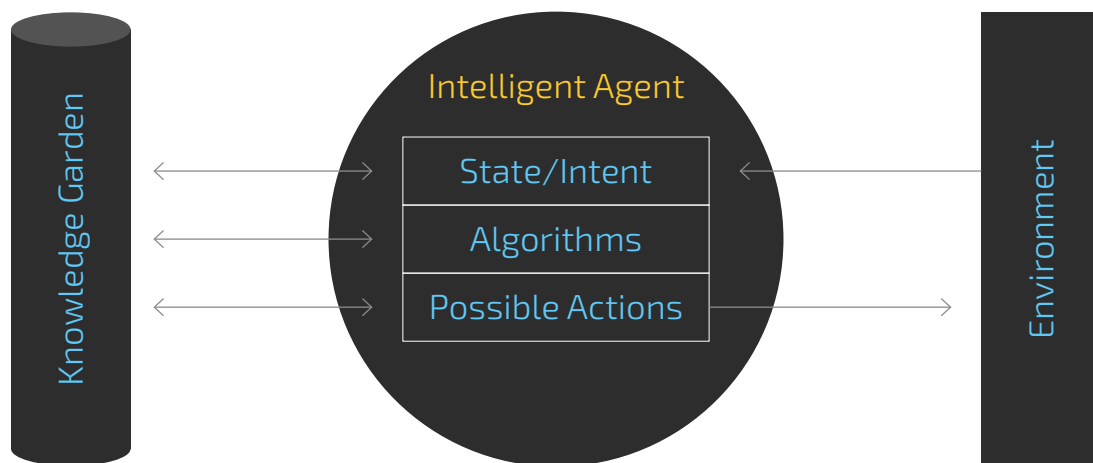
These value-added services provided by Virtual Agents can gradually build up trust from customers which is important for creating a positive customer experience while at the same time driving out inefficiency and cost for EA.

MyWave Enterprise Intelligent Agent Platform

A standard Virtual Agent is typically comprised of five key features:

1. The ability to connect to the Environment around it through channels such as Mobile, Voice, Internet of Things Devices and Data Analytics Engines
2. The ability to understand customer Natural Language Intent or take action from new data insights provided by the Environment
3. Algorithms to determine the best conversation or course of action to take based on input from the Natural Language Intent and the Environment
4. The ability to present the best possible actions based on the Intent or data input back to the Environment
5. A Knowledge Garden or 'Brain' to continuously build up context and information on it's user so that the Intelligent Agent is able to be more personalised and useful in the context of the user's everyday life. This feeds every part of the Intelligent Agent.

Components of an Intelligent Agent



Process Recommendations

A typical implementation from MyWave involves the following components and is designed to enable:

- A customer centric business approach
- Customers not being locked into a single product silo, but able to seamlessly connect across multiple product silos to achieve the outcomes they are after
- Memory, pro-active prompting and assistance and a level of personalisation across all product verticals
- Simple, easy and intuitive user experience for the customer
- Doing as much of the heavy lifting as possible for the customer and employees, automating tasks and removing cost and effort

The general process is as follows:

1. Define and design the VA experience

In MyWave's experience, when designing a VA, there are some considerations to take into account such as: Who is the primary target audience? How does the business currently talk to them? What's the tone? Does it change depending on the domain?

Once identified, these can be used as input when defining the persona of the VA to make sure that the Virtual Agent experience is a true representation of the enterprise and its brand and provides a single cohesive experience.

The persona: Who is she/ he?

Creating a detailed persona is also another vital piece in defining the VA's personality, tone and language; particularly

when designing the conversation-based experiences. Who is she/ he? Where did they grow up? What level of education do they have? Where they live, their aspirations, likes and dislikes etc. Understanding who the VA is, coupled with the brand and its target audience will ultimately shape the conversation and create an experience that is engaging and reflective of the enterprise and its brand.

Another important factor in designing a VA experience is to understand who they 'belong' to; the customer or the business. They both have their benefits and challenges.

If the VA represents the business: **Trust** – it is important that each interaction should build the user's trust in the agent, acting in the interest of another party over the user will destroy their trust in the Intelligent Agent. Help the user understand that even though the VA represents the business, its purpose is to help the customer get an outcome.

If the VA belongs to the customer: **Knowledge** – it is important that the customer feels comfortable and is confident in the agent's ability to know and understand the domain, thus giving the best possible user-centric outcome to the customer.

Finally, no matter who it belongs to, there are some key characteristics a VA should demonstrate in its interactions: **Mutual value** – each interaction or experience should provide mutual value for the user as well as the enterprise. This incentivizes the user to share more information. **Permission** – While the agent may be doing things for the user with minimal effort, it is important that the

Agent always asks for permission before performing a significant task, so the user never feels out of control. **Respect** – it is important that each interaction is respectful and with a purpose intended to benefit the user, no spam, no random pop ups.

2. Explore and identify scenario(s)

One of the most effective ways to explore which scenarios or customer journeys to tackle first would be through a Design Sprint. This would involve a 5-day workshop with the relevant people to identify the best potential customer journey(s) based on areas where a VA would offer the most benefit for both the business and customer. The aim of the sprint would be to have a 'wireframe prototype' of an MVP at the end of the 5 days to test and validate with customers and possibly worked into a live product.

<http://www.gv.com/sprint>

3. Feasibility study and cost benefit analysis

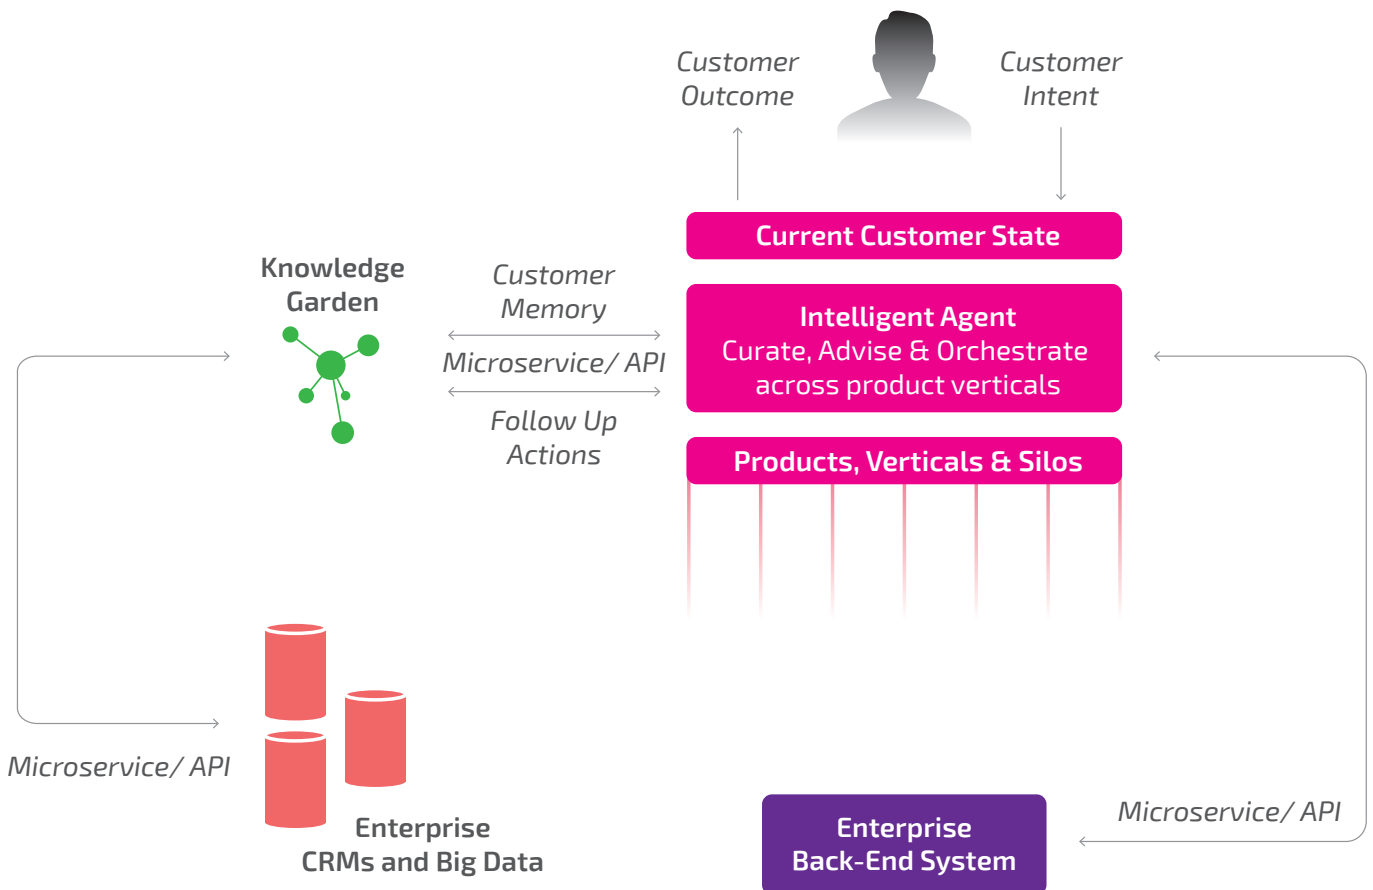
Once the scenarios have been identified, an analysis of the required EA backend system should be undertaken to identify the best way to integrate them with the VA.

If any required alterations have been identified, implementation costs can be estimated and applied to a cost benefit analysis within a business case to ensure the Use Case is economically viable for EA.

4. Implementation

Before implementing any conversations, we typically recommend starting with required integration points such as the authentication system, customer databases and other back-end systems. Software features should be scheduled and implemented iteratively and incrementally. Interactive feedback from customers within the implementation process is key to getting the best user experience and frequent delivery of changes is recommended. We suggest implementing the simplest conversation first to allow to a shorter feedback cycle and quicker learning. Conversations do not have to be limited to current processes but can span multiple verticals/departments to get new outcomes for customers.

Living Intelligent Layer - Generating Living Little Data and Customer Centred Outcomes



It is important that the implementation of the virtual agent does not negatively impact EA's customer experience. To limit the potential for poor customer experiences at MyWave we have found releasing initial implementations to a small targeted customer segment or to internal EA employees only is a good way to counteract this risk. This approach gives EA

the chance to iteratively learn from customer feedback before releasing the Virtual Agent to their wider customer base. It is also useful for gaining additional training data for the Natural Language Processing algorithms and ensure we have a desirable accuracy rate in understanding the customer's intent.

Natural Language Processing

Natural language processing is a very broad topic. Sentiment analysis, natural language parsing and entity tagging are just some of many tools in natural language processing.

For a Virtual Agent, the ability to understand and extract the user's intent from natural language is key in its ability to provide an intelligent experience for the customer. This ability is called Natural Language Understanding.

With current technologies, there is no universal algorithm to understand all human languages. Algorithms can only be trained for limited set of intents with good quality training data consisting of examples and metadata of their meanings (controlled learning).

The more training data the Virtual Agent has access to, the less likely it will be to hand over to a call centre agent to complete the customer's intent as it will understand more queries.

There may be cases that the training data does not cover, and therefore a short fall-back conversation should be designed as part of the implementation to guide the customers to express their intent in a way the VA can understand and guide the customer to the correct conversation.

The level of natural language processing for different languages varies greatly. If EA were to support a language that is not provided by any existing tools in the market it may be quite costly. Similarly, supporting more languages will increase

the hardware requirements and may slightly increase the cost of hosting the service. We would recommend investigation into the predominant languages spoken by EA customers and what tools are available to support each language.

MyWave can support multiple languages as required, however we recommend starting with English and building out from there.

Customer Knowledge Model

Usually the Customer's Knowledge Model (i.e., how data is named, organised, and stored) is designed according to its purpose and usage. For example, concession card numbers are stored as customer reference number in Centrelink but stored as concession card number in EA billing system.

A well-designed customer knowledge model for the Virtual Agent is essential for creating an intelligent customer experience.

Customer Centric Knowledge Model

The Customer Knowledge Model in a Virtual Agent such as the one used by MyWave (Knowledge Garden) is centred on the customer. It will be different from the one in the EA customer database which was designed from a business perspective.

Let's take the service address as an example.

In the EA database, it would be stored as service address under an account, that service address is then linked to a home address and that address has a customer name associated with it.

In the Virtual Agent this data model is flipped so that the account is centred on the customer, the customer may have **one or more properties** linked to their account and each property has an address and metre number.

Benefits of Customer Centric Knowledge Model

Since the Virtual Agent remembers information about a customer and applies this information wherever possible, the data the Virtual Agent learns can be used in multiple unrelated conversations to ensure the customer never starts from zero again and each conversation builds on the context of what they have done previously.

For example, the payment details collected in a **pay my bill** conversation can be used to purchase energy products such as a Tesla Powerwall by simply selecting the payment details previously collected by the Virtual Agent during the customer's energy bill payment. This data strategy also makes up-selling much easier as the Virtual Agent remembers which product was purchased previously and is able to suggest new products of higher value to the customer.

This data model can also be used in scenarios that go beyond EA for example in the PI we tested linking Centrelink data with EA customer data to generate a Concessions outcome in three steps and a few minutes versus the lengthy complex process today.

Voice Application Considerations

To interact with the Virtual Agent by voice, customer speech needs to be converted to text (Speech to Text) and the responses of the Virtual Agent to be converted to speech (Text to Speech). The abilities of speech translation (text to speech and speech to text) depend on the device.

Voice has limitations, the user can only answer one thing at any time, the Virtual Agent is limited to asking customers one thing at a time. For example, if the Virtual Agent needs to ask the first name and last name of the user, it will have to ask for the information in two different questions.

Delivery Channels

Voice Assistant

Voice assistants are gaining popularity. There are currently three major platforms available:

- Google Assistant - Google Home, Google Home mini
- Amazon Alexa - Amazon Echo, Echo Dot
- Siri - Apple HomePod

Currently the speech translations are performed on each respective platform and the language the devices support depends on their settings.

Due to Apple's restrictions it is not possible to integrate with Siri. MyWave and EA's Virtual Agent could however integrate

MyWave Enterprise Intelligent Agent Platform

with the Google Assistant and Amazon Alexa platforms. The limitations of a pure voice interface make rendering complex information in a useful way almost impossible. Voice can however be a good channel to start a conversation (such as find me the best energy plan) but eventually it will have to pass the conversation back to another channel where information can be presented visually unless the Speech device also has an integrated screen (e.g., the Virtual Agent tells the user briefly about the saving of the plan and asks if it should send the plan details to the user's email address).

Mobile App

All new smartphones are equipped with speech to text features like the dictation on iOS. It is possible to ask the user to engage the dictation feature without implementing any specific functionality.

However, third party speech translation services on the device will need to be tightly integrated with the speech translation into the mobile app (such as provide a button for the user to speak to the app and read responses automatically).

There is a cost for the device's third-party speech translation services and the corresponding library must be integrated in the mobile app. The languages the speech translation is optimised for may also be limited, this depends on the provider.

The Mobile app does however provide a more flexible and rich user interface with the ability to combine voice and visual tools together. Outcomes serviced only by a Voice Interface may provide a worse experience due to the limitations of the channel. Therefore, even though there are speech translation services

integrated with the mobile app, we do not recommend that the conversation should adopt a pure voice only approach.

Web Application

Due to security constraints, it is currently not possible to have speech translation solely within a web application. Therefore, speech translation service availability depends on the operating system of the customer's computer. For instance, dictation and VoiceOver accessibility features on Mac OSX allow the customer's voice to be translated to text and read the text displayed on the browser.

Similarly, to a mobile app, a voice only conversation design is not recommended due to the limitations outlined previously.

Conversation considerations

The nature of the conversation is an important factor to determine whether it is suitable for voice. If the conversation involves sensitive data, the conversation occurs out of the home or there is no ability to verify the user, then the conversation is unlikely to be suitable for voice.

The complexity of the conversation is also a factor to consider. There is no help available for users during the voice only conversation and it is hard to keep context in mind. If the flow is too long or complex, users will feel tired, lost and confused very quickly. They may also not give an accurate answer within the time the device expects them to speak. Similarly, a conversation is not suitable for voice if the conversation requires users to gather documented information (such as finding information from their energy bills) as they will likely take a long time to reply.

Target audience

According to the research, younger people and people working in tech are more familiar with innovative technologies. They are more likely to be the early adopters and use voice assistant services (e.g. Google Assistant).

If most of the users are expected to be mature adults (e.g. age over 50), voice is unlikely to be the best medium via which to deliver the Virtual Agent experience.

Supported languages

The most comprehensive voice supported language is English. Support for other languages varies. Even if a language is supported by a device, there are no guarantees as to the quality of the service. We would advise testing any voice language services before committing to support a new language.

Legal / regulatory requirements

Since voice is a limited medium in which to present data, it may fail to fulfil some regulatory requirements.

For example, if the Government requires energy plan details to be shown to customers in writing before they can select a new plan, it will not be possible for a voice only conversation to be used to sign up new customers.

Checklist

We have prepared a short checklist for determining whether it is viable to present a conversation through a voice only channel:

- Who are the customers for the conversation? Does research tell us they will likely work with the Virtual Agent through voice?
- What languages do they speak?
- How many steps are required to complete the conversation?
- How many pieces of information are exchanged in the conversation?
- How long will it take for an average customer to find any required input?
- Is there any sensitive information the customer needs to provide via voice?
- Is there any sensitive information the Virtual Agent tells the customer?
- Is there any information they customer needs to know that is difficult to communicate via voice alone?
- Are there any regulations that prevent us from providing the service via voice only channels?
- Does the customer need to switch to another channel to complete the conversation? If yes, at which step does the switch happen and how do you bring the customer back to the conversation?

Recommendation

Voice assistants are handy for simple tasks but there are many limitations with voice only channels. It is recommended to use the voice assistants to start a conversation for a quick outcome. If the conversation requires more than that, it is recommended to provide an intermediate outcome in the voice conversation and take the customer to other channel, such as web, to continue the conversation.

The voice support brings some additional cost and complexity. It is recommended to implement conversations without voice in the first phase of the Virtual Agent releases. Once the Virtual Agent services are built, EA can extend the services to the voice channels. Similarly, we also recommend initially providing English as the first language and then moving to other languages.

Handover to Live Agents

The Virtual Agent may not be the preferred channel for all customers. Some customers may need more help, some may have complex queries and some customers just prefer talking to a person. Under these circumstances, handover to live agents within a conversation is necessary.

Live Agent channels

Traditionally, customer calls are routed from one call centre agent to another. The customer then needs to repeatedly explain their intention to each person. The context of the conversation is not typically handed over. This is not only a poor user experience but also increases the time and the cost to handle the customer query.

For a positive user experience, it is essential for the Virtual Agent to hand over the customer seamlessly to the live agent with the conversation details and context. MyWave and Eva would enable this experience.

In EA, there are 2 channels for live agents to communicate with customers.

Live Chat

With Live Chat, the conversation can continue within the existing Virtual Agent user interface providing a natural and seamless handover experience.

However, the seamless handover experience may also blur the boundary between the response of the machine (which is quick and responsive) and the response of human (which can be slower). The customer may feel that the agent is being less responsive which may decrease customer satisfaction.

When the customer is handing over to a live agent, it may take seconds or minutes (depending on the workload in the service centre) for a live agent to be available. This may result in the customer believing that the conversation has "hung".

Therefore, the customer expectation should be managed carefully. The conversation should present a message that the customer is being handed over to a live agent and it may take a few minutes.

Phone call

Live agents can also call the customer and continue the conversation over the phone. Unlike live chat, the phone conversation will not be captured within the Virtual Agent.

When the customer service centre is busy, the customer has to wait for a live agent to be free to take their call. Using the VA to schedule a call back from a live agent can save a customer's time and improve the overall user experience.

Further, with MyWave the call centre agent will have the full context of where the customer has got to, improving the customer experience and average handling time to resolve the query or generate an outcome.

Customer Tolerance

Understanding a User's request

Natural language processing needs to be trained and initially may not know every possible query a user might ask. Customers may be frustrated when the Virtual Agent fails to understand their request and replied "Sorry, I don't understand it. Please try again."

Instead of saying 'Sorry, I don't understand' a fall-back conversation should be implemented to guide the user to express their request when the natural language processing failed to identify customer's intent. If the Virtual Agent doesn't understand a user's intent, we recommend that it should report this to the business and that this data be used to further train the natural language algorithm to better understand customer queries.

Examples of valid requests on the user interface will also make it easier for customers to express their requests and learn how to talk to a Virtual Agent.

Stuck in Conversation

The Virtual Agent will keep asking for the same piece of information if the response from the customer is invalid. Good input validation and validation error messages should be integrated in the conversation to allow the customer to easily correct their mistakes.

The design of the conversation flow should also avoid the possibility of 'infinite loops'. If infinite loops cannot be avoided, the customer should be able to break the loop by requesting a chat with a live agent.

Customer service centre system integration

Handing over conversations to live agents with conversational context is important for a frictionless user experience.

Systems used by the live agents should be integrated with the Virtual Agent so that the whole conversation, and where the customer has got to, is presented to the live agents during the live chat session or the call.

For live chat, the live agent chat interface should be integrated with the Virtual Agent so that the live agent response will be sent back to the Virtual Agent and presented to the customer within the conversation. Furthermore, there should be a mechanism to avoid overloading live agents on the live chat

channel. Having customers waiting for half an hour on the live chat channel is a very poor user experience.

EA should factor in workload of their live agents to the automatic handover customers to live chat and call back. Customers should be offered a call back when the customer service centre is under high workload to prevent long service wait times.

UX Considerations

Checklist

- What are the conditions under which the Virtual Agent should present various options for handing over to live agents?
- Do customers prefer one live agent channel over another?
- Should the customer have the choice between live chat and call back? Or it should be determined by the Virtual Agent based on the workload of the customer service centre?

EA Back End Integration

In the EVA Program Increment, the Virtual Agent implemented on MyWave's Myia is integrated with EA APIs (Customer and account details, apply concession & Centrelink concession eligibility) provided by the integration team.

The collaboration between the integration team and MyWave worked extremely well and it is expected that the same approach will be adopted to integrate with other EA systems such as the Knowledge Portal going forward.

Customer Data Synchronisation

The Virtual Agent retrieves information from the customer and account database in EA via an API to minimise the effort from the customer and creates an intelligent Virtual Agent experience with a Single View of Customer.

During conversations, the Virtual Agent will also collect new or updated information from the customer which will be used to refresh the customer & account databases across EA.

Customer Authentication

The customer authentication approach was investigated in the EVA PI. It is recommended to have a logged in experience with the Virtual Agent where the customer initially signs up and verifies themselves using personal details and identifiers from their bill. Once the signup is successful, the customer can log in into the Virtual Agent with their log in credentials as determined by EA.

Back End API Gateway

In the EVA PI, the Virtual Agent was integrated with a EA Back End API gateway. The communication was JSON over HTTPS. This is standard industry practice.

Cost & Time

In our experience, legacy backend systems take more time to change. This may affect the costing and the delivery timeline of the project. Limitations of existing systems may also limit the ability of the Virtual Agent to deliver certain outcomes. Evaluation of

the feasibility of the backend systems in delivering certain functionality should be conducted before finalising the conversation design.

EA "Vault" System

There was a concern about the feasibility of the Virtual Agent be able to integrate with the existing EA VOLT System.

As far as the Virtual Agent is concerned, it typically communicates with backend systems via API gateways. As long as the API gateway has been implemented or there is a way for MyWave to access the data, the Virtual Agent can make use of the service provided by the API and improve the customer experience of using the service.

Security and Privacy

Since the Virtual Agent collects and remembers personal customer information, it is important to protect the privacy of EA customers.

It is recommended to host the Virtual Agent on EA's premises making use of their security and firewalls.

If the Virtual Agent is hosted elsewhere such as in AWS, it is recommended to only host it in the Australia region with proper firewall rules applied. Any data going out of Australia can be intercepted by foreign governments. The communication between the AWS deployment and the EA backend API gateway should be protected with VPN.



*Hello,
how can I help
you today*