

WHITEPAPER

Conversational AI in Banking: 7 Lessons From The Field

How to avoid the most common pitfalls and unlock customer and business value – *fast*



In Brief

FIRSTLY, WHAT IS CONVERSATIONAL AI?

The term conversational AI (CAI) refers to the underlying set of intelligent technologies that enable software systems to interact with humans using natural language processing (NLP). This involves the ability of software to understand the intent behind what a human is saying and respond in an intelligent, conversational way. In the last decade, technologies and use cases have evolved so rapidly that we have seen a deluge of terms enter circulation like chatbot, virtual agent, voice assistant and conversational UI to name a few. For further explanation, see a deeper dive on page 9.



WHO SHOULD READ THIS WHITE PAPER?

In [part one](#) of this two-part series, we explored how conversational AI will revolutionize banking as we know it—and how senior executives and customer-focused leaders can make this new channel a fundamental part of their banks' wider customer engagement strategy.

This white paper is a practical guide for both business leaders and technologists with customer-facing responsibilities in banking, whether you:

- Are at the start of your conversational journey and are looking for best practice on securing investment and building strategies, platforms and experience capabilities
- Have secured investment and are looking to maximize ROI on current conversational AI programs or take your bank's customer experience to the next level

WHY READ THIS WHITE PAPER?

We have reached an intersection of digital engagement and technological maturity that makes the role of conversational AI—and what it can achieve—impossible for any bank to ignore.

In this white paper, we will help you navigate the latest developments in the field and biggest challenges faced by businesses today so that you can create a conversational-first business, saving you time and money while providing a better experience for your customers.

Key Takeaways

Transform Your Customer Experience

The power of natural language and emotional awareness in an interaction should not be underestimated. CAI is therefore a critical part of modern customer experience strategy, helping your bank stand out from the crowd, build new relationships and strengthen existing ones.

Slash Customer Service Costs

With cost pressure relentlessly increasing and every bank's budget squeezed year over year, it's clear to see where CAI can make the biggest impact: the bottom line. Increased automation, reduced wait times, personalization, consistency and more successful service outcomes can all drive significant efficiency and revenue gains. Expect simple automation to deliver around a 30% reduction in customer service queries and up to 80% for more complex and sophisticated experiences that integrate with your back end systems.

Deliver Early Value

Always start by understanding the problems you're trying to solve. CAI is not a silver bullet for everything, but there are almost certainly quick wins to be gained by identifying customer interactions that will deliver maximum value with the lowest effort.

Prepare for the Customer Service Revolution

CAI will continue to evolve and influence our digital ecosystems far into the future. Outstanding customer service will almost certainly depend on a fine balance of CAI and human interaction. Your bank must transform its customer service vision and start preparing now.

Banks have the opportunity to be at the forefront of the next wave of transformation, seamlessly integrating conversational AI to solve the complex challenge of serving and assisting time-challenged customers at every step of their journey.

\$110 billion:

Global spend on AI is forecast to double, growing to more than \$110 billion in 2024, with banking and retail predicted to spend the most (IDC)¹

\$7.3 billion:

Operational cost savings using chatbots in banking globally by 2023 (Juniper Research)²

8.4 billion:

Voice assistant units will have been sold by 2024 (Statista)³

40%

 of 18- to 44-year-old consumers are ready to bank through social platforms and voice assistants (EPAM)⁴

¹ <https://www.businesswire.com/news/home/20180919005045/en/Worldwide-Spending-on-Cognitive-and-Artificial-Intelligence-Systems-Forecast-to-Reach-77.6-Billion-in-2022-According-to-New-IDC-Spending-Guide>

² <https://www.juniperresearch.com/press/press-releases/bank-cost-savings-via-chatbots-reach-7-3bn-2023>

³ <https://www.statista.com/statistics/973815/worldwide-digital-voice-assistant-in-use/>

⁴ <https://www.epam.com/2020-banking-report>

Contents

- A Deeper Dive on Conversational AI..... 5

- Creating a Conversational Bank: 7 Lessons from the Field 8
 - Identify the Problems You’re Trying to Solve10
 - Align the Organization on a Conversational AI Vision12
 - Think Strategically about Your Conversational Platform Architecture15
 - Secure Funding & Momentum for Your Conversational Strategy17
 - Staff the Right Talent: Starting with a Conversational Analyst19
 - Create a Persona for Your Conversational AI Early21
 - Optimize Your Virtual Agent through Agile Design & Delivery 23

- Meet our Contributors 26

- Get in Touch27



A DEEPER DIVE ON

Conversational AI

Understanding Conversational AI

Conversational AI technologies aren't new. However, over the last decade or so, the technologies have evolved so rapidly that we have seen a deluge of terms circulate and find use in the industry—chatbot, virtual agent, voice assistant and conversational UI to name a few. That—combined with the complex nature of these technologies—has made them hard to keep up with, invariably resulting in confusion for businesses and consumers alike. To fully realize the promise of conversational AI, we must first agree upon some key terms and a structured way of thinking about this exciting new technology.


SOLUTION

There are two types of conversational AI solutions:

1. Virtual assistant: personal assistants designed to carry out common tasks for an individual, such as setting an alarm, turning devices on or off or finding information from the internet. Examples include Amazon Alexa, Google Assistant or Siri.

2. Virtual agent: a representative of a group or business, commonly used in a commerce or customer service context, designed to engage with customers to serve business goals (think Capital One's Eno).

'Hey Alexa, what's the weather today?'



'Hey Eno, what's my rewards balance?'



INTERFACE

It is important to note that while several terms are used interchangeably, many are used to tell you more about the interface the conversation happens in (not the actual solution). For example, a chatbot converses with customers through a chat channel while a voice assistant communicates using voice-enabled technologies. A consumer could even use more than one of these interfaces to complete a single interaction. For instance, you could start a conversation with your smart speaker and pick it up later using an app on a mobile device.

As technology continues to advance, more interaction channels and interfaces—such as augmented reality and holography—are likely to emerge, bringing new terms into the mix. It's therefore important that we distinguish interfaces from the solution.

TECHNOLOGY

The term conversational AI refers to the underlying set of intelligent technologies that powers these solutions, including:

- Machine learning algorithms give the software the ability to learn from the data provided, gather insights and make predictions on previously unanalyzed data.
- Natural language processing (NLP) and natural language understanding (NLU) techniques use machine learning algorithms to allow software to comprehend intents through the extraction of information and meaning from human language.
- Together with machine learning algorithms, acoustic modeling enables software to hear and speak using speech recognition and synthesis.


Advances in these technologies over the last few years has led to rapid adoption by businesses, but many have failed to deliver a ROI. In this white paper, we'll help you to cut through the hype to deliver real customer and business value quickly.

“The catalyst for this explosive transformation toward AI assistance hasn't come from an individual technological innovation but from multiple innovations driven by a very human need to help us cut through the ever-increasing digital noise in our lives. Devices, emails, alerts, accounts, passwords, online orders and shipping notifications are increasingly piling up to the point that managing it all has become an overwhelming task in of itself. The recent rise of conversational AI has come from a desperate need for an assistive layer between our lives and our digital services; this need will drive the next transformative wave in digital.”



Tom Moran

Associate Director,
Experience Design



CREATING A CONVERSATIONAL BANK:

7 Lessons From The Field

At a Glance

At the start of your conversational AI journey, or ready to take it to the next level? Sections one through four are a great place to start. Experiencing a poor return on investment from current programs? Jump to sections five through seven.

1. Identify the Problems You're Trying to Solve

Understanding what is driving customer interactions is the first step to identifying which ones can—and should—be automated. Do customers find it difficult to navigate your help and support options on your website? Are there long wait times in your contact center? In this section, we help you determine which use cases to automate first. **Page 10**

2. Align the Business on a Conversational AI Vision

Given rapid growth in market maturation and adoption, it's important that your bank embrace conversational AI as a critical means of interacting with—and delivering services to—customers along many complex touchpoints. Learn how bringing cross-departmental stakeholders on the journey with you is fundamental to developing a conversational bank. **Page 12**

3. Think Strategically about Your Conversational Platform Architecture

Here, we explain why it's critical to have clear visibility of your wider contact strategy and a vision for automation when at the start of your conversational AI journey. **Page 15**

4. Secure Funding & Momentum for Your Conversational Strategy

There has never been a better time to invest in conversational AI. However, how you paint a picture of customers' needs against a backdrop of a rapidly-evolving landscape will be key to gaining stakeholder buy-in. Here, we give our top tips on creating a business case that will convince even the most skeptical decisionmakers. **Page 17**

5. Staff the Right Talent: Starting with a Conversational Analyst

As with any new technology, it takes time for good design principles to catch up. We're now in a position where best practices and principles exist, so it's critical to ensure you're capitalizing on these when forming your team. Here, we explain why hiring new roles—such as conversational analysts and conversational designers—are crucial to the success of any conversational AI team. **Page 19**

6. Create a Persona for Your Conversational AI Solution Early

As humans, we're wired to anthropomorphize everything from pets to cars—and technology is no exception. If you don't think about your solution's personality early, you're likely to create one that doesn't represent your brand. Learn why spending time on designing a persona will build trust and lead to a significant increase in your performance metrics. **Page 21**

7. Optimize Your Virtual Agent through Agile Design & Delivery

To deliver a successful conversational AI solution, your design and delivery teams must adopt an agile mindset and embrace design thinking. From rapid prototyping to defining the right metrics, here are our top tips for increasing performance through design and delivery. **Page 23**

1. Identify the Problems You're Trying to Solve



KEY LEARNING

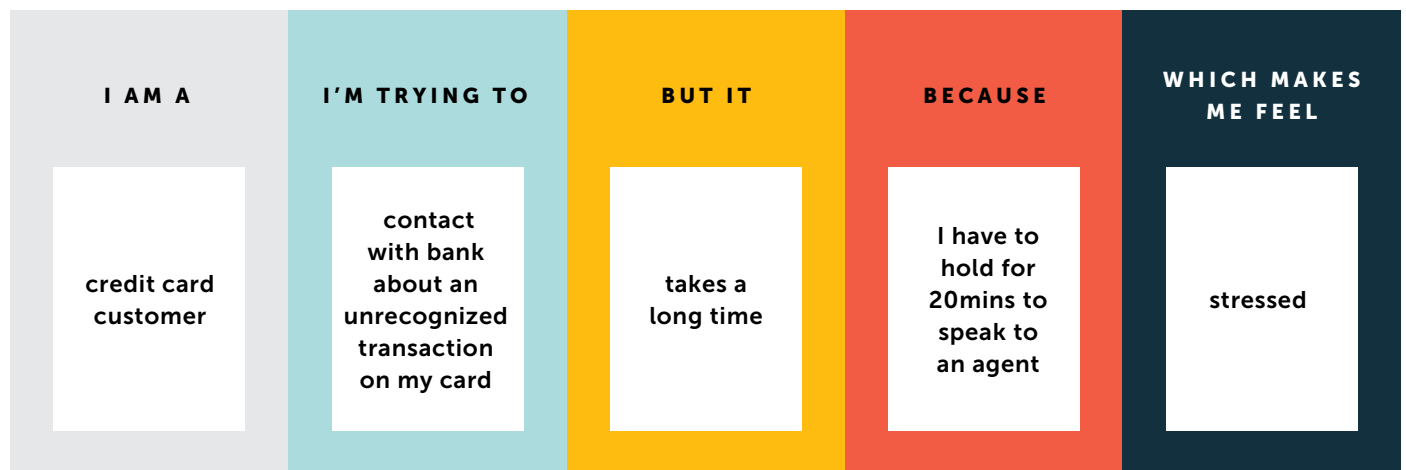
Many Conversational AI projects fail as a result of poor discovery done at the beginning of the project. Make sure you spend time upfront checking what you're building is desirable for customers, technically feasible and has business viability.

The business value for chatbots is often apparent when you look at the numbers; however, being led too much by cost reduction or other drivers without adequately considering your customers' needs is a mistake. The vital thing to get right when deciding where to start or what to do next is balancing business and customer value with effort and complexity.

Many banks will spend months building so-called "MVPs" only to find that customers have no interest in what is delivered. To avoid these situations, start with a strong focus on customer needs and frame problem statements. Then, ideate and create hypotheses for how you think a virtual agent can solve these problems. The next step is to use real MVPs and experiments to test ideas and concepts early. The best teams do this within days, not months.

If you want to deliver actual customer value, it's essential to have a good understanding of why and how customers contact your bank. Look to call and chat data, digital usage data (especially FAQ webpage visits) and complaints data. Don't just look at individual interactions but overall journeys and customers who have to move between channels to resolve their queries. Common customer problems include high wait times or customers being confused by the myriad of help and support options available to them.

FRAMING PROBLEM STATEMENTS



There are loads of ways to test your ideas early with customers. For example, you can quickly create rapid prototypes with conversation design tools, like Botmock or Botsociety, and get these into customers' hands in a research setting.

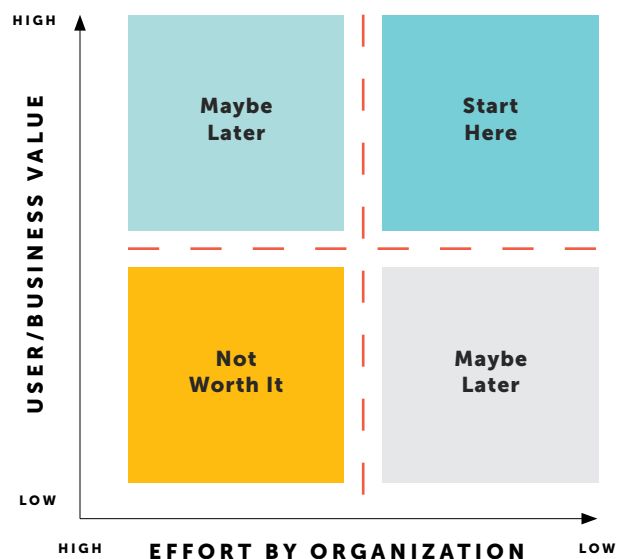
Or run a Wizard of Oz MVP. Just like in the film, *The Wizard of Oz*, where the all-powerful-wizard turns out to be a man behind a curtain pulling some levers, a Wizard of Oz MVP allows you to create an illusion of a virtual agent to gain an understanding of how your product will be received by your customers, without requiring development.

Many of our clients use Wizard of Oz tests in live environments where their best agents act out the role of a virtual agent to test hypotheses and validate risky assumptions they've made. For example, "Would a customer want to switch their mortgage deal using a chatbot?" or, "Would a customer feel secure receiving their account balance from their Alexa?" Trust us; it can save months of wasted design and development time in the long run.

Finally, make sure to balance user and business value with organizational effort. To decide which use cases to automate first, think about those that will deliver maximum value with the lowest effort by mapping each use case on a value versus complexity matrix.

High Value: High-value efforts will be use cases with your significant volume drivers. They can also be customer queries prone to high peaks and troughs in demand like Individual Savings Accounts (ISA) queries around tax year-end, which add pressure on your operations team. Other high-value intents are those prone to staff error, such as mortgage or account opening conversations where agents must read compliance scripts word for word.

Low effort: The most straightforward use cases to automate don't require authentication or back-end integrations. Those that do require back-end integrations will be easier if you can already tap into existing APIs or RPAs without building them yourself. For a bank, another key effort factor is the risk involved in launching some use cases. For example, making payments or changing addresses will appropriately have more red tape. Also, opening new channels, such as WhatsApp or Google Assistant, are relatively easy technically but have many more legal, security and privacy considerations. Involve colleagues from these specialist teams as early as possible to understand any risks or issues to be aware of upfront.



Below are some example use cases to consider:

1.Simple FAQ and triage: A virtual agent that can answer simple customer queries that don't require back-end integrations is a great place to start. Many customers don't know what self-service options are already available to them, so use your virtual agent to teach them how to do this, transferring them to human agents only when necessary. If you have a high call or chat transfer rate in your operation, simply triaging customers to the right person the first time can reduce anywhere from 10% to 40% of contact times.

2.Balance and transactional queries: This is probably the most simple use case to launch when integrating with your enterprise systems. The real value here comes from surfacing this information in new and accessible channels, such as smart speakers or coupling the data with richer information than customers currently get through other digital channels, to solve queries like unrecognized transactions.

3.Fraud Notifications: Fraud notifications currently sent via SMS from banks generate millions of inbound calls from customers who ironically question their legitimacy. Integrating your virtual agent with your fraud system can be a secure and trusted way to follow up with customers and validate suspicious activity. Reducing these inbound calls can allow your fraud teams to reinvest in more fraud checking activity and stop more fraudsters.

4.Spending analysis, budgeting and investments: Check out UK-based fintech Plum, which couples conversational AI with machine learning algorithms to provide spending insights, budgeting and automatic savings and investments, helping millennials to save money.

5.Cross-selling and conversational commerce: When you're at scale across your customer base, advanced analytics and integration with your CRM will allow you to understand customer needs more deeply and promote the right services and products in conversations. This will drive new revenue in your business. Look to other industries, such as telcos that are even doing end-to-end sales journeys using conversational AI.

2. Align the Organization on a Conversational AI Vision



KEY LEARNING

To ensure your conversational AI program isn't viewed as a digital side-project, create a shared vision for CAI across the bank and ensure it's a vital pillar of the broader contact strategy.

As we explored in '[Why Conversational AI will Revolutionize Banking as We Know It](#),' now is the time to make conversational AI a fundamental part of your bank's wider transformation strategy.

Whether you're experimenting with your first conversational interface or developing more sophisticated platforms and experience capabilities, organizational collaboration and alignment is fundamental to developing a successful conversational AI program. Given rapid growth in market maturation and adoption, it's important that stakeholders across the business embrace conversational AI as not just an interface type but as a means for achieving larger organizational goals.

There are three core areas that need to be proactively addressed to ensure the business is aligned around a common vision and organizationally prepared to make that vision real:

1. Co-Create and Articulate a Realistic—Yet Ambitious—Future Vision for the Business

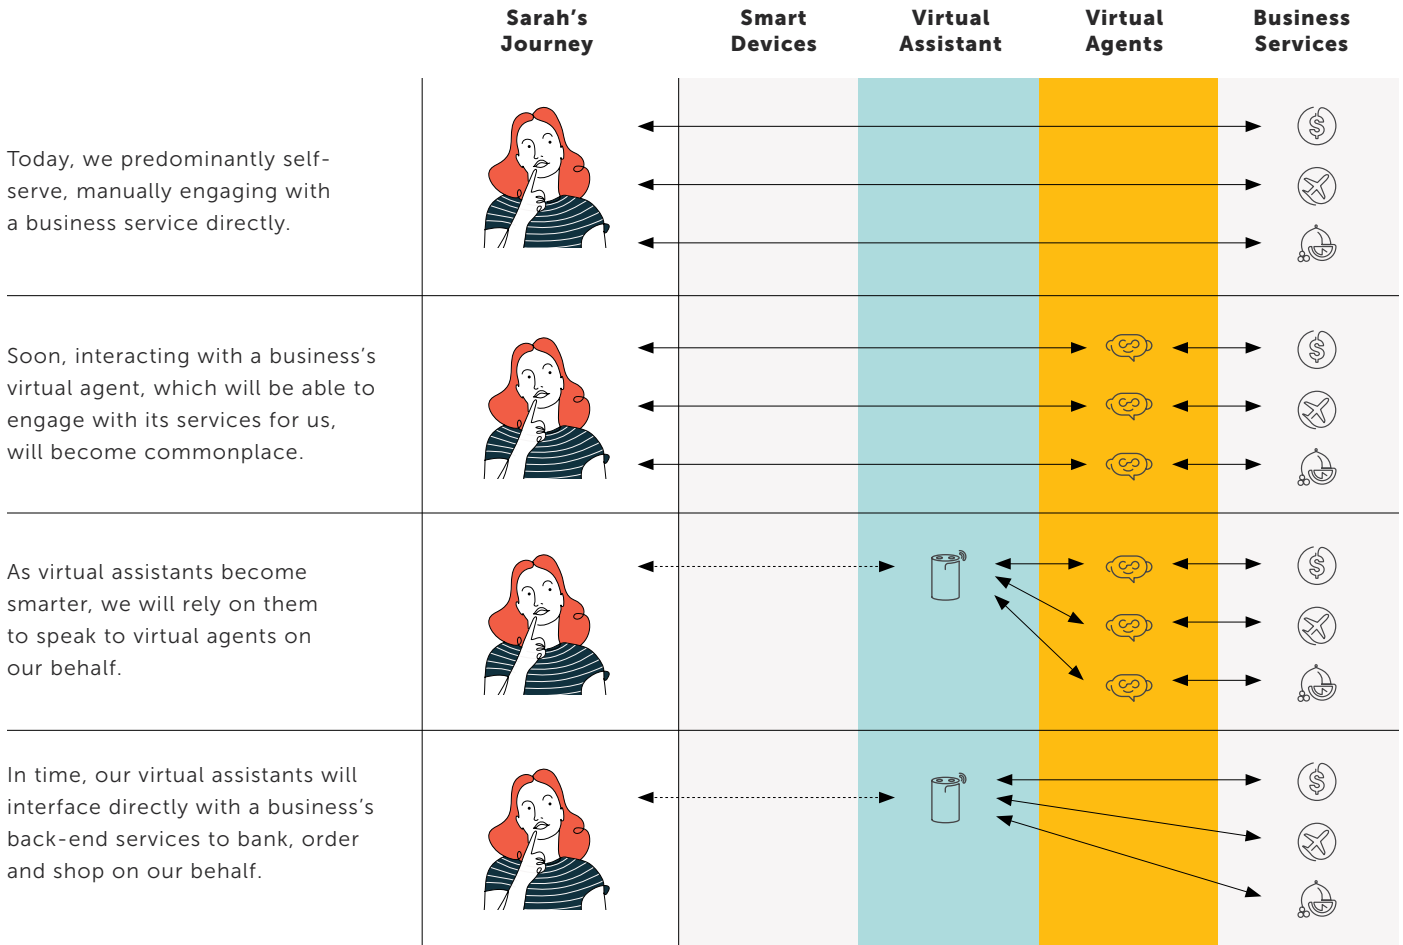
Your conversational AI strategy should clearly articulate how your new initiative will help to solve key challenges faced by employees and customers—both now and into the future.

To do that, you need to adopt approaches and methodologies that will allow stakeholders to fully explore the art of the possible. The key is to ensure that stakeholders feel inspired and empowered to contribute, ultimately leading to a clear consensus based on informed decisions regarding the organization's future direction. By uniting senior stakeholders across the organization towards a common goal, you can create one conversational platform and framework that can be reused by all business units, for every scenario.

On the following page we take a look at where Conversational AI is headed, and key methodologies used.



HOW CONVERSATIONAL AI WILL EVOLVE OVER THE NEXT 10 YEARS



EPAM CONTINUUM'S METHODOLOGIES

EPAM Continuum utilizes a range of methodologies for shaping future visions, including:

- Backcasting:** Working together, we define a desirable future vision for how the business will leverage conversational AI, and then we work backwards to identify near-, medium- and long-term steps that need to be taken to make this vision a reality.
- Future Framing:** For this methodology, we ask senior teams to consider a range of potential future scenarios driven by conversational AI before agreeing and aligning on how the business will respond to each scenario—if and when it becomes a reality.

2. Address Organizational Siloes

One mistake we often see is that businesses try to deliver conversational AI programs within siloed organizational structures. For example, a mortgage team may launch their own virtual agent without thinking about how it will communicate with virtual agents you may have working across other customer journeys. The vision described earlier is fundamental to breaking down these barriers and should serve as a “North Star” for all teams. Without first addressing these silos, businesses will inevitably deliver fragmented experiences that have a high risk of failure.

An additional dimension to breaking down siloes is ensuring strong alignment between your conversational AI teams and your operational teams. Often, operational teams are engaged either too late or too little in the process, leading to disengagement among stakeholders or even projects being pulled.

Any good conversational strategy needs close partnership between these teams from the outset. It is a partnership that is mutually beneficial—you can’t launch a conversational experience that isn’t underpinned by human support. From an operational standpoint, you can help them to drive their efficiency.

3. Know When to Pivot

While creating and aligning on a conversational AI vision and strategy at the start of your journey is fundamental, so too is knowing when to refresh and pivot—and having guardrails in place to do so. As we outlined earlier, it’s clear that this technology—as well as the adoption of it—is accelerating rapidly and there will be many developments along the way. That’s why it’s crucial to build future roadmaps with flexibility.

So... how do you do that? Here are some key questions to ask yourself:

- *How can we create multi-disciplined agile teams that can continually adapt to customer and business needs?*
- *How can we scale this agility up to the executive level to ensure we retain organizational alignment while pivoting?*
- *How can we test and experiment early to learn if solutions are desirable to customers, feasible to build and viable for our business?*
- *How should we build our solution today to allow for switching to different—or a combination of multiple—AI platforms and services in the future, so that we get best capabilities and price even when the ecosystem landscape changes?*
- *Can we discover unmet market needs for a new product or service using AI technologies?*
- *Is our solution scalable if more business groups want to replicate our success?*

In addition to thinking through these questions at the start of your planning, it’s essential to put in place a continuous process of external observation so you can adapt to emerging trends across technology, customer behaviors and broader societal shifts.

Over the past year alone, we have seen the creation of the most powerful language model ever (GPT-3 by Open AI). Google publicly launched Meena, reportedly the most human-like open domain chatbot yet, and Microsoft proposed a patent that has the potential to reincarnate a lost loved one, with the ability to copy their exact styles and communication patterns.

Re-evaluating your strategy and solution design in light of a new technological development or industry trends can simplify the build or uncover new possibilities in what you can do with a conversational AI solution.

“Outstanding customer service usually derives from intuitively knowing how to serve customers at every step of the journey. Translating that understanding into a holistic service strategy and seamlessly integrating current and emerging technologies will allow your bank to rapidly adapt to changing consumer needs and expectations. Starting small is recommended, but with the end goal constantly in mind, giving yourself the flexibility to evolve as the world moves towards more ubiquitous experiences. Success will come from finding the perfect balance between automation and fundamentally human moments.”



Panos Archondakis

Global Head, Banking & Wealth Management

3. Think Strategically About Your Conversational Platform Architecture



KEY LEARNING

Beware of the sales pitch from organizations that say implementing this technology is quick and easy. Despite advances in Natural Language Generation (NLG), implementing and training conversational AI is relatively manual and time-consuming, if you're not using the right approach. A flexible architecture with the right building blocks and a data-driven approach is key to automating as many processes as possible and delivering value fast.

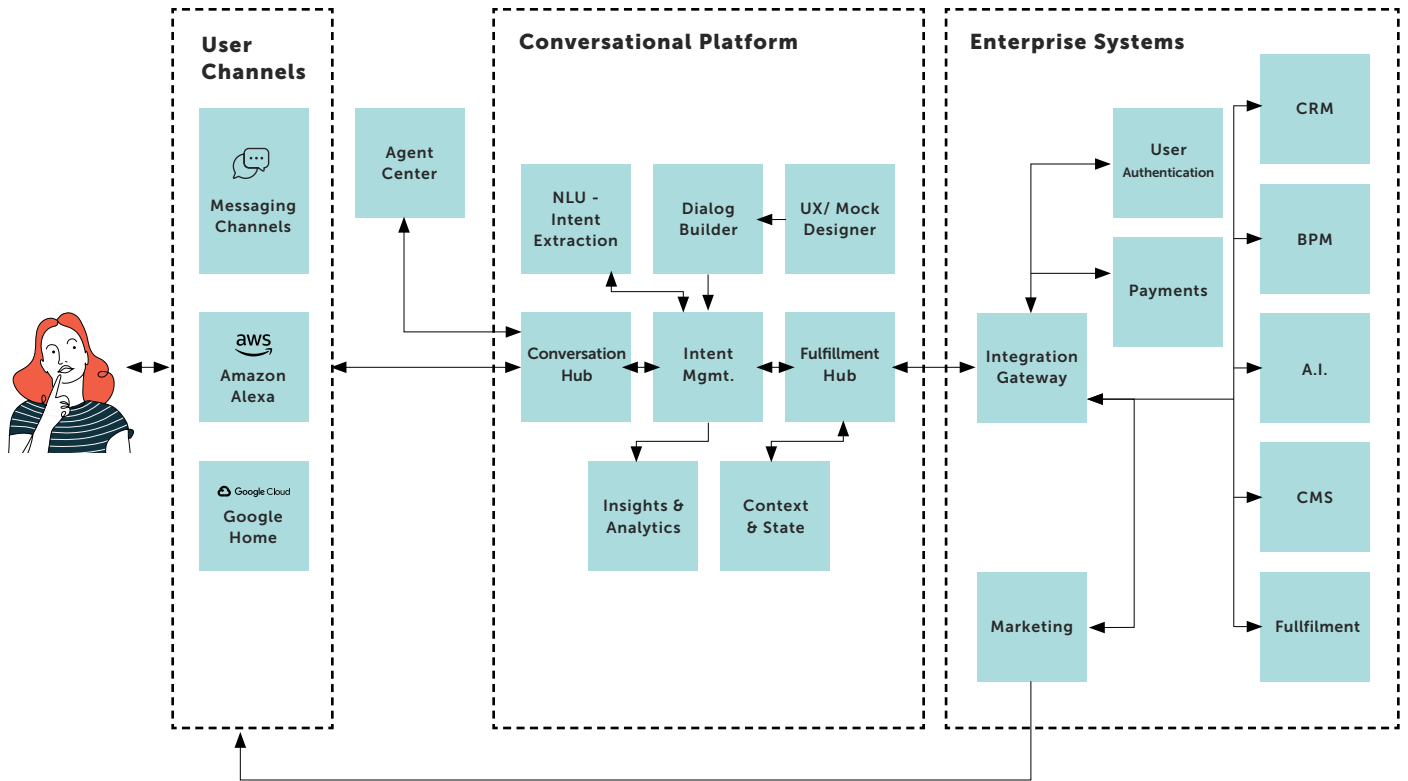
As you progress, far-reaching and long-lasting architecture and solution decisions will have to be made to put principles, guidelines and foundations in place. Where possible, leverage existing investments and consider the needs of all business units that may want to deploy conversational AI solutions in the future.

It is helpful to think about virtual agents as having two distinct architectural phases:

Phase One: Implementation of the architecture for your virtual agent should create a conversational platform that integrates with critical communication channels and can seamlessly handover to human agents within those channels. It doesn't have any integrations into back-end enterprise systems, but it can already deliver significant value. For example, one of our customers managed to contain 30% of its conversations within this kind of architecture. Their virtual agent simply answered frequently asked questions and directed customers to self-service options. Keeping the focus on these critical initial building blocks will enable you to get value quicker. However, try to create a platform that's able to be scaled to handle future needs to avoid a costly rework later.

Phase Two: Focus on delivering authentication and hooking what you've already built into your back-end enterprise systems to unlock end-to-end use cases, such as transactional queries. We've seen banks take well over a year to wade through the governance and processes involved in delivering this phase. Start early and try to progress in parallel with phase one activity in order to incorporate new back-end capabilities once they are available.

REFERENCE TARGET ARCHITECTURE — CONVERSATIONAL COMMERCE & CARE



As part of your conversational platform, you'll need to select the right NLU, e.g. IBM Watson, Google Dialogflow or Amazon Alexa, and dialogue building tooling if the platform providing the NLU does not provide one that meets your needs. As part of the selection process, be mindful that these platforms' key features, pricing and performances vary and might impact the delivery and design of the solution. There is no best in market—the best solution for you will depend on your requirements, broader ecosystem and technology and cloud providers. It's best to get hands-on in the platforms yourself and run technical proofs of concepts. There are many conversational AI platforms out there, but few will be able to handle the scale required for an enterprise-level organization.

If you're looking for the best in class for your overall architecture, it will mean partnering with multiple suppliers for various components. Given the ever-changing landscape, it's also a good idea to stay away from tying your business to one provider or a monolithic solution. We'd advise you to have an easy-to-implement migration strategy if you want to switch in the future.

How you integrate and manage different channels is key because it can be onerous if you have to build and maintain these integrations from scratch. Consider the best architecture to deliver use cases that are both inbound (at the customer's request) and outbound (initiated by events or marketing). It is also critical to have clear visibility of the broader contact strategy when making decisions here. For example, banks with a clear omnichannel strategy will need to consider early how their virtual agent will integrate with multiple channels and store and pass context between these channels seamlessly.

Finally, build your architecture with security and privacy in mind and involve stakeholders from these teams early in any technical designs or solutions. Be aware that all conversation data—even redacted ones—are classified as Personally Identifiable Information (PII). Not having the proper controls in place around platform user access management, segregation of duties for deployments, your data pipeline, data storage and redaction or content publishing will hinder your ability to scale.

4. Secure the Right Funding and Generate Momentum for Your Conversational Strategy



KEY LEARNING

Conversational AI is a journey not a destination. It is critical to have a realistic and transparent business case. There are clear benefits to be had but not all will be realized immediately; continuity and commitment are key to success.

You may be convinced that creating a conversational AI program is critical to digitally transforming your bank, but now it's time to get everyone else on board. While some may be put off by the time and money needed to implement and run a successful conversational AI program, it's easy to start with a low initial investment, prove the benefits and then scale up to millions of customers.

Here are our top tips for building a strong business case to key decisionmakers:



1. Think Bigger Than Just Cost Reduction.

Often banks think of conversational AI programs as a means to cutting back on operating costs, but they can do so much more than that. First and foremost, conversational AI should be about transforming your customer experience. Here are some questions to get you thinking about other ways it could benefit your bank:

- Can this technology help you leverage data to understand your customer better than before?
- How might your staff benefit from virtual agents?
- How could conversational AI help your business boost revenue?
- Are there key complaint drivers that could be solved with automation?

- How can this technology be utilized to leverage other forms of AI and automation to really solve your business and customer problems?
- If you free up capacity from human agents, how can this be reinvested?
- Can a virtual agent enable you to launch or scale more channels that customers want?
- How could you execute proactive marketing or engagement via a virtual agent?
- How could conversational AI help you reduce fraud or ensure compliance?

By thinking through these questions, your business can benefit in a number of ways that goes beyond cost-cutting.



2. Include Competitor Research as Evidence.

It's important to research the landscape to understand how advanced your competitors are in conversational AI. A good place to start is news articles, analyst reports and annual results, which often include stats on virtual agents. Look to successful financial service virtual agents already in-market (Bank of America's Erica recently surpassed 10 million users, as an example⁵), but also, take learnings from other industries, too. Telcos, such as Vodafone or Verizon, have transformed their businesses with conversational AI. There's a big risk to being seen as behind the curve—drawing on these facts and figures will help you build a compelling case.

⁵ <https://newsroom.bankofamerica.com/press-releases/consumer-banking/bank-americas-ericar-surpasses-10-million-users-introduces-new#:~:text=Company%20Also%20Reveals%20Nation's%20Most%20Mobile%20Banking%20Cities&text=Bank%20of%20America%20today%20announced,requests%20in%20the%20coming%20weeks>



3. Paint a Picture of the Shifting Market Landscape.

Increased adoption in this technology is undeniable. For years, customer expectations have been set by tech giants, such as Apple, Amazon, Facebook and Google. All of these players have made big bets on conversational AI solutions, and, as a result, they are becoming commonplace in our lives. What's more, a recent report from Zendesk reported "half of customers say they like contacting support over the channels they use with family and friends" and "28% of people expect a reply on chat in under 5 minutes."⁶ Without a virtual agent strategy, having a brand presence in channels with rapid response times is far too costly. As part of your business case, it's also important to identify the key risks of doing nothing at all. The gap between your customers' expectations and the service you offer will widen increasingly over time—and it's likely that customers will respond with their feet.



4. Be Explicit on Where Conversational AI will Cannibalize Lower Cost Channels.

One of the biggest objections we hear against launching a virtual agent is that organizations fear they will syphon traffic away from existing web or app journeys. For example, adding a chatbot on your help and support webpages might result in new conversations that your FAQ webpage could have solved—and any conversations not solved by the chatbot will have to be picked up by human agents. However, this isn't always a bad thing. For example, when Vodafone launched its sales chatbot, it actually doubled its website conversion, making a strong case for cannibalization.⁷ If you want to avoid this altogether, we've seen some businesses invest in a behavioral tagging solution for their websites that ensures only customers who would have ended up calling receive an invitation to chat with their virtual agent.



5. Give Yourself Some Breathing Room.

When forecasting benefits, give yourself some breathing space. You're building a new platform from scratch, so expect the unexpected along the way. You won't receive all of the benefits on day one as these solutions take time to train and improve. You will, therefore, need continuous optimization of what already exists while you move onto something new. Also, start small and scale up—we see a lot of businesses trying to run before they can walk. You're forming new teams, and it will take time to get a handle on the velocity at which they can deliver. Ensure you have an achievable set of financial targets at the beginning to avoid low staff morale and risks to the program before it really begins.



6. Build the Right Business Case.

Securing funding for teams and outcomes is a much better approach than building a business case around projects and outputs. The teams must run discovery workshops and identify the correct use cases to deliver the desired results. Coming up with a three-to-five-year plan for how to automate every call your contact center might receive might feel like a great way to engender support, but it takes all the autonomy away from your team and leaves very little room for innovation.



7. Secure Cross-Departmental Budget.

Most likely, there will be other teams in your organization who could re-use your platform and framework. If you're working in a retail bank, is there a business banking area that might want to team up? Or could the IT or HR departments benefit from automating many of the repeatable conversations they have daily? These areas might have budget and benefits to help, so work with them as early as possible. This not only avoids unnecessary duplication but will drive economies of scale with suppliers.



8. Redistribute Funding to Help Drive Innovation.

Finally, could you redistribute funding from other areas of the business to deliver more benefits using conversational AI? For example, redistribute FAQ funding to create a simple FAQ virtual agent or use IVR/telephony funding to build a telephony-integrated voice assistant. Push for a small initial investment first, so you can prove value over existing help and support options.

⁶ <https://www.zendesk.com/blog/difference-chat-messaging/>

⁷ <https://econsultancy.com/vodafones-chatbot-is-delivering-twice-the-conversion-rate-of-its-website/>

5. Staff the Right Talent: Starting with a Conversational Analyst



KEY LEARNING

One of the most important learnings is that the roles and skillsets you need to deliver great conversational experiences are different to what you have in web or app teams. Expect to have to build new role profiles and hire externally.

The size and shape of your team will depend entirely on where you are in your conversational AI journey. The hard part is finding talent with relevant experience in this field when they are in such high demand across the industry. You will also need to create new role types that may not currently exist within your organization, such as a conversational analyst and a conversational designer.

What is a Conversational Analyst?

Conversational analysis is not a new topic. Created in the 1960s by sociologists Harvey Sacks, Emanuel Schegloff and Gail Jefferson, conversational analysis is the study of social interaction—whether verbal or non-verbal.⁸

In your conversational AI team, your conversational analyst will use NLP techniques and machine learning algorithms to understand what customers are asking and how they are asking it, helping to identify your intents, visualize conversational flows and select the training phrases you should use to train and test your virtual agent. They can even use sentiment analysis to understand customers' feelings and predict their likelihood to switch to another bank.

One of the biggest mistakes companies make is basing big decisions on tiny data sets, not knowing if what they're seeing is impacting one customer or many. Over time, you'll pack so much into your designs that they will become over complicated and unwieldy. A conversational analyst will help you avoid that by leveraging insights from hundreds of thousands of conversations.

What is a Conversational Designer?

While your conversational analyst will uncover key insights and make sense of them, you'll need to figure out what experiences to make and how to implement them. That's where conversational designers come in. Part copywriter, part UX designer, conversational designers create the conversational flows, write the virtual agent dialogue, utilize rich features—such as quick replies, buttons and carousels—and train your bots.


Often, companies will hire technical people for this role since the platform may require some coding knowledge. This is a mistake. Don't hire technical people and teach them how to write well. Hire a conversational designer and eliminate the need for them to code using no-code interfaces.

Over the page we break down what an ideal team shape is for Discovery and MVP phases.

⁸ https://pure.mpg.de/rest/items/item_2328034_3/component/file_2328033/content

THE IDEAL TEAM FOR DISCOVERY AND DELIVERY OF AN MVP

Discovery Team



Product Manager

Business Analyst

Solutions Architect

Visual Designer

POC/MVP Team



Product Manager

Business Analyst

Solutions Architect

Conversation Analyst

Conversation Designer

Lead Developer

Senior Developer

QA Engineer

Hire from Your Operations Teams

Take your partnership with your operations team a step further by offering secondments into your conversational AI teams. It's a great way to create buy-in and invest in growing the key skills needed internally—and save time. New hires will already be experts in processes and how to speak to customers. As a result, you'll find a steep drop off in the amount of time needed in the discovery and design phase when they join the team.

There are many bad conversational AI programs and chatbots out there, and that's because—as with any new form of technology—it takes time for good design principles to catch up. We're now in a position where best practices and principles exist. Ensure you're capitalizing on these when recruiting your team.

6. Create a Persona for Your Conversational AI Early



KEY LEARNING

EPAM's own research tells us that customers trust their banks—don't damage that relationship. The ultimate goal is to create a conversational AI agent that is a brand ambassador for your organization.

"Let's make sure everything's factually correct and think about personality later." We've heard this too many times. Please don't make this mistake—your conversational AI program needs a personality, and it's crucial to build it upfront. As humans, we're wired to anthropomorphize everything from pets to cars—and technology is no exception. There is no such thing as a virtual agent without a personality. If you don't think about your solution's personality early, you're likely to create a personality that doesn't represent your brand.

Building Personality

To create an engaging persona, start by thinking about the demographics and psychographics of your customers. Use customer personas if you have them or create them from scratch if you don't. Then create a backstory for your conversational AI program. Think about age, gender, ethnicity, family background, experience, job title, likes, dislikes and personality traits. Can you create a character that resonates most with your target market? It's important to note that your conversational AI solution won't be proactively sharing a lot of this information with customers—the backstory you create is mostly a guide for your conversational designers.

Another common mistake we see organizations make in this area is thinking that personality is only equal to tone of voice. When running customer testing, you'll see a multitude of factors that go into a customer's perception of personality—the size, style and layout of the entry point and interface; the use of rich features; the length of messages being sent; and the delivery speed of text bubbles to name just a few. Keep these in mind when designing your experience.

"When it comes to conversational AI solutions' brand and personality, some companies are taking it quite seriously. Recently, Bank of America hired talent from Disney to help them create their chatbot Erika, and Capital One hired an anthropologist, a journalist and a filmmaker from Pixar to help in their endeavors⁹. From the likes of Pixar and Disney, we can learn how to build emotional connection with characters. It begins with an authentic relationship, builds with positive experiences and culminates in genuine loyalty."



Phil Wilce

Creative Director, EPAM Continuum Europe

⁹ <https://www.reuters.com/article/us-usa-banks-digital-personality-idUSKBN28H0V9>

Building Trust

Persona is important from an engagement point of view, but it's also a way to engender trust with your customers. Eventually, your conversational AI program might live in a multitude of channels, including those owned by third parties, such as WhatsApp, Google Assistant or Amazon Alexa. Regardless of channel, it's important to keep your brand identity consistent to put customers at ease. For example, on a recent project we learned that a segment of customers was more comfortable speaking to their bank's virtual agent via Google Assistant or Alexa if it had a different voice than the assistant itself.

If you spend time getting the design and persona right at the beginning, expect a significant increase in your performance metrics. Alongside improving customer experience scores, one bank saw the number of customers willing to engage with their virtual agent double thanks to improvements to their tone of voice. It's the low-cost way of driving the most benefits from your virtual agent.



7. Optimize Your Virtual Agent Through Agile Design and Delivery



KEY LEARNING

Be flexible. You will not achieve your CAI goals in one pass, nor will you get it completely right the first time. Be prepared to adapt and evolve, using all the agile tools available to listen, learn and continually refresh.

To deliver a successful conversational AI solution, your team needs to adopt an agile mindset and embrace design thinking. Here are some of our tips for making that a reality:

Streamline through Rapid Design and Automated Testing & Deployment

Many conversational AI teams are still heavily reliant upon process mapping tools, like Visio or Lucid Chart, to create designs. The conversational designer will create the flows, write the dialogue, add the logic and then do a handover with the development team, which will then build it in their conversational AI platform. Only after it's built can it be tested with customers. There will likely be a wealth of feedback at this point, resulting in updating the original design in Visio before briefing development again. The whole process can take weeks from start to finish which leads to a lengthy delivery and refinement approach. Customer feedback is captured too late in the process. Mistakes are made, things are lost in translation, and often the flows themselves aren't kept up to date once live.

Take our advice: Throw away your Visio designs and instead start designing in a no-code, rapid prototyping conversation design tool. This allows your designers to mock up designs quickly; you can even interact with prototypes using natural language. The most powerful benefit of this is the ability to test your virtual assistant with real customers in hours and shortcut your learnings, totally independent from your development team.

You can even leverage these designs within your continuous integration and continuous delivery (CI/CD) pipeline. The best teams will actually empower designers to be able to automatically export and deploy these conversational designs into their chosen conversational platform with the click of a button, ensuring any bugs or issues are spotted much earlier in the design process and removing the need for any handover at all with a developer.

Building as much automation into your deployments as possible will help you scale up without doubling your team size for every new use case. Functional testing of your bot, for example, can become very resource-heavy and onerous. However, with current tools and techniques available, it's already possible to automatically generate and run test scenarios for every possible conversation within your virtual agents, including automatically selecting test phrases from existing transcripts. Not only will this save time, but you'll spot expensive errors that would otherwise have made it into production releases.

Continuously Test, Learn and Improve through Optimization

Many companies launching their first conversational AI tool neglect to think about optimization. It's good to bear in mind your virtual agent will never be as bad as it is on day one. Consider launching to a small number of customers, and even take it down after a few hours to improve before re-launch. If this even feels too risky, consider hiring a crowd testing firm. By using testers under NDA, you can see if your solution is working as intended without risking any bad publicity.

Once your virtual agent performs as expected, scale it up by launching new entry points, brands and use cases. But remember, it's crucial to optimize what you have already alongside starting anything new. Any business that forgets this fact will notice a decay in quality over time as the way customers interact with it changes; this is what we call "bot rot."

Alignment on Metrics Gives a Clear Sense of Purpose

Businesses often struggle to determine appropriate KPIs for measuring the success of their conversational AI programs. Below are key metrics to track when getting started or optimizing what you already have:

- **Volume of conversations:** the total volume of conversations handled by your bot
- **Containment:** the percentage of conversations your bot can service without escalating to a human or directing the customer to another channel
- **Handovers:** the percentage of customers escalated to chat, messaging, branch or call center support and which channel they are redirected to
- **Bounce rate:** the percentage of conversations where a customer opens a conversation but never sends a message

- **Abandonment:** customers who leave the conversation before reaching a point in the conversation where they've received help
- **Not handled:** the percentage of utterances where the bot hasn't matched an intent
- **Sentiment Score:** a numeric score that quantifies the feelings of the customer at conversation or utterance level
- **Net Promoter Score (NPS) or customer satisfaction (CSAT):** customer experience scores typically captured via surveys at the end of conversations but which can also be automatically predicted for conversations without surveys to give a holistic view
- **First Contact Resolution (FCR):** percentage of queries resolved by your virtual agent at first contact without the need for them to contact another channel within a defined time period

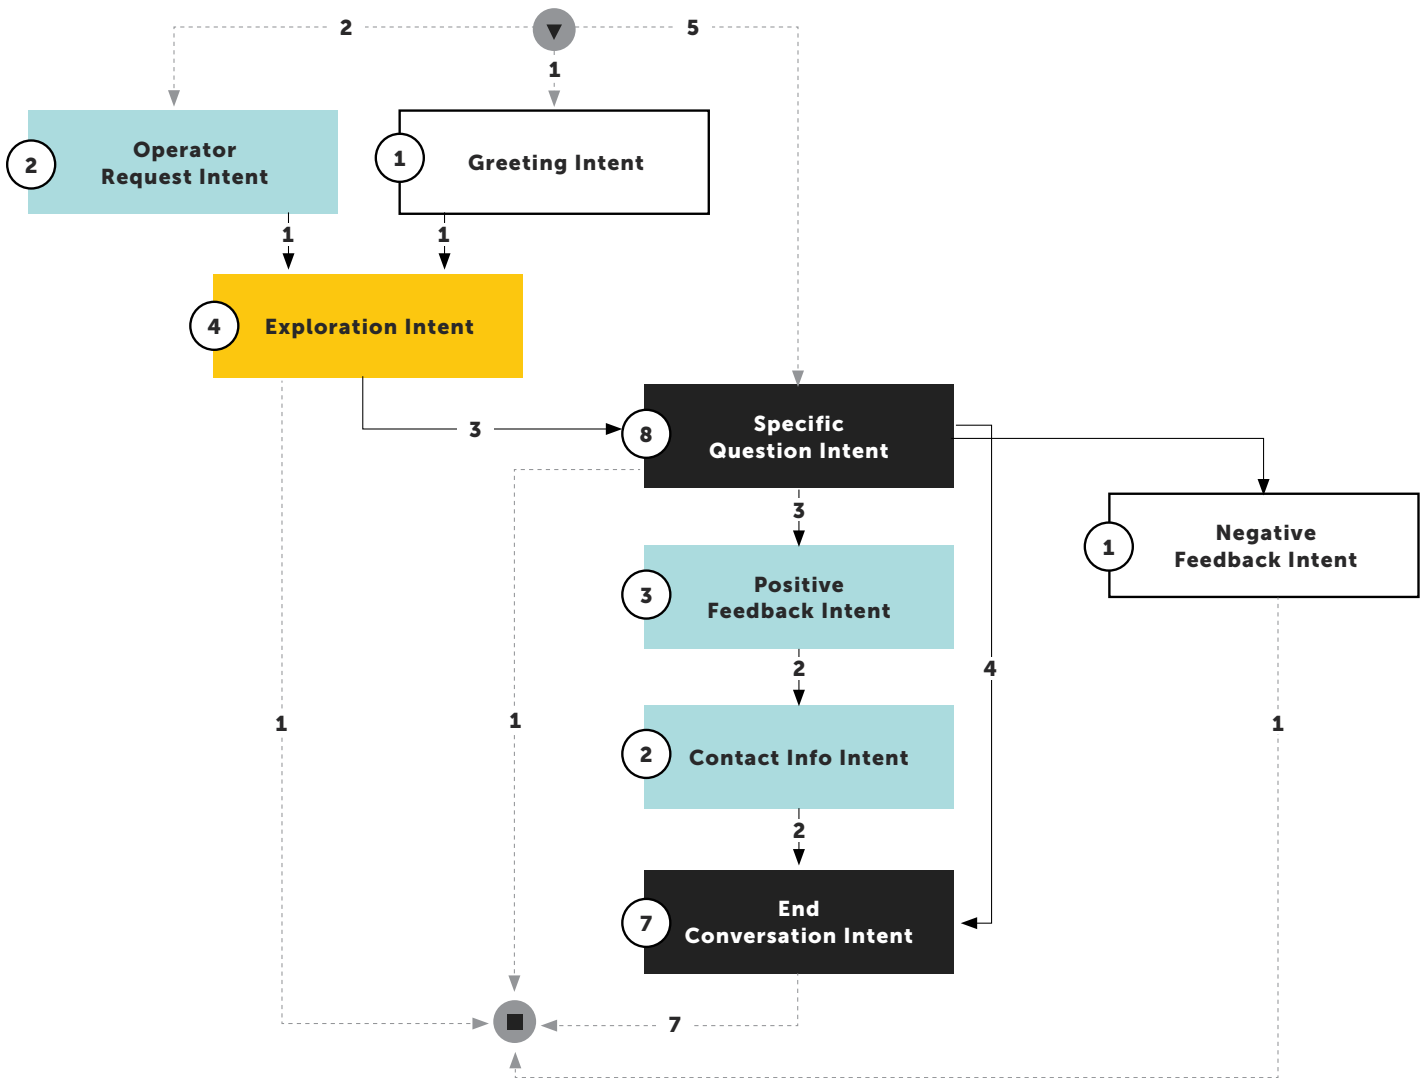
These metrics are crucial to understanding whether the changes you've made have led to improvements and better ROI. As well as viewing them at a macro level across all conversations, it helps to tracking them for each intent, channel and entry point. It's also vital to continually run A/B tests when deploying changes to understand if it's the work you're doing moving the needle or other factors.

Finally, process mining conversations and visualizing the data in conversational flows is a great way to see common paths and where customers are dropping out. Your conversational analysts will be able to identify critical areas of concern for the team to focus on.



BELOW IS AN EXAMPLE OF A SESSION FLOW

Conversion ID	Time Stamp	Utterance	Intent
3	2018-04-02 14:03:02	Can you please give me quotes of mortgage rates	Specific Question Intent
3	2008-04-02 14:03:05	I am interested in 15 years only for refinance	Exploration Intent
3	2008-04-02 14:03:08	Excellent, I would be interested to proceed	Positive Feedback Intent
3	2008-04-02 14:03:12	My email is x@x.com	Contact Info Intent
3	2008-04-02 14:03:15	Thank you, I will be waiting to hear back	End Conversation Intent



Once you've launched your conversational AI program successfully, keep monitoring it and make improvements as necessary so that you are meeting—and exceeding—customers' evolving expectation.

Meet our Contributors



Rachel Bimbi

Product Manager and Conversational AI Specialist

Rachel is a Product Manager and Conversational AI Specialist at EPAM. She has worked with some of the biggest brands in the world to deliver market leading Conversational AI, with a particular focus on Financial Services. She has spent ten years working in the retail banking space, and believes that Conversational AI is the future of banking.



Panos Archondakis

Global Head, Banking & Wealth Management

Panos is responsible for shaping and developing EPAM's strategy for banking and wealth management, tailoring the company's services and solutions specifically for the industry. Panos also works directly with clients to customize their business models, develop digital strategy and lead multi-disciplinary solution teams.



Kshitij Sharma

Senior Manager, Intelligent Automation Consulting

Kshitij (K) is an AI specialist in the Intelligent Automation Business Consulting practice at EPAM where he helps clients define their AI strategy and execute AI and Automation led transformations. K has been advising clients on Conversational AI and helping build successful products and services in Financial Services and other industries since the earliest stages of conversational technology ecosystem.



Tom Moran

Associate Director, Experience Design

Tom has a wide-ranging career in digital and design with a specialization in finance, thanks to his background working with FinTech start-ups and traditional banking platforms. At EPAM, he conducts user research and testing, leads design workshops and helps to create beautiful and meaningful design experiences.

Get In Touch

Conversational AI will take center stage in the next wave of transformation in banking. Whether you're looking for a partner to co-create a vision, define the best architecture for your platform, augment your design and delivery teams or simply discuss where the future is headed, we'd love to hear from you.

Let's build the future of banking together.

EPAM-CONTINUUM.COM

enquiries@epam.com

We fuse integrated consulting with EPAM's engineering expertise to accelerate breakthrough thinking into meaningful impact.

THE FUTURE. MADE REAL.