CELENT

OLIVER WYMAN

THE COGNITIVE CORPORATE BANK

HEADING TOWARD THE SUMMIT

Alenka Grealish
19 March 2018

# CONTENTS

EXECUTIVE SUMMARY

KEY RESEARCH QUESTIONS

1 *What's the state of AI in corporate banking?*

2 *Who's on the sharp end in customer-facing applications?*

3 *Who's summiting in the middle and back offices?*

If one were to build a cognitive corporate bank, what would it look like? Over the past three years, the use of advanced artificial intelligence (AI) in corporate banking has been accelerating, and projects are moving into production. This progress enables us to envision what a cognitive corporate bank could look like, that is, one that optimally applies advanced AI as part of a broader digitization strategy involving other automation technologies (our definition of "corporate" is broad, including small business and midsize companies). A cognitive bank is not one filled with robo-bankers but rather one that effectively applies AI when it is the optimal solution and enables bankers to focus less on administration and more on adding value. Over the next decade, AI will not only dramatically change the nature of bankers' work as well as that of their commercial customers but also will reshape competitive differentiators.

In this report, we discuss how advanced AI can improve customer engagement, reduce operational inefficiencies and risk, enhance data analytics, and ultimately drive revenue growth. Figure 1 outlines AI use cases from the front office to the back office by development stage — from base camp (vision and experiments) to the camps in between (proofs of concept) and the summit (in production, battle-tested). It is important to note that not all AI summits are created equal, with some being higher than others and harder to attain.

To provide concrete examples within corporate banking, we examined the advanced AI offering and implementations of over 20 third party providers to banks and corporations. We selected 18 to profile based on the scope and maturity of their AI experience: Ayasdi, Brighterion, Genpact (acquired Rage Frameworks), HighRadius, IBM, iGTB (Intellect Design), Infosys, Intel (Saffron), Mphasis, Oracle, Pegasystems, Pelican, Previse, Sage, TradeShift, Wipro, and WorkFusion. The 18 represent a wide spectrum of providers from tech giants building AI platforms to address multi-industry challenges to single use case providers.
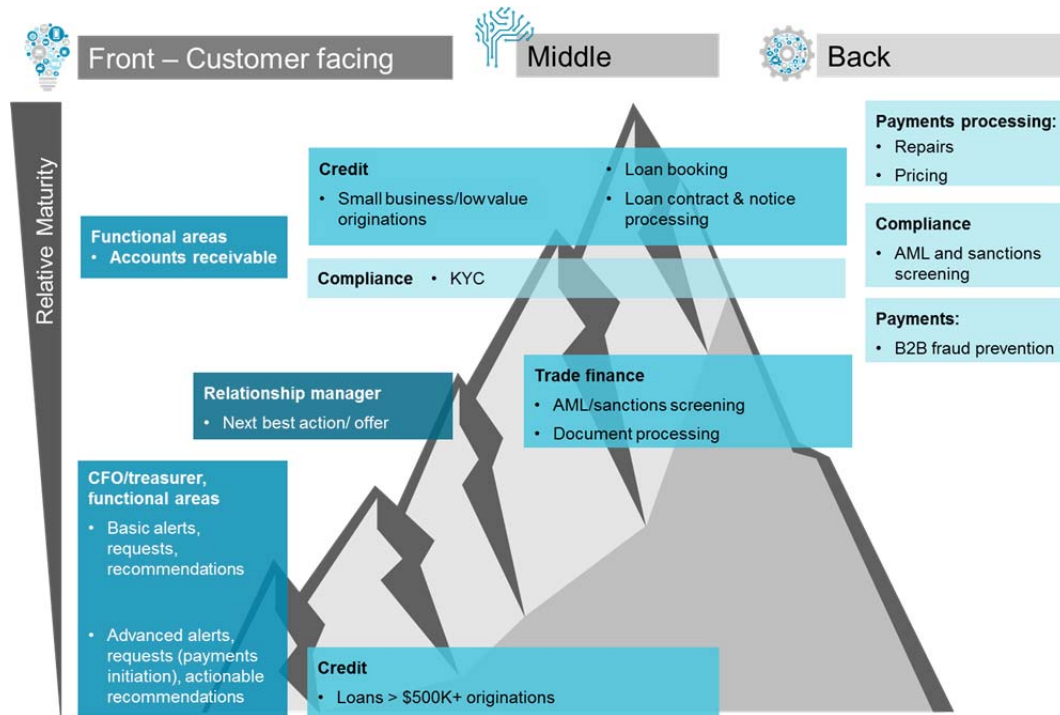
The advances made by these providers and their bank and corporate partners bode well for significant improvements in corporate banking customer engagement and operational efficiencies over the next decade. The banks and corporations that incorporate AI into their long-term digital strategic plans will succeed at reaching AI summits unthinkable even five years ago. Like the mountain climbers attempting to summit all 14, they will gain from each experience including the failures and have increasing advantage over those that have yet to leave base camp.

Cognitive Computing and Advanced Artificial Intelligence — Report Focus

Celent uses IBM's definition of cognitive computing: "systems that learn at scale, reason with purpose and interact with humans naturally. Rather than being explicitly programmed, they learn and reason from their interactions with us and from their experiences with their environment." (https://www.ibm.com/blogs/internet-of-things/iot-cognitive-computing-watson/)

We focus on five advanced AI technologies: machine learning, natural language processing (NLP), natural language understanding (NLU; aka computational linguistics), natural language generation (NLG), and visual recognition. NLU differs from NLP in that it captures the context and intent of words by approaching text as language as opposed to data. These technologies are not mutually exclusive but overlap to varying degrees. The dominant one is machine learning, which has been employed for several decades, but has evolved significantly. Its most advanced form, deep learning (aka deep neural net), is driving improvements in NLP, NLU, and NLG. These technologies power virtual agents and bots (often referred to as smart, intelligent, or cognitive), as well as, expert systems. The most advanced models apply adaptive learning to avoid model obsolescence and are capable of modeling at the entity level (e.g., specific logic is developed for each bank customer).

For an overview of AI in banking, see the Celent report *Artificial Intelligence in Banking: Where to Start?*, July 2017.

For a further discussion of the prerequisites and examples of successful AI implementations, please see the upcoming Celent report, *Vision to Reality: AI in Action.*

AI SUMMITS IN CORPORATE BANKING

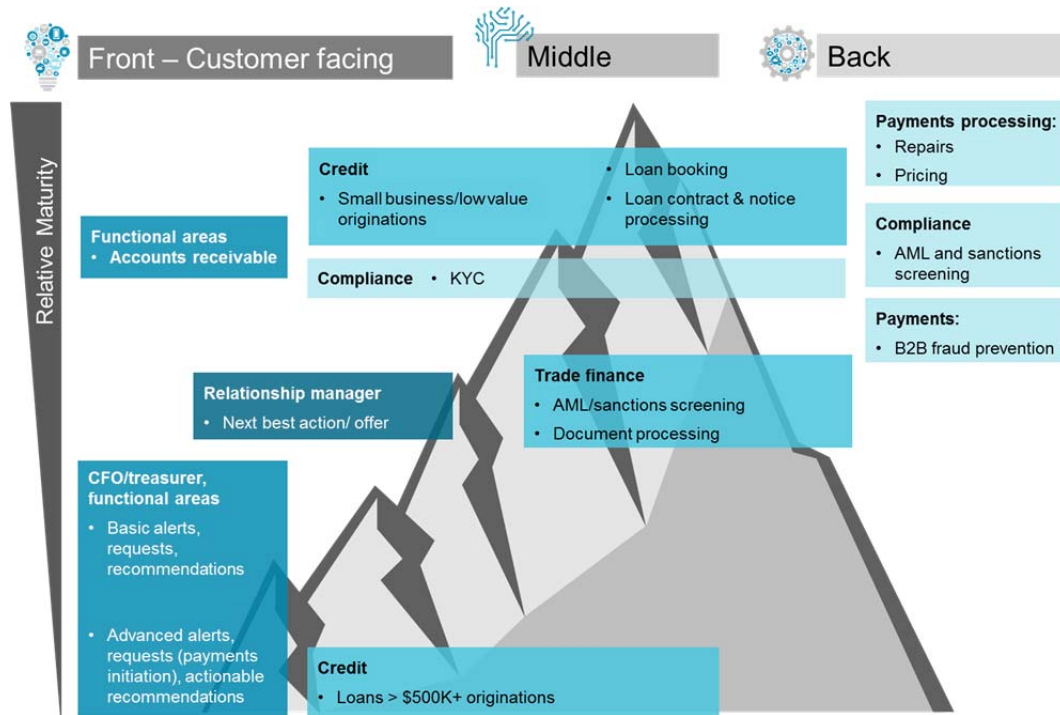*What's the state of AI in corporate banking?*

AI applications have proven themselves in select use cases and are showing strong promise in others. Virtual assistants supporting clients are moving from vision to reality.

While some advanced AI techniques have been around for decades, the technological infrastructure to fully harness them had been lacking. A confluence of forces is enabling advanced AI applications in banking to reach new heights. Among the leading drivers are cheaper, faster computing power (thanks to graphical processing units or GPUs and field-programmable gate arrays or FPGAs) and data storage — most importantly in the cloud, newly accessible data sources, advances in machine learning, open source initiatives by the AI giants, and third party providers developing bank-specific applications.

Thanks to these advances, banks and their third party providers have hit a few AI summits and are advancing toward others. As Figure 2 displays, their current position varies significantly by use case. To date, the strongest returns on investment have been realized in the back office but the front and middle offices are catching up.

To date, much of the energy and success has been concentrated on scaling AI in payments processing and fraud detection and compliance (know-your-customer, anti-money laundering, and sanctions screening). Payments processing is the most mature (e.g., AI-based repairs). AML stands out as garnering the strongest interest among banks, and third party solution providers are uniformly reporting strong interest as well as sales. It is also among the most challenging implementations reported by banks due to the complexity involved in terms of the number of systems, data sources, and regulatory requirements which vary across jurisdictions. The use of AI in detecting business-to-business (B2B) payments fraud is budding. Due to the unique traits of B2B payments, it has proven harder to develop AI than it was for low value/high volume payments (e.g., cards). Trade finance, the quintessential paper and manual-driven back office operation, will eventually be transformed by AI.

Front – Customer facing | **Middle** | **Back**

Relative Maturity

Functional areas
- **Accounts receivable**

Credit
- Small business/low value originations

Compliance • KYC

- Loan booking
- Loan contract & notice processing

Payments processing:
- Repairs
- Pricing

Compliance
- AML and sanctions screening

Payments:
- B2B fraud prevention

Relationship manager
- Next best action/ offer

Trade finance
- AML/sanctions screening
- Document processing

CFO/treasurer, functional areas
- Basic alerts, requests, recommendations
- Advanced alerts, requests (payments initiation), actionable recommendations

Credit
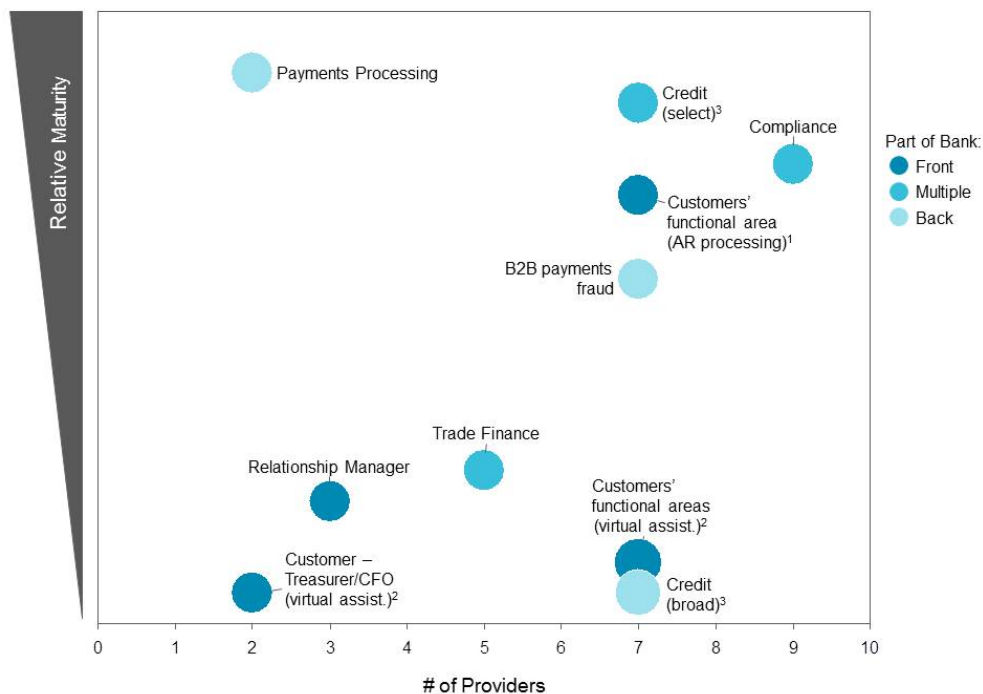- Loans > $500K+ originations

Source: Celent interviews and analysis
Note: Customer-facing includes direct support for CFOs, treasurers, and functional areas: procurement, expense management, accounts payable, and invoicing/accounts receivable.

In the middle and front offices, a few AI summits have been: accounts receivable processing, small business credit underwriting, commercial loan booking, and monitoring. Summits are coming into view for additional applications: virtual assistants responding to basic questions and relationship managers (RMs) being guided by AI-generated next best action/offer suggestions. In the long run, advanced AI will dramatically change customer engagement. Banks will engage with corporate customers through virtual assistants and deliver daily recommendations and enable actions to be initiated verbally. The work of RMs will become much more interesting.

THE MOUNTAINEERS IN ADVANCED ARTIFICIAL INTELLIGENCE

Celent scanned the third party provider landscape to determine which ones have advanced AI applications in corporate banking (at minimum in proof of concept). Figure 3 shows the number of providers by use case and by relative maturity. Not surprisingly, the most crowded competitive field is compliance, with nine providers addressing KYC, AML, and sanctions screening. In contrast, payment processing, the most mature use case, is relatively uncrowded. Payments fraud detection — -specifically business-to-business payments — is a budding application with seven providers and is expected to gain adoption as real-time payments grow. AI support of corporations' functional areas is another popular area, attracting seven providers. While it is relatively immature (except for accounts receivable), AI shows strong potential for both virtual assistants geared toward improving operational efficiencies. Trade finance has garnered resources of five providers. While relationship manager-related use cases have attracted fewer providers, we expect the success of early movers to attract new entrants.

Figure 3: Use Case Maturity and Competitor Concentration

Figure 4 arrays the 18 providers which stood out in eight areas. We differentiate the providers by solution maturity: in production, in implementation or released but not installed, and in pilot/proof of concept or scheduled for release in 2018. In addition, we place the providers in three broad categories. "Multi-industry" includes providers which have built AI technologies that they are tailoring to bank use cases. "Bank focused" comprises providers which have built AI technologies specifically for banks. "Financial supply chain focused" providers have roots in supporting functional areas and have enhanced their offering with AI or have built their offering with AI. We include leaders which are delivering advanced AI directly to corporations (HighRadius, Mphasis, Previse, Sage, and Tradeshift) because it is important for bankers to understand how nonbanks are setting AI precedents and, in the process, raising corporations' expectations.

In the following sections, we delve into each use case and provider solutions.

Figure 4: AI Mountaineers in Corporate Banking

Legend: ✓✓✓ = in production
✓✓ = in implementation or released but not installed
✓ = pilot, proof of concept or scheduled for release in 2018

Use Case / Company	Customer Facing		Relationship Manager (NBA/NBO[1])	Credit[2]	Trade Finance (docu-mentary)	Compliance (onboarding, KYC, AML)	Payments	
	Treasurer/CFO	Functional Areas					Fraud	Processing
Multi-Industry								
AYASDI						✓✓✓	✓✓✓	
Brighterion (mastercard)						✓✓✓	✓✓✓	
IBM	✓	✓	✓✓	✓✓		✓[3]	✓✓✓	
Infosys NIA		✓✓✓	✓	✓	✓			
intel Saffron						✓✓		
Mphasis		✓		✓	✓			
ORACLE						✓✓✓	✓✓✓	
wipro holmes						✓✓✓		
WorkFusion				✓✓✓	✓✓✓	✓✓✓	✓✓✓	✓✓✓
Bank focused								
genpact[4]				✓✓✓				
iGTB	✓✓	✓✓			✓			
NUMERATED				✓✓✓				
Pega			✓✓✓					
Pelican						✓✓✓ ✓✓[5]	✓✓✓	✓✓✓
Financial supply chain focused								
highradius		✓✓✓						
previse				✓✓				
sage		✓✓✓						
TRADESHIFT		✓✓						

Source: Celent interviews and analysis
1. Next best action and/or offer.
2. Includes all types of credit; e.g., trade finance (open account) related credit, such as, supply chain financing.
3. Includes Financial Crimes Insight for Watson—Due Diligence and Financial Crimes Alerts Insights with Watson. Outside of the areas covered herein, IBM Watson is in production in Regulatory Compliance Management and Analytics.
4. Genpact acquired Rage Frameworks in March 2017.
5. Released but not yet installed, Trade-Based Money Laundering.

CUSTOMER-FACING: CFOS/TREASURERS AND FUNCTIONAL AREAS

Who's on the sharp end in customer-facing AI applications?

2

There are a few pioneers bringing advanced AI, including virtual assistants to commercial customers.

Advanced AI will have a profound impact on how humans do their work and engage with machines. It will transform bank-customer engagement. As consumers, we are becoming increasingly comfortable talking to machines thanks to Apple's Siri, Amazon's Alexa, and Google Assistant. This comfort level will gradually flow over to the business world as it did with mobile banking. Treasurers will ask a virtual assistant for FX exposures, risk assessment, and hedging options. Procurement will ask a virtual assistant to generate forecasts, and accounts receivable will ask for a chart showing days sales outstanding for a specific customer vertical. Relationship managers will reach out to clients based on next best action recommendations generated by a machine learning model.

The majority of commercial customer-facing AI applications only recently left base camp, which is not surprising given the complexities involved in scaling on the commercial side compared to the retail side. For most of the leaders attempting the customer engagement ascent, camp 1 is reporting and answering basic questions (e.g., who can help with problem xyz), camp 2 is providing insights/alerts regarding basic financials (e.g., if a customer checks a cash balance online daily, the virtual assistant can proactively relay the balance daily and alert if balance is nearing zero) and enabling payments initiation, and camp 3 is answering more sophisticated questions (e.g., cash flow forecast based on customer's account data and bank's external data). The summit is making bespoke recommendations for responding to an alert (e.g., a cash shortfall is expected; here are three options to cover it, and here are their costs if any).

Figures 5, 6, and 7 chart the third party providers at the sharp end of customer-facing AI. The framework is from a commercial customer lens: the treasurer and his/her team and staff in relevant functional areas (procurement, accounts payable and receivable). Figure 5 lists providers in order of maturity level, and Figures 6 and 7 list them alphabetically within maturity level.
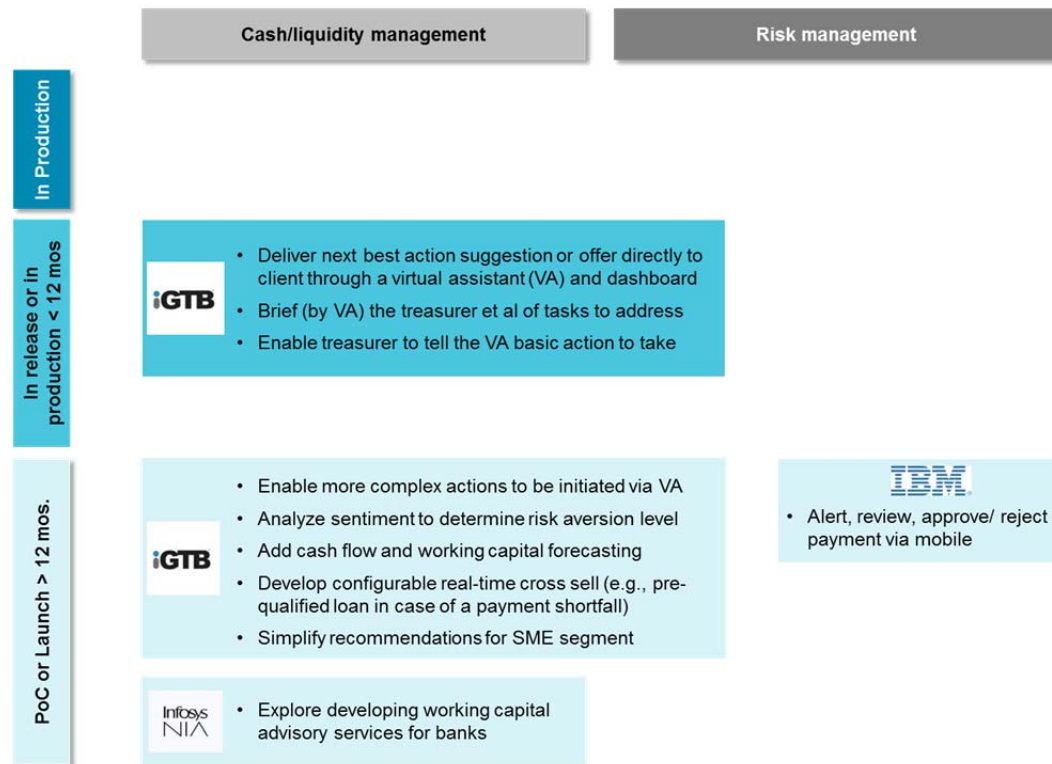
Figure 5: At the Sharp End in Customer-Facing AI Applications — Overall

Use case / Company	CFO/Treasurer		Functional Areas				
	Cash/ Liquidity	Risk Mg't.	Procure-ment	Accounts Payable			Accounts Receivable
				Invoice processing	Risk Mg't.	Payment initiation	
sage			✓✓✓	✓✓✓	✓✓✓		✓✓✓
highradius							✓✓✓
Infosys NIA							✓✓✓
iGTB	✓✓	✓✓				✓✓	
TRADESHIFT			✓✓	✓✓	✓✓		
IBM	✓	✓				✓	
Mphasis				✓			

Legend: ✓✓✓ = in production
✓✓ = in implementation or released but not installed
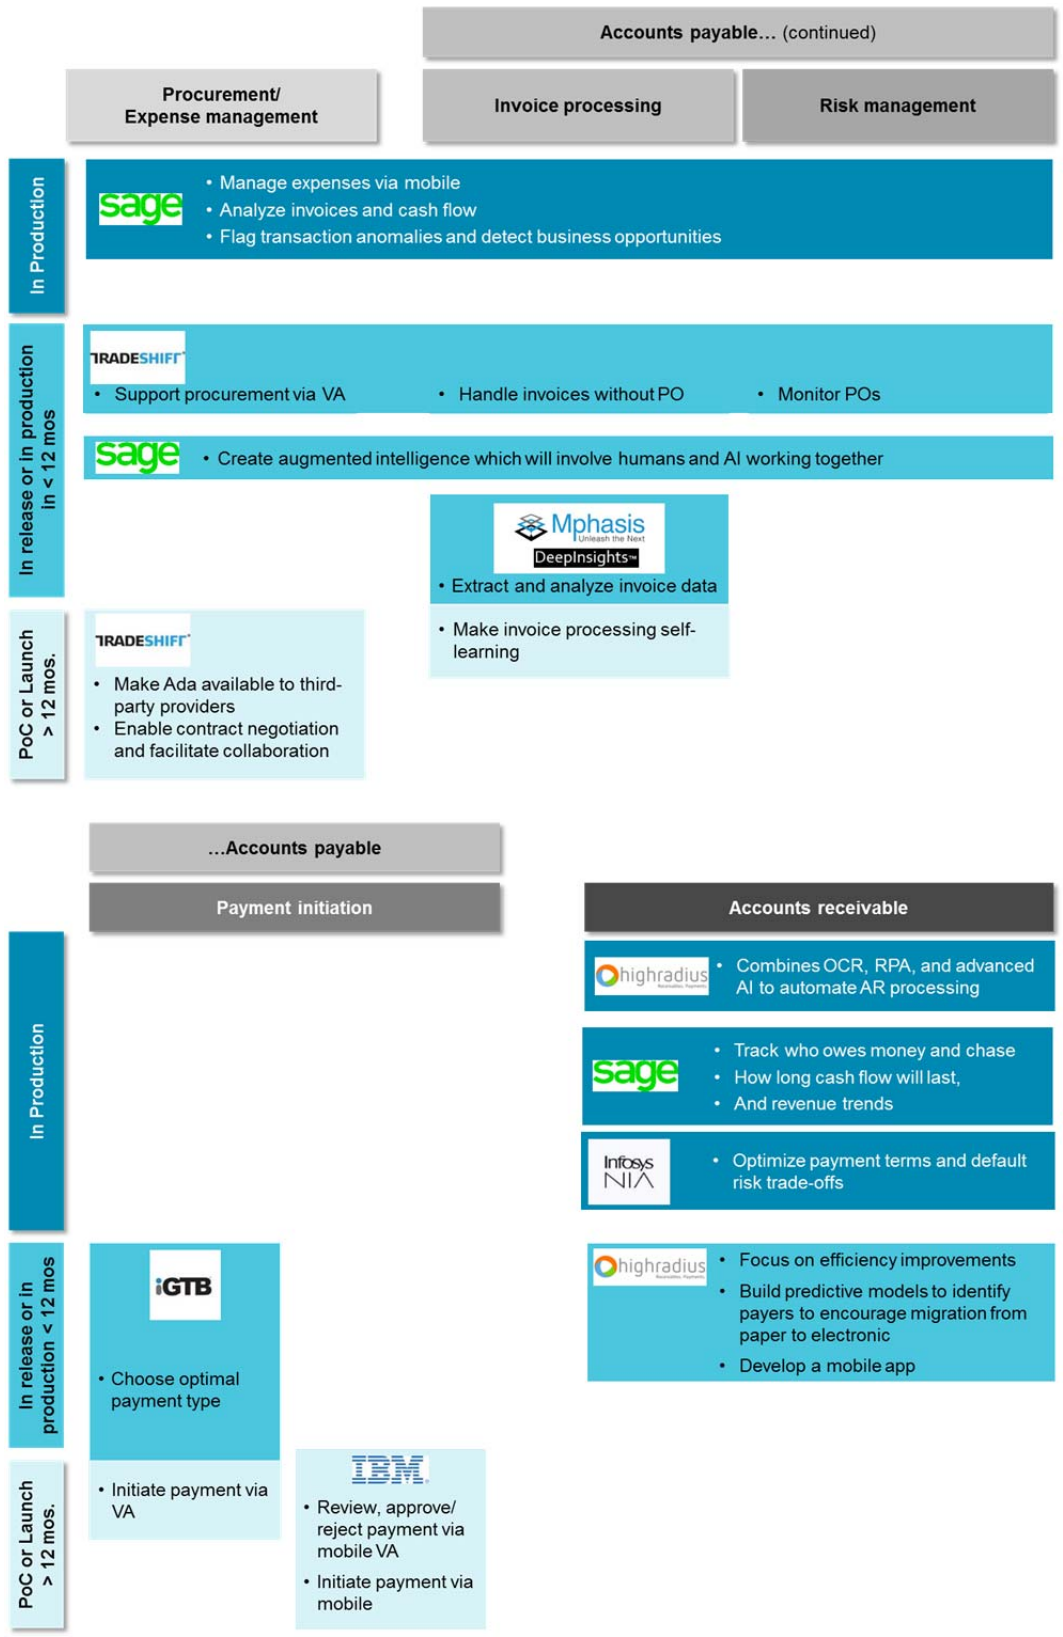✓ = pilot, proof of concept or scheduled for release in 2018

Source: Celent interviews and analysis

Figure 6: At the Sharp End in Customer Facing AI Applications— CFOs and Treasurers

Source: Celent interviews and analysis

Figure 7: At the Sharp End in Customer-Facing AI Applications — Functional Areas

Accounts payable… (continued)

	Procurement/ Expense management	Invoice processing	Risk management

In Production

sage
- Manage expenses via mobile
- Analyze invoices and cash flow
- Flag transaction anomalies and detect business opportunities

In release or in production In < 12 mos

TRADESHIFT
- Support procurement via VA
- Handle invoices without PO
- Monitor POs

sage
- Create augmented intelligence which will involve humans and AI working together

Mphasis DeepInsights™
- Extract and analyze invoice data
- Make invoice processing self-learning

PoC or Launch > 12 mos.

TRADESHIFT
- Make Ada available to third-party providers
- Enable contract negotiation and facilitate collaboration

…Accounts payable

Payment initiation

Accounts receivable

In Production

highradius
- Combines OCR, RPA, and advanced AI to automate AR processing

sage
- Track who owes money and chase
- How long cash flow will last,
- And revenue trends

Infosys NIA
- Optimize payment terms and default risk trade-offs

In release or in production < 12 mos

iGTB
- Choose optimal payment type
- Initiate payment via VA

highradius
- Focus on efficiency improvements
- Build predictive models to identify payers to encourage migration from paper to electronic
- Develop a mobile app

PoC or Launch > 12 mos.

IBM
- Review, approve/reject payment via mobile VA
- Initiate payment via mobile

Source: Celent interviews and analysis

HIGHRADIUS

HighRadius, a third party provider of end-to-end integrated receivables applications, stands out as a leader in applying advanced AI coupled with OCR and RPA to automate accounts receivable processing. HighRadius has historically sold its solutions directly to companies. In 2016, it began partnering with banks to bring its offering to the banks' customers.

At most companies, the cash application process is paper-based and hence manually intensive and error-prone, making it ripe for transformation. AR operators deal with numerous challenges from managing a myriad of remittance types received/accessed through a variety of channels to matching a payment with an invoice and dealing with exceptions (e.g., the remittance has inadequate information; the payment includes an unearned deduction; a payment has no remittance). The challenge has been growing with the rise of e-payments. While the migration to e-payments (ACH, virtual card, and wire) has benefited buyers, for suppliers, an e-payment often travels separate from the remittance, resulting in remittance retrieval and matching headaches and the need for more AR operators. As a result, growth in demand for AR automation is accelerating.

To tackle these challenges as well as others that suppliers face (e.g., in credit extension, e-invoicing, and collections), HighRadius has built an AI platform, Rivana, and a reporting virtual assistant, Freda. In this report, we focus on how Rivana is expediting and lowering the cost of the cash application process, including deduction handling. Figure 8 illustrates how HighRadius' Cash Application is integrated into a bank's operation, ingests a variety of AR-related documents, and relays relevant data to the bank customers' systems.

Figure 8: HighRadius Cash Application Automation – Bank Example
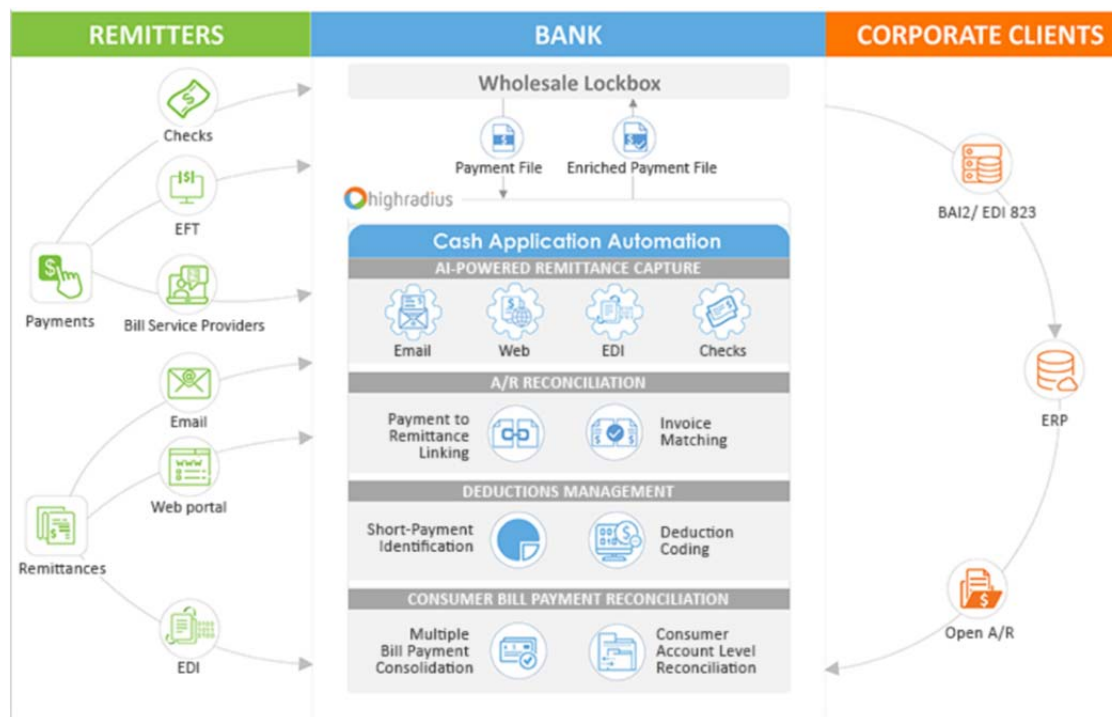
Source: HighRadius.

HighRadius Cash Application has an automation toolbox comprising OCR, RPA, and advanced AI. OCR is used to digitize paper documents. RPA is used to pull remittance data from payers' web portal, to massage data (e.g., append or remove leading zeros as needed), and to append data from the AR system (e.g., store number on deductions). Machine learning is applied in three areas: digitization of remittance detail, automation of exceptions handling, and remittance generation. First, machine learning is used to capture and normalize remittance information from a variety of sources (e.g., email body, email attachments, and check images in various formats — structured and unstructured) and eliminates the need to build and maintain templates.

For deduction handling, HighRadius uses a rules-based approach to identify and code a deduction and machine learning to predict the probability of a deduction being valid or not, which helps an AR operator handle the exception faster.

The effectiveness of HighRadius Rivana is seen in the sharp improvements in the average remittance data lift rates from emails: 77% three years ago to 92% today.

In 2018, HighRadius plans on focusing on efficiency improvements and on building predictive models to identify payers to encourage migration from check to ACH and from paper invoices to e-invoices. It is also developing a mobile app for payments and invoices.

IBM WATSON

IBM Watson Conversation provides AI tools to implement natural language interfaces between a bank and their customers; that is, customers can chat or talk (via speech) to their bank. A bank can add Watson Conversation APIs to its existing bank applications. Unlike some of the simpler request/response chatbots, Watson Conversation can remember context and carry on a bona fide conversation. In addition, Watson Conversation allows a bank to give its virtual assistant a personality, which distinguishes the way it responds.

Figure 9 outlines the high-level architecture of Watson Conversation for Payments. End users can use a variety of channel interfaces. APIs call Watson to make a variety of "translations" (e.g., speech to text and vice versa) to enable Watson to determine the end user's intent and respond accordingly. Responses could be voice/text or a graphic. Graphical answers are generated by Watson Analytics. It examines the data provided by the user (data can be uploaded to the cloud or be relayed via a direct connect to an on-premise data source), predicts ways to interpret the data, and displays the results in a graphical format. Users have the option to see customized graphs. A cognitive orchestrator built by the bank connects to the bank's payments services hub to execute the payment.

Source: IBM

In the corporate banking context, treasurers could use Watson Conversation to respond to fraud alerts, check on the status of their cash position, and approve large payments that require additional levels of approval. For example, when IBM Safer Payments detects a fraud, a text alert could be sent to the treasurer. Then Watson Conversation could be used to allow the alert recipient to take action using natural language. Other examples include a small business contacting customer service or filling out a loan or mortgage form in a simpler question-and-answer format. In addition, business users can type in their questions and Watson will interpret them and respond with graphical answers. Users are able to customize the graph or ask the question differently, resulting in a different graph.

The results have been positive: one bank which put Watson Conversation into production to handle common retail customer service questions reduced the number of incoming calls to the call center 50%.

IBM Watson Conversation is currently in production in a wide variety of industries (e.g., insurance and wealth management) and is starting in banking (e.g., RBS is using chatbot technology and Watson to support its consumer contact center agents). IBM is continuing to enhance Watson Conversation tooling and also putting together "Developer Journeys" on the IBM Cloud for Financial Services to teach developers how to build their own Watson Conversation chatbots. A promising harbinger of payments initiation is IBM's successful proof of concept executing a Zelle payment via speech in Q3 2017.

iGTB

iGTB (Intellect Global Transaction Banking, part of Intellect Design) is embarking on delivering next generation customer engagement by rebuilding its entire platform to be API-first and AI powered. Figure 10 contrasts a traditional transaction-centric platform with Intellect Contextual Banking eXperience (CBX) Platform. iGTB embarked on its replatforming journey two years ago with the vision of adding an AI layer whose goal is to understand what the end user's intent is, develop recommendation logic, and apply intelligent routing. The user interface has no logic but rather has microservices which make calls via APIs.

Figure 10: iGTB Contextual Banking eXperience Platform Compared to Traditional Platform
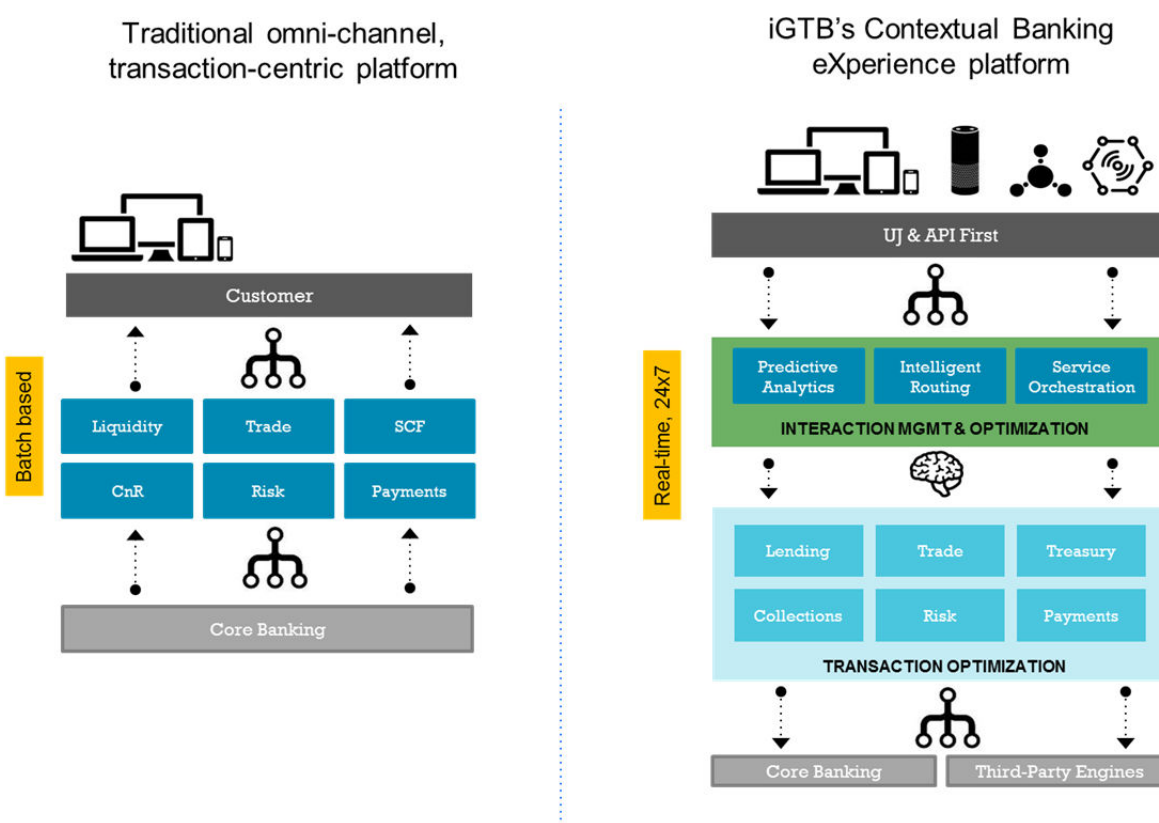
Source: Intellect Design
Note: UJ = user journey

Figure 11 shows the iGTB dashboard which contrasts sharply with traditional cash management dashboards. At the top, it calls out the components of beginning and ending cash balances. In the middle, it charts out expected changes in balances over a week. On the right-hand side, it provides alerts regarding action items and not only makes recommendations but also enables the end user to take action at a click of a button. Machine learning will power recommendations regarding FX exposure hedging, liquidity optimization, liquidity repair (shortfall), optimal payment type, and in the case of a payment shortfall, options to cover the shortfall.
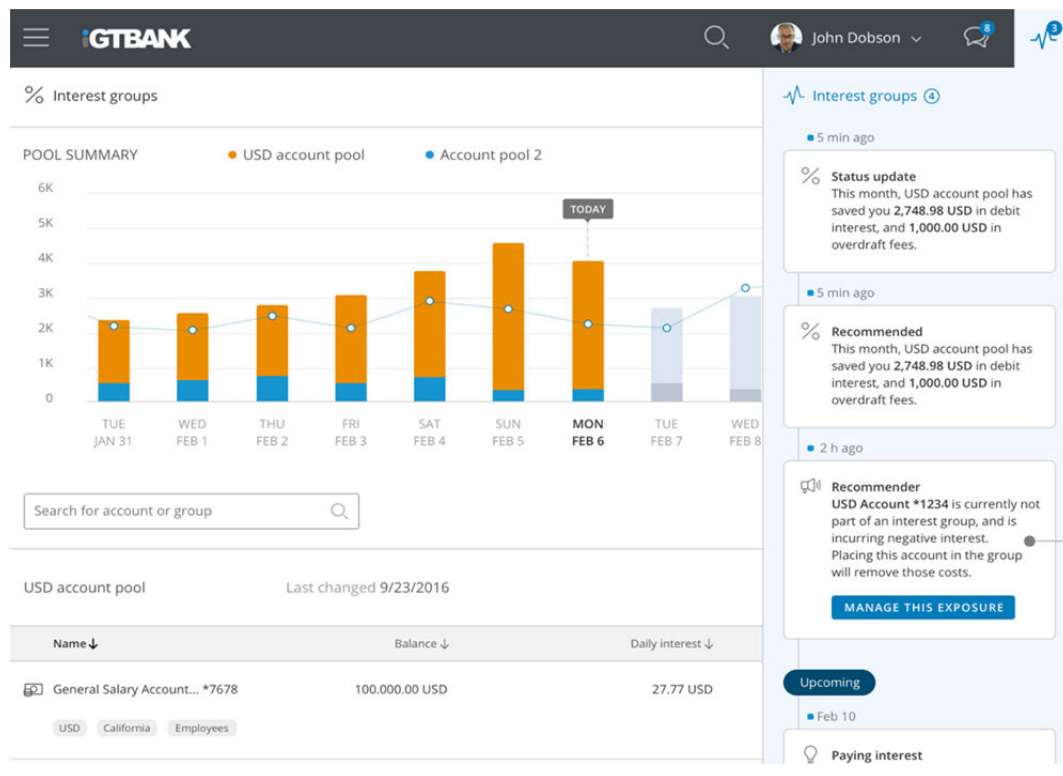
Source: Intellect Design

For payment routing optimization, five metrics are gauged: speed, cost, risk/trust, traceability, and revocability. Risk is the most complex to measure and takes into consideration who, what, and why. For the "who," it examines the type of transaction (e.g., cross-border cross-currency, B2C, or B2B) and counterparty attributes, and who executes the payment (e.g., human, application, or device). For the "what," it determines the context, including the nature of the goods or services being purchased and, if there is a trade agreement, the terms of sales and procurement, including any expected warrantees or financial instruments (e.g., letter of credit). The "why" is the frequency of the payment (e.g., recurring or one-off) and the purpose (e.g., payroll or vendor). These risk factors drive recommendations such as chose a revocable payment type, set up a letter of credit, and/or hedge your FX exposure. In addition, if there are insufficient funds in an account to make a payment, the platform makes suggestions regarding how to cover the shortfall, such as, partial payment, transfer funds from another bank account, make an intracompany loan, and/or draw a credit line.

To demonstrate how its platform can readily plug into an open virtual assistant, iGTB built a demonstration with Alexa and had a couple of Alexa developers train Alexa to understand transaction-related questions which a treasurer handles every day. Figure 12 is a transcript of the dialogue between a CFO and CBX18 for Alexa.
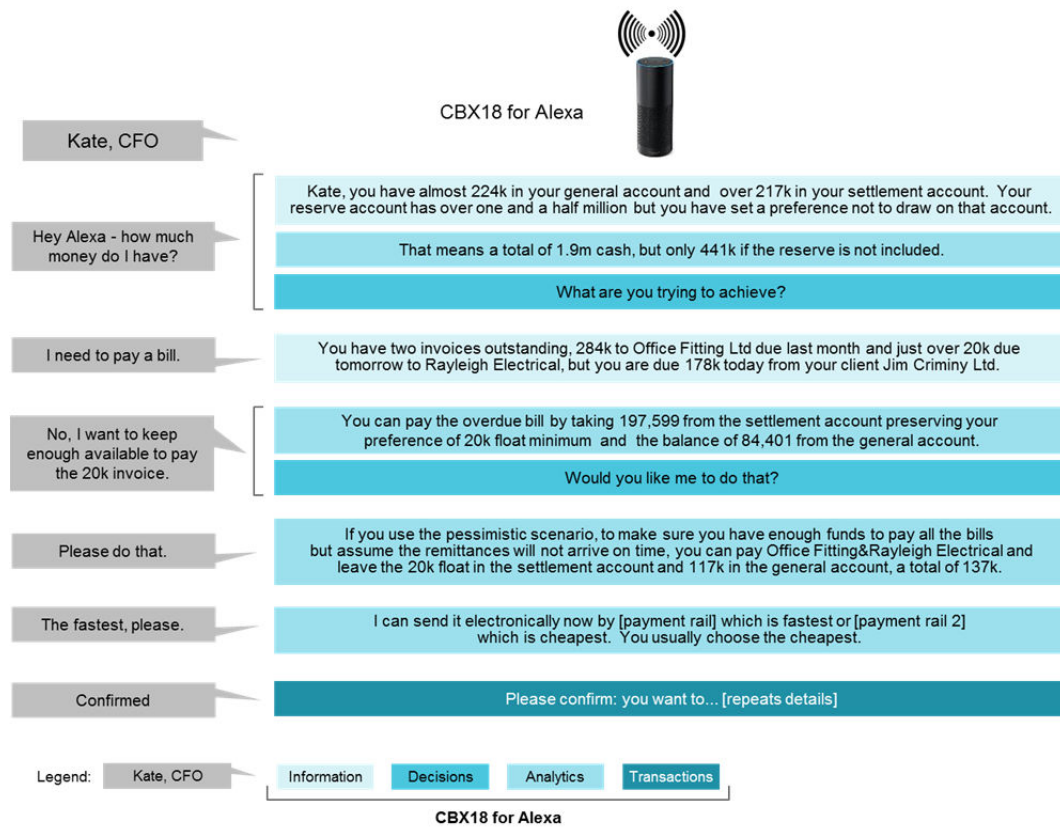
iGTB plans to be in production with CBX by this summer, with the Alexa app following two months later. Initially, the Alexa interface will allow inquiries only. Once Alexa for Business authentication SDKs are more mature, it will allow for payment initiation.

On its roadmap for 2018 are:

- Cash flow and working capital forecasting based on limited or no data available from ERP or accounting integration.

- Configurable real-time cross sell capabilities based on next best offer (e.g., prequalified invoice factoring or loan in case of a payment shortfall).

- Further simplification of recommendations geared toward the SME segment.

INFOSYS NIA

Infosys NIA is an AI platform which collects and aggregates data into a self-learning knowledge base and then automates repetitive business and IT processes. NIA can be applied to a number of use cases in a company's functional areas as well as commercial/corporate banking, including working capital advisory services, prospecting, cross-selling and retention (next best product, product migration, and attrition prediction), and credit. To customize NIA applications to specific industry/customer segment needs, it has subject matter and functional experts that bring critical conceptual understanding to the build, test, and implement processes. Figure 13 outlines the NIA architecture at a high level.

Figure 13: NIA Architecture
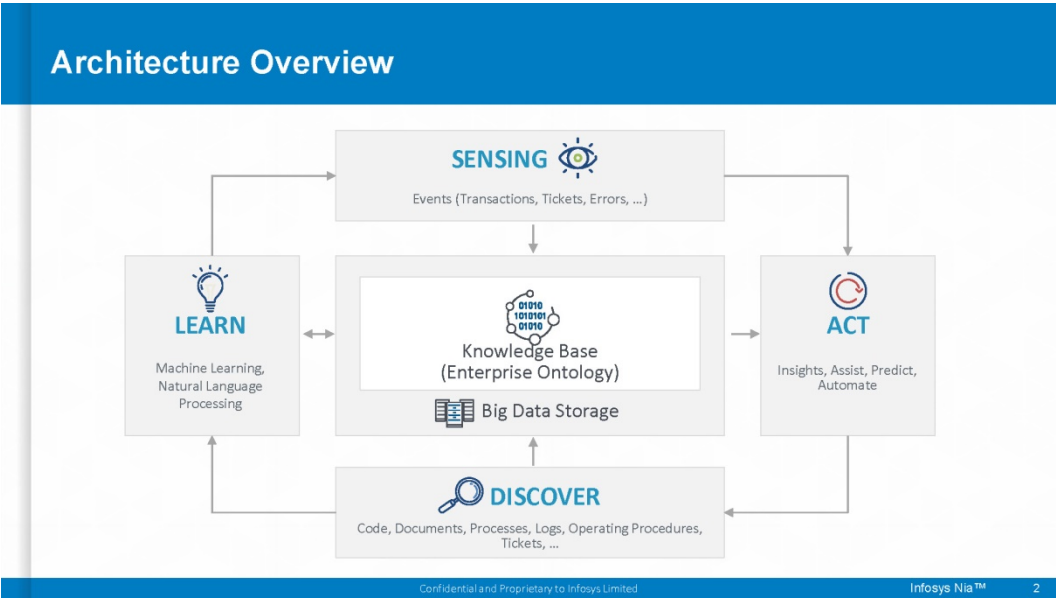
Source: Infosys

NIA differentiates itself based on its ability to facilitate and expedite model iterations. Typically testing, iterating through machine learning models, consumes significant programming time. Once a data set has been defined, NIA can automatically generate different model types. For each model, NIA produces a receiver operating characteristic (ROC) curve which displays the predictive strength of a model based on true positives (in machine learning parlance, probability of detection) compared to false positives (probability of false alarm) as shown in Figure 14. The greater the area under the curve, the stronger the model is.

Figure 14: Example of NIA-Generated ROC Curves
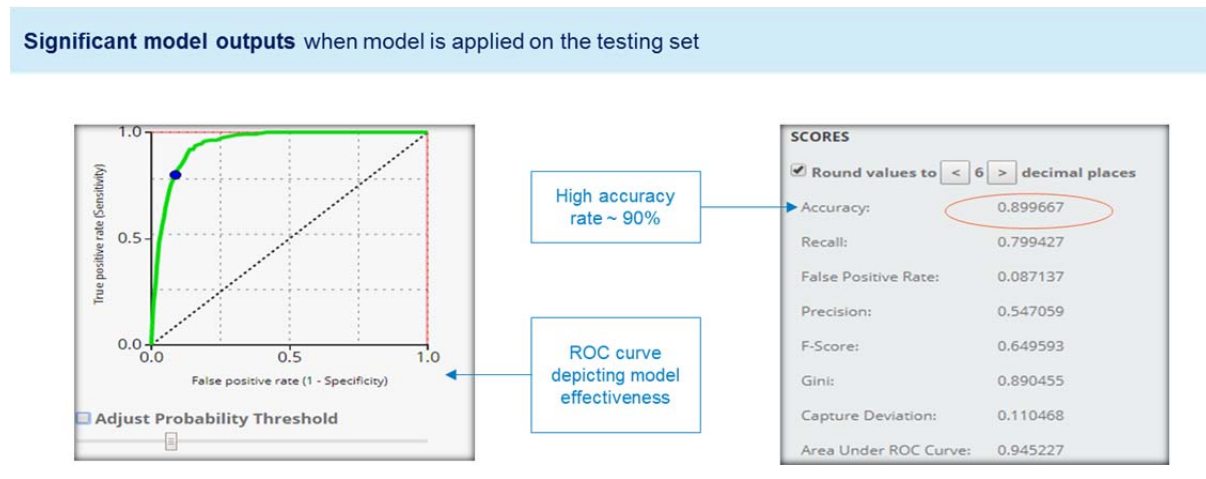
Source: Infosys.

NIA has been applied in a corporation's accounts receivable to benchmark payment terms and gauge default risk. As a starting point, NIA built a model based on both internal data (e.g., payment terms, historical payment patterns, default amounts, and risk ratings) and external data (e.g., financial ratios, cost of capital, and country risk) which generated receivables exposure ($) by default risk clusters. Next, it was used to analyze tradeoffs between payment terms and default risk and answer critical questions as outlined in Figure 15. NIA was used to detect patterns in days sales outstanding by payment terms and default exposure, patterns in payment terms, and country/product risk ratings. The output is used to determine how to modify terms to optimize working capital while mitigating risk.

Figure 15: NIA Applied in AR to Optimize Payment and Default Risk Tradeoffs
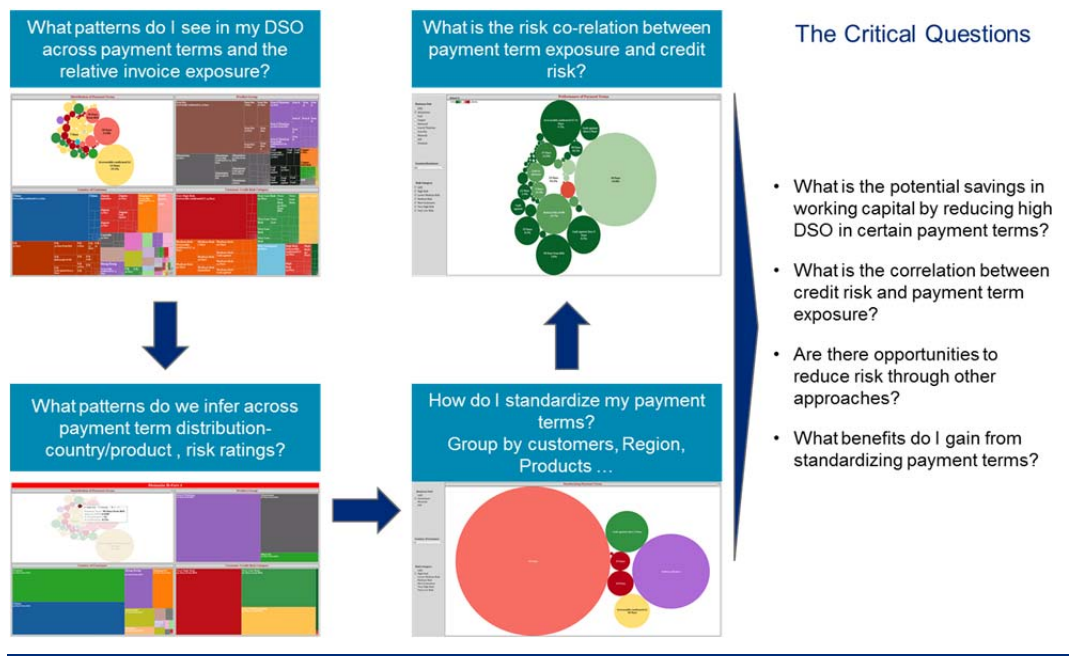
Source: Infosys.

As discussed in the next section on AI supporting relationship managers, NIA can be used by banks to provide advisory services to their customers not only in AR but in working capital optimization overall.

MPHASIS DEEPINSIGHTS

Mphasis is building a cognitive computing platform, DeepInsights (patent pending), which comprises smart data ingestion, cognitive analytics, visualization, and interactions as outlined in Figure 16. Mphasis has developed domain expertise in three corporate and corporate banking use cases: invoice processing, financial statement analysis (credit underwriting), and trade finance document verification (the latter two are discussed in the Credit and Trade Finance section).

Figure 16: DeepInsights Overview — From Smart Data Ingestion to Visualization and Interaction
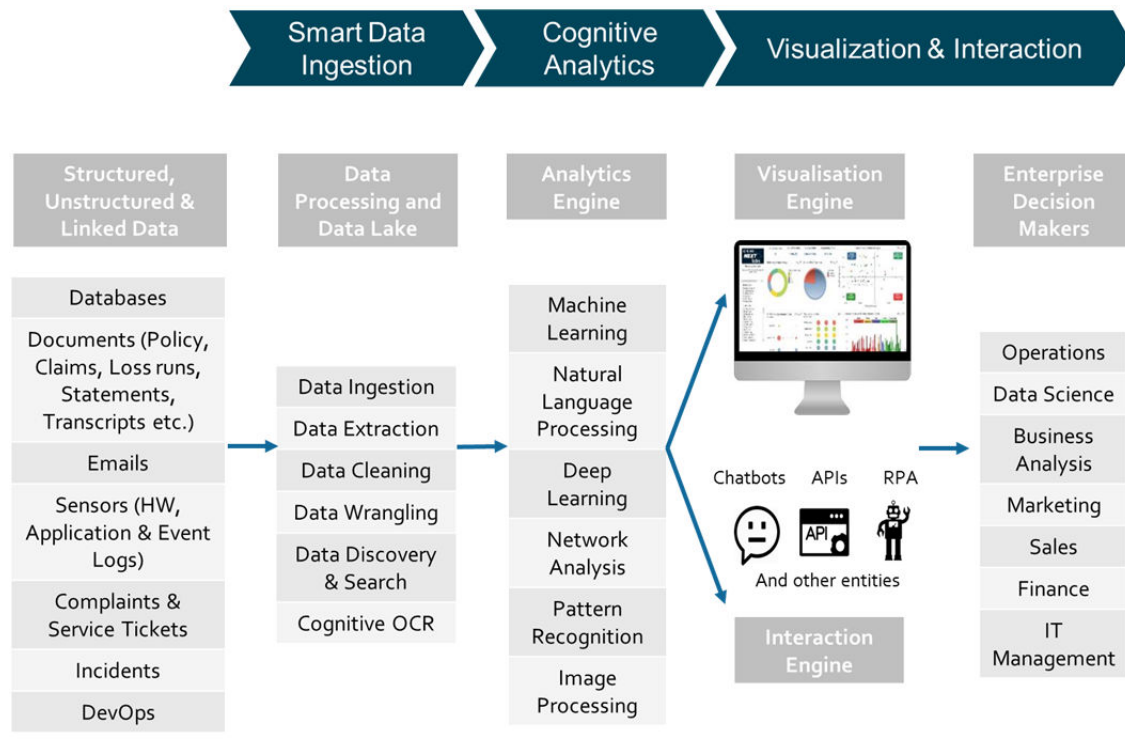
Source: Mphasis

As part of its smart data ingestion, it has developed cognitive OCR, which uses advanced image processing and machine learning methods to pre-process images and to intelligently identify characters, words, tables, logos, check boxes, and other useful information in the image.

For its analytics engine, Mphasis is building pretrained proprietary models based on select client data. It then customizes these models based on individual client data. Most of its clients are in the top 15 of their industry and usually only have to supply 15 days of data to train the models. It is developing visualization and interaction engines which are user interface agnostic and can deliver information and data charts to a variety of end points, including robotic process automation systems.

Mphasis has found a prime use case in invoice processing, which is typically fraught with inefficiencies, including handling paper or PDF invoices and manual matching to purchase orders and goods received. The automation challenge is compounded by the fact that a myriad of invoice formats exist and often contain unstructured data. To automate the process, DeepInsights applies its cognitive OCR and analytics engine to convert paper/PDF into a machine-readable format and interpret the data. The results of recent pilots for a European auto manufacturer and the largest healthcare company in the US are compelling: 95% accuracy of data extraction, 90% reduction in time to digitize an invoice (1 minute from 10), and 40% decrease in invoice processing costs.
On the 2018–2019 horizon, Mphasis plans on extending DeepInsights to making invoice processing self-learning using reinforcement learning.

SAGE

Sage (a leading accounting/ERP software provider) is in the vanguard when it comes to leveraging advanced AI to simplify a business's financial management. It was the first of its peers to launch a virtual accounting assistant, Pegg (July 2016), which automates and expedites financial management via mobile devices. Demonstrating its value, within six weeks of launch, Pegg had users in 135 countries. Internally at Sage, Pegg has already answered over half a million queries related to customer support and is continuously improving its knowledge by learning and observing human support agents.

Sage has wisely taken an open approach to encourage adoption. While Pegg can be used with Sage Accounting (formerly Sage One), it can also be used stand-alone and hence could act as entry product for Sage Accounting. In addition, Sage has connected Pegg to popular communication platforms including Salesforce and Slack. In 2018, Sage plans on integrating voice technology. Using Pegg is as simple as adding Pegg as a contact; there is no app to install or login credentials to enter.

Currently Pegg is focused on supporting accounts payable and accounts receivable. It uses machine learning to detect end users' expense patterns and make personalized recommendations. The more conversations Pegg has with end users, the smarter it gets and the better it can determine how to support its end users. Among the most popular use case is expense filing and management. Other use cases include invoicing and AR management. Figure 17 displays a common conversation.

Figure 17: "Conversation" with Pegg
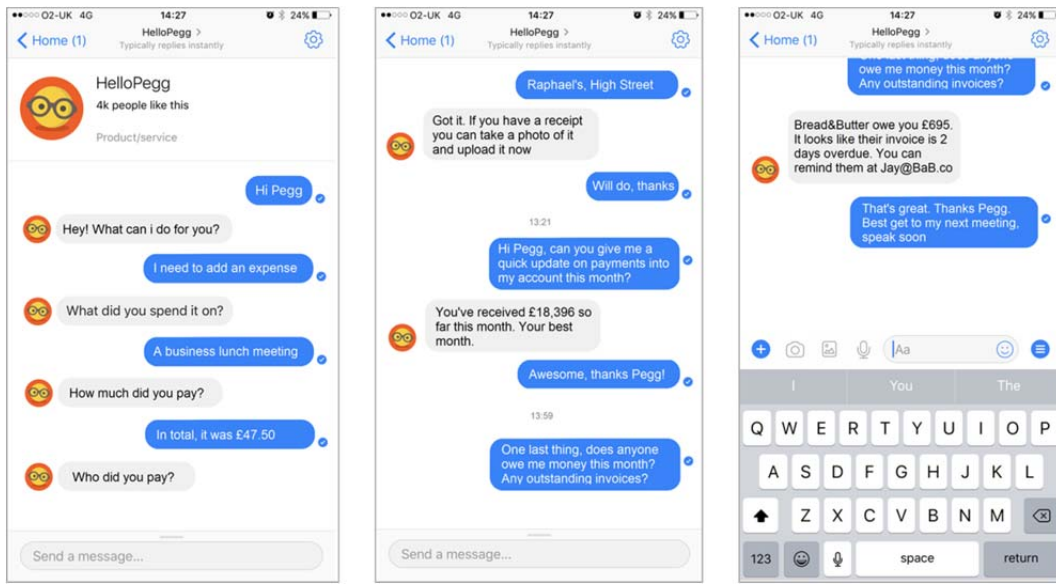
Source: Sage

Sage plans on extending Pegg's capabilities in 2018, creating augmented intelligence which will involve humans and AI working together based on their comparative advantage. For example, in seconds, AI can give accountants an analysis of millions of transactions, detect anomalies, and automate intelligent alerts. Armed with this data and take-aways, humans can then deliver their strategic input.

TRADESHIFT

Tradeshift, a cloud-based business network providing procure-to-pay solutions, is approaching the summit for virtual assistant supporting procurement and accounts payable. It plans to launch Ada, a conversational interface, across its platform in Q3 2018. Ada sits on top of the AI layer that Tradeshift is building, as outlined in Figure 18. Ada responds to questions by drawing upon any relevant data sources and workflows on the platform, and running this data through a variety of algorithms to arrive at an answer. Ada can also highlight relevant parts of the user interface for the end user to take action based on the answer. It can also connect to Slack and act on behalf of an end user (e.g., to ask for approval of a purchase order).

Layers					
UI	Collaboration		Tradeshift Apps		Customer & 3rd party apps
Conversation	Intelligent NLP				
AI	Transaction Analytics + AI	Product Ontology + AI	Graph AI	Organizational Ontology	Decision AI
Services	Transaction Engine	Product Engine	Network Engine	Org Structure Service	Workflow Service
Data	Documents & Transactions	Products	Companies & Connections	Users	Decisions

Source: Tradeshift

Ada will be able to handle a variety of common use cases. For example, across the financial supply chain, it will be able to support vendor selection, purchase order monitoring (did this person make the same order two weeks ago?), invoice processing (how should I code this invoice which has no purchase order? Can you resolve this PO/invoice matching error?), and payment (can you get me a card for this purchase? what's the optimal time to pay this invoice?).

Tradeshift is planning to make Ada available to third party providers which offer value-added services in marketplaces, procurement, invoicing, and payments, including solution providers to specific verticals and which offer specialized solutions (e.g., for direct spending or strategic sourcing). In addition, it is building machine learning algorithms to provide network-based trust/risk information to the buyers and suppliers on the Tradeshift network. It expects its bank partners to begin testing Ada's capabilities in 2018.

Further out on the horizon, Tradeshift envisions buyers and suppliers both using AI engines to negotiate contracts and facilitate collaboration. Parties on Tradeshift will be able to negotiate contracts on the network itself, and the platform will enforce them. Machine learning could be used in aligning and mapping these new contracts based on agreements currently in place with a specific client. In addition, machine learning could provide insights and highlight opportunities for improvement and collaboration both within a company's supply chain and outside of it. For example, it could highlight which supplier on the Tradeshift network is best suited to handle a particular need. Then built-in network trust gets a relationship going relatively quickly, while the live collaboration makes that relationship productive, and the on-platform agreement helps to keep it on track.

FRONT OFFICE: RELATIONSHIP MANAGERS

Advanced AI is being applied to improve the tools available to relationship managers (RMs) and marketing. RMs could be more proactive in customer service when provided "next best action" and "next best offer" (NBA/NBO) recommendations generated by a decisioning system with predictive and adaptive analytics. Currently solutions are at camp 1, providing RMs with basic NBA (e.g., customer is approaching credit or covenant limit, here are options to avoid hitting the limit) with camp 2 coming into view and including basic alerts regarding changes in behavior (e.g., drop in payment transactions month-over-month). Further out, camp 3 is prospecting, and the summit is sophisticated NBA based on sentiment analysis. In addition to RMs, marketing could improve segmentation accuracy and targeting through machine learning models (e.g., clustering algorithms).

Figure 19 and Figure 20 outline the leaders in advanced AI applications supporting RMs. Pegasystems is a clear leader here, having invested in enhancing its customer relationship management (CRM) and Client Lifecycle Management (CLM) solutions with advanced AI. Note that iGTB CBX will deliver next best action and offers directly to the client, which is covered in the prior section. IBM Algo Credit Manager v5.4, which includes sentiment analysis which RMs can use to monitor clients, is covered in the Credit section (page 33).

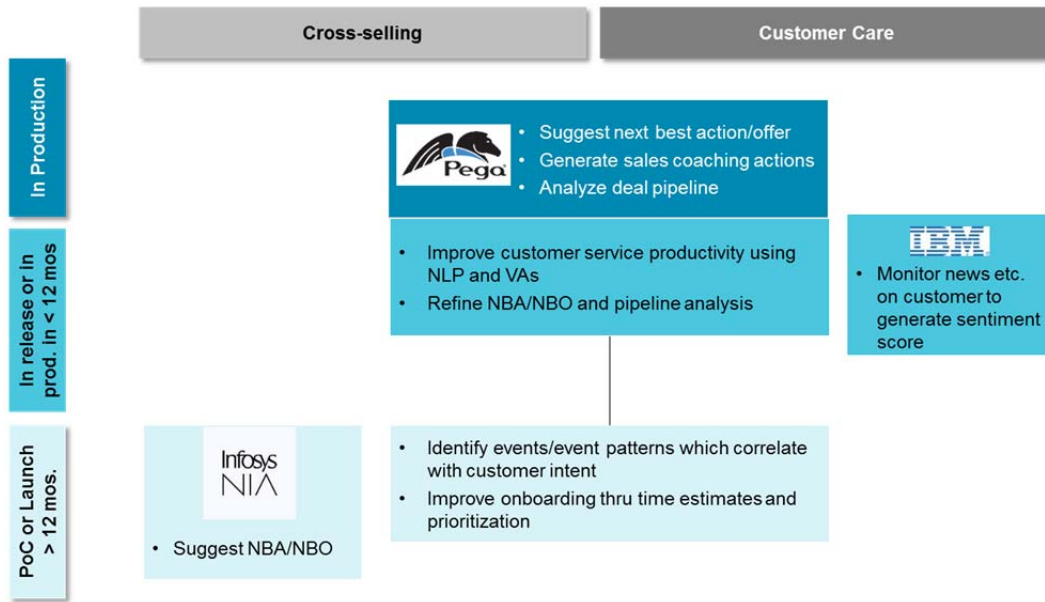Figure 19: At the Sharp End in Relationship Manager Support — Overall

Company	Cross-selling	Customer Care
Pega	✓✓✓	✓✓✓
IBM		✓✓[1]
Infosys NIA	✓	✓

Legend: ✓✓✓ = in production
✓✓ = in implementation or released but not installed
✓ = pilot, proof of concept or scheduled for release in 2018

Source: Celent interviews and analysis
1. IBM's solution is covered in the Credit section.

Source: Celent interviews and analysis

INFOSYS NIA

Infosys NIA (which is overviewed on page 19) has several use cases in supporting commercial and corporate bankers. It can be used to identify NBA/NBO for RMs and to support a bank's advisory services in the area of working capital optimization and risk mitigation. It is currently working on a few early stage proofs of concept in corporate banking at top tier banks in the US, Europe, and Asia. For example, to increase cross-selling and retention for a US-based global bank, Infosys is using clustering and classification techniques on prospective and existing client buying behavior data and financial data (e.g., working capital ratios) and identify NBO for each cluster. In addition, it is working on modeling product migration and client attrition.

In the area of advisory, it recognizes that banks have both more expertise and experience in managing credit and payment terms risk than their customers and hence can be useful advisors. The challenge for banks is how to differentiate their advisory services and make them more impactful. One potential means to differentiate is through superior data analytics powered by advanced AI as discussed above in the accounts receivable context (page 20).

PEGASYSTEMS

Pegasystems has been developing and incorporating advanced AI into its CRM and CLM solutions. In its CRM, it is developing AI-driven NBA/NBO recommendations, customer attrition likelihood alerts for RMs, sales coaching, and pipeline analysis. It is implementing machine learning (specifically predictive and adaptive decisioning) and NLP into its Customer Decision Hub (CDH), which is a core part of its CRM solution, which supports marketing, sales, onboarding, and servicing. As part of its CLM application, Pega is developing AI use cases to provide product recommendations, estimate onboarding times, and estimate customer lifetime value.

Figure 21 shows the CRM dashboard for a sales executive and the key components: NBA/NBO, sales coaching, and pipeline analysis. To generate NBA/NBO, CDH analyzes customer data, customer communications, and transaction data. It processes events in real time to catch critical opportunities (e.g., customer service call, unusual transaction behavior, deadline approaching to close a deal, document missing). When implementing CDH's machine learning capabilities, Pega undertakes A/B testing, comparing its predictive results with those from the bank's predictive models. Pega is drawing upon its learnings from implementing NBA/NBO in consumer banking to extend to commercial banking (e.g., as it is doing with RBS).

Figure 21: Pega Sales Automation Dashboard
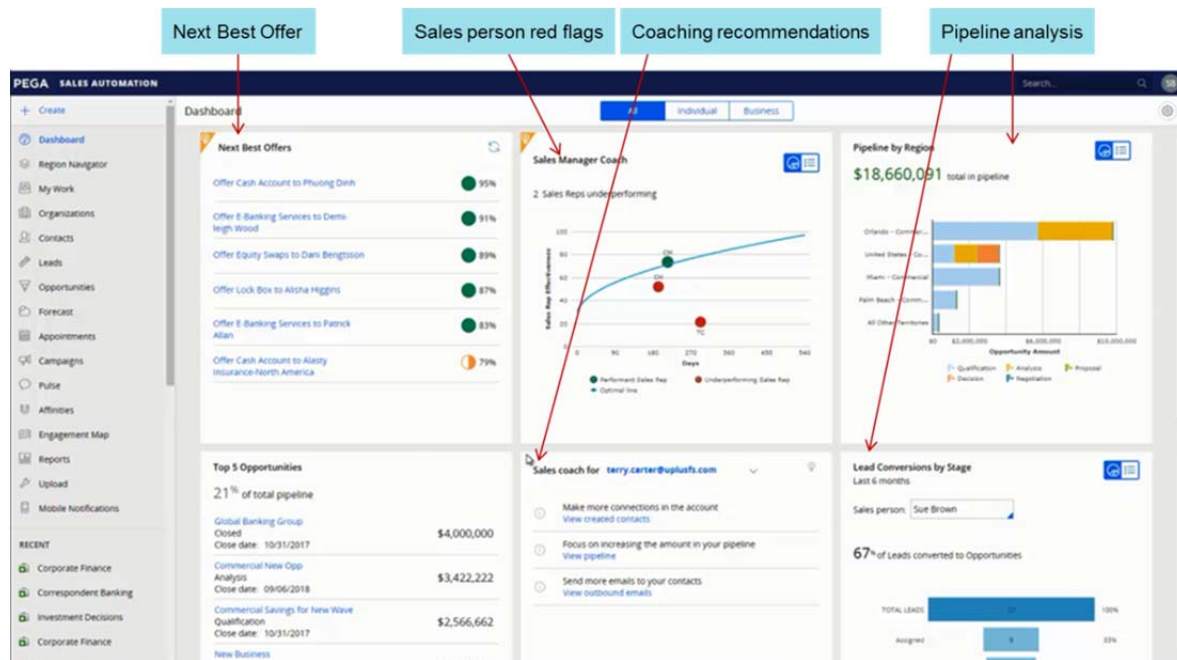
Source: Pegasystems

Pega is a pioneer in AI-based sales coaching. Its adaptive model predicts the success of a new sales rep within a defined time period and charts the optimal performance trend line with continual updates. Beginning on the sales rep's 90th day from hire, the model predicts effectiveness, with the definition being configurable by the bank. The model has 50 predictors pulled from sales and human resources data and is trained using actual sales rep performance data. In addition, the system helps the sales executive coach a sales rep (for example, to close a deal faster).

For pipeline analysis, given the relatively smaller dataset, Pega uses predictive analytics to estimate the probability of a deal moving to the next stage, of ultimately winning the deal, and the most likely quarter based on current sales activity and benchmarks for similar deals.

Building a data-driven, action-oriented dashboard with a friendly user interface for RMs is only half the challenge. The other half, convincing them that it is useful, can be the greater challenge. Commercial and corporate banking are nuanced relationship-driven

businesses, and the best RMs pride themselves on not only their subject/client matter expertise but also their intuition. To convince RMs of the credibility of the AI-based recommendations and make them explainable, Pega has developed a visualization tools. Figure 22 shows the "radar chart" which maps the predictors that impact a sales representative's performance. The color of the predictor icon indicates good (green), neutral (gray), and bad (red).

Figure 22: Pega Makes AI Transparent — Sales Manager Coach
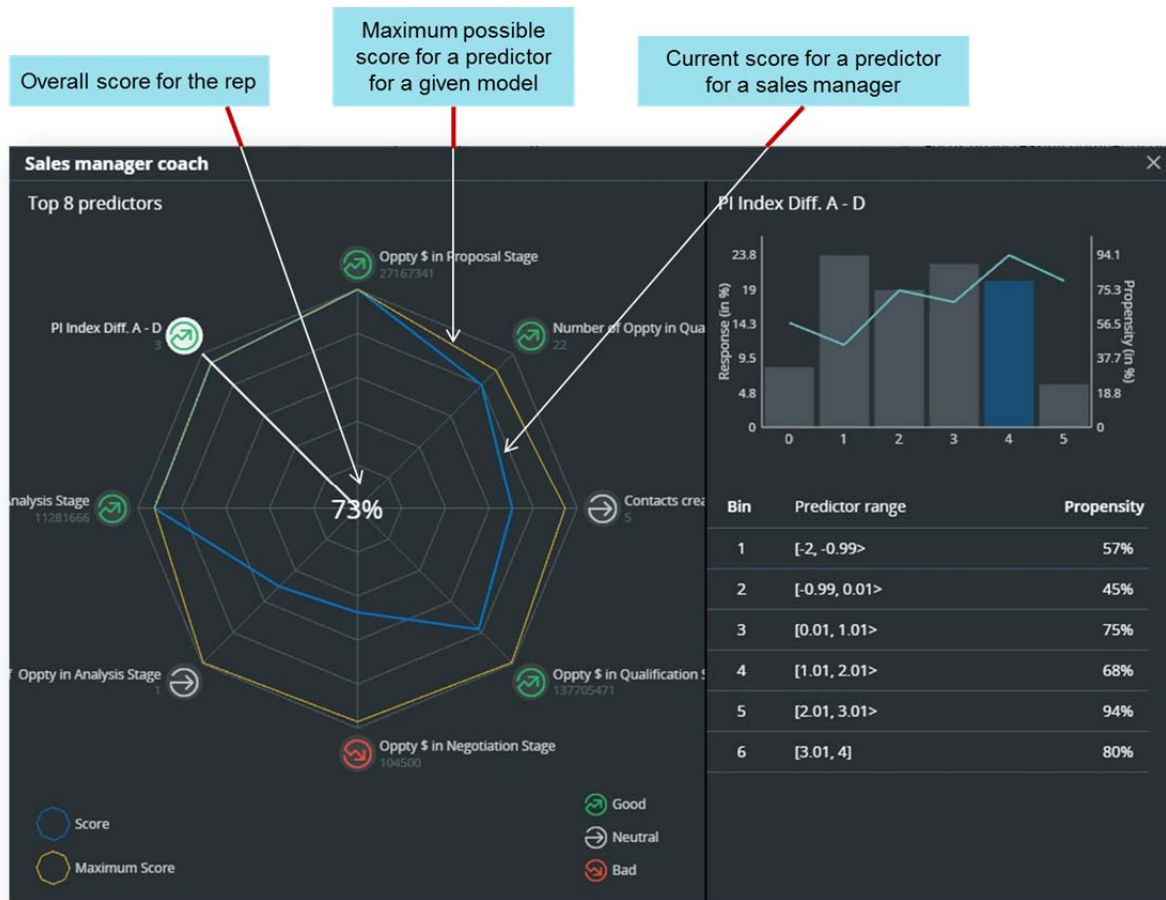
Source: Pegasystems

Currently, Pega is working on increasing customer service productivity. It is exploring ways to apply NLP to text or voice-to-text across a variety of channels (e.g., IVRs, chat/text, and email) and use a virtual assistant to respond when appropriate and/or smartbot to route queries to the right team for resolution.

For 2018–2019, Pega will continue to refine its NBA/NBO models to, for example, identify gaps and inefficiencies in the underlying components of a NBA strategy, and then recommend alternatives likely to lead to higher responses. In addition, Pega is working on identifying the events and event patterns which have a high correlation with customer intent, and dynamically integrating the attendant insights into decision strategies without

human intervention. For its pipeline analysis, it is working on grading deal attractiveness based on risk tolerance and relationship value to name a couple variables. In its CLM application, it is exploring the use of AI to increase onboarding time transparency for the customer (a model estimates onboarding time), to help prioritize onboardings based on estimated customer lifetime value, and enhance product selection through product recommendations.

CREDIT AND TRADE FINANCE

Key Research Question

3

Who's summiting in the middle and back office?

Numerous mountaineers are reaching challenging peaks, including credit underwriting and compliance.

Credit and trade finance exhibit a variety of pain points which can be solved by advanced AI coupled with traditional automation technologies. Both are plagued by a myriad of paper inputs and manual processes. In addition, credit relies on data analytics, which is a prime candidate for machine learning. Small business credit, in particular, with its relatively high origination costs, has attracted significant advanced AI investment (for details, see the Celent report, *Small Business Credit: Exemplars of Innovation*, September 2017). The summits that can be reached are a function of numerous factors ranging from the models themselves (e.g., loss prediction strength and adaptability) to a bank's willingness and ability to let a machine run part or all of the process instead of a human.

On the credit side of the business, standard underwriting and approval processes are excellent candidates for machine learning because they involve consistent, repetitive processes. Machine learning can improve small business credit scoring and pricing through its ability to process both internal and external structured and unstructured data (as credit-oriented fintechs have been touting) and adapt to changing credit environments. Moreover, it is showing potential in loan collection strategies. In larger size commercial credit, AI is being put to effective use in contract review and credit monitoring.

In trade finance, AI is being used to reimagine Byzantine document workflows across multiple internal and external data sources and automate manual processes and tackle fraud and compliance challenges, in particular, AML and sanctions screening (compliance is covered in the next section). The summit is digitizing the myriad of documents, including financial documents (bill of exchange), commercial documents (invoice, packing list), shipping documents (bill of lading, airway bill), insurance documents and validating and matching them per a variety of rules. Not only are the paper, manual processes expensive for banks, any delay in identifying any documents that are not compliant with the terms puts the bank at risk of financial loss.

Figure 23 and Figure 24 outline the leaders in AI applications in credit and trade finance and show a vibrant competitive field with clear plans over the next two years. Figure 23

lists providers in order of maturity level, and Figure 24 lists them alphabetically within maturity level.

Figure 23: Mountaineers in Credit and Trade Finance – Overall

| Use Case / Company	Credit			Trade Finance (document processing)
	Sales	Application/ Underwriting	Booking & Servicing	
NUMERATED	✓✓✓ (small business loans – end2end)	✓✓✓	✓	
genpact		✓✓✓ (parts of the process)	✓✓✓	
WorkFusion		✓✓✓ (parts of the process)	✓✓✓	✓✓✓
previse	✓✓ (supply chain financing – end2end)	✓✓	✓✓	
Mphasis		✓ (parts of the process)		✓
IBM			✓✓	
wipro holmes				✓

Legend: ✓✓✓ = in production
✓✓ = in implementation or released but not installed
✓ = pilot, proof of concept or scheduled for release in 2018

Source: Celent interviews and analysis
Note: Oracle has Trade Finance module but it currently does not have any advanced analytics. The next release will include OCR and text analytics.

Figure 24: Mountaineers in Credit and Trade Finance
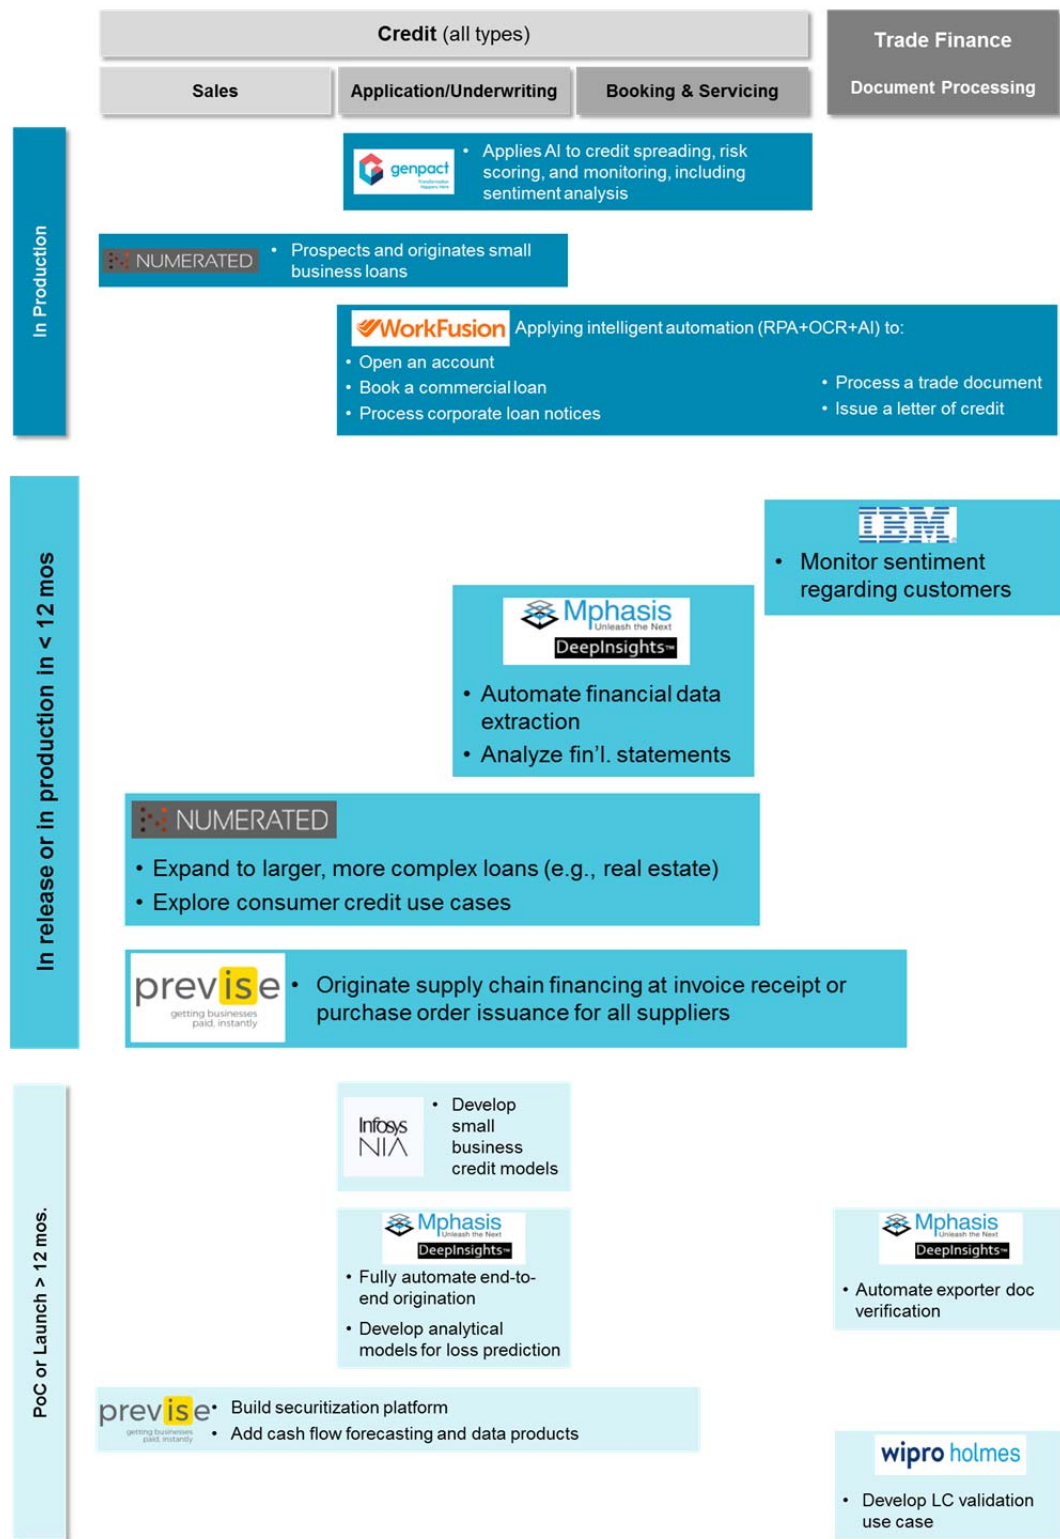
Source: Celent interviews and analysis
Notes: For information on Infosys NIA, see page 19. It is currently working on an early stage proof of concept in the area of small business credit. For information on Wipro, which is planning to apply HOLMES to documentary trade finance, see page 52.

GENPACT

Genpact (a global professional services firm) acquired Rage Frameworks in March 2017 and has incorporated Rage's AI into its AI platform, Genpact Cora, launched in July 2017. It has two AI powered solutions related to commercial lending: Credit Assistant for underwriting and credit risk monitoring and Cora LiveSpread for analyzing client financial data for underwriting and risk management. Credit Assistant has been battle-tested by eight North American banks and three banks in other regions.

Credit Assistant was developed by Rage Frameworks, which was founded 10 years ago based on a vision that any knowledge process could be automated by configuring building blocks of software. It began by tackling the credit underwriting process (its first incarnation was as ecredit.com). Early on, it found that much of the data in enterprises today is unstructured (e.g., contracts and footnotes in financial statements). To normalize this data, it uses NLU models (computational linguistics) combined with patented extraction capabilities. It has become a leader in applying contextual and traceable deep learning, which is paramount for bankers which need to be able to explain the "why" of their model outputs.

Credit Assistant includes a proprietary credit rating methodology, which runs multiple machine learning models to arrive at a consensus. In addition, it uses input from its credit monitoring system, an early warning system that mines a million news articles each day to identify signals that will have implications on the financial performance of a borrower.

The power of Credit Assistant is seen in the results of one bank's implementation to support loans ranging from a few thousand to tens of millions:

- Implementation in 77 days with connections to 50 internal systems, 200 screens, and 250 end users across 24 markets in 10 languages.

- Addition of new products without programming and new underwriting rules or scorecards implemented in hours.

Genpact built Cora LiveSpread, applying AI technologies to extract and interpret financial documents and using a built-in knowledge base of 60,000 rules across 22 sectors, 19 languages, and GAAP rules for 12 countries. It has reduced costs up to 70%. In addition, it institutionalizes global and local compliance.

IBM WATSON

To improve credit portfolio monitoring and anticipate potential issues, IBM released Algo Credit Manager v5.4 in 2017, which includes the introduction of sentiment analytics based on unstructured data using Watson's natural language capabilities. Watson captures sentiment (negative, neutral, positive) directly from unstructured data sources (e.g., news, social media, customer feedback sites) and translates this data into a sentiment number and tracks the trend of the sentiment over time (Figure 25). If there is a change in the trend, a bank can proactively manage a customer and for example, if the sentiment turns negative, start a process of customer re-evaluation.
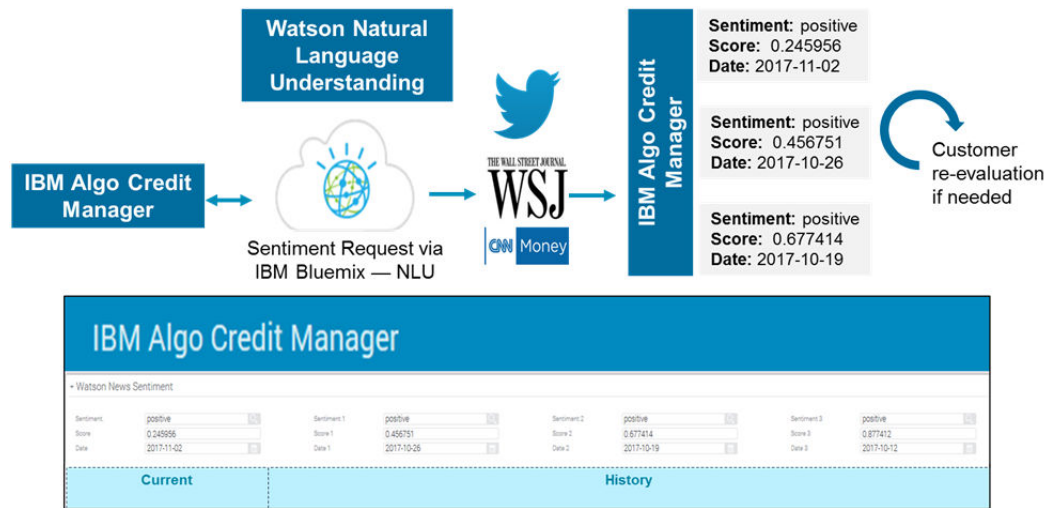
Source: IBM.

MPHASIS DEEPINSIGHTS

Just as Mphasis DeepInsights (outlined on page 19) is being used to digitize and analyze invoice data, it can be used to extract data relevant to underwriting credit and to digitizing trade finance document processing. It extracts structure and unstructured data from digital and nondigital sources (e.g., images, news articles, annual reports, emails, and customer complaints). It then analyzes this data to determine next steps in a specific workflow.

For credit underwriting, DeepInsights has been used in a pilot to automate financial data extraction and apply advanced AI to analyze the data. The results are compelling: 95% accuracy of data extraction, 70% reduction in time to extract relevant data, and 50% decrease in operating costs (the clients are a leading global bank, a top tier European bank, and a top 5 European insurer).

In trade finance, DeepInsights can be used to digitize a variety of trade finance document processing and AML and sanctions screening. To date, it has been used in a pilot to automate seller/exporter document verification for a top tier US bank and French retail bank. In addition, it provides real-time alerts and a reporting dashboard (part of its visualization engine) to enable swift exceptions handling. The results from its pilot are compelling: 75% reduction in time to digitize a trade document (from 40 minutes to 10), and around a 30% decrease in operating costs.

On the 2018–2019 horizon, Mphasis plans on extending DeepInsights to fully automate loan origination from document collection and processing to approval. For trade finance, DeepInsights will be applied in AML and sanctions screening. It will provide an integrated solution that screens documents for contraband items for custom clearance and OFAC sanctioned entities for compliance and flags potential money laundering activities.

NUMERATED

Numerated Growth Technologies' platform is powering a new credit model at several banks. It enables not only a bona fide end-to-end fast process for small loans (a single session as short as five minutes for current customers) but also direct marketing and prospecting at a relatively low cost. Numerated was incubated in Eastern Labs within Eastern Bank and was spun-off in 2017 (Venrock, Cultivation Capital, and FIS are investors). Hence, it was built from a banker's lens with a strong understanding of the realities and challenges bankers face as well as the need to accommodate a bank's credit policy. Numerated and FIS have a strategic partnership to bring Numerated's real-time lending and marketing automation platform to more US banks. Numerated's integration with FIS core systems provides an out-of-the-box solution for banks to more rapidly implement lending automation and offer real-time lending to their SME customers.

Driven by a revenue growth objective, Numerated built an end-to-end, omnichannel platform (over half of all traffic comes through mobile devices). Its offering goes beyond a digital workflow to include sophisticated upstream prospecting and marketing. Therein lies its secret sauce: a single session process. It enables a bank to email prequalified small business customers an invitation to apply for a loan, and interested customers can click within the email to be taken to a personalized five-question application, receive an instant decision, review and sign loan documents (these too are streamlined to under five pages), and receive funding in their account, provided the borrower is a bank customer.

Underlying its instant decisioning engine are various AI and data science techniques and tools, including machine learning systems to categorize and verify business identities by synthesizing multiple public data sources, NLP to determine key business characteristics from Internet data, and predictive payment likelihood and loan performance models using bank transactional and tax return data in combination with credit bureau information. Numerated has also created new variables that banks can use in decisioning, including corporate structure and transaction analysis of a business's deposit account. In addition, it offers flexible, dynamic workflows and the ability to make more than one offer. For example, if a bank wants to route certain segments of applicants to underwriters for a deeper review, it can. If a bank wants to route an applicant to an SBA loan offer, it can. If it wants to route a rejected applicant to an underwriter for a second review, it can. If it wants to incorporate share of wallet into the analysis and offer a discounted interest rate, it can.

For prospecting, Numerated has trained natural language processing and machine learning text-recognition systems to cleanse and classify every banking transaction and then identify a potential applicant's key financial providers and sources of expense and revenue. For example, payments to loan providers are recognized and classified. This information feeds and enhances prospect lists used by RMs, who can offer refinancing opportunities to potential applicants. The engine also allows bankers to see a business's "network," that is, with whom it is transacting and whether its counterparties are bank customers. When counterparties are officers of businesses known publicly through regulatory filings, these relationships can also be exposed to identify unrealized opportunities beyond the current customer base. The platform can also identify payments

and transfers to third party financial institutions and providers and make this information searchable and sortable for easy sales opportunity targeting.

Proof of the effectiveness of Numerated's platform is seen at Eastern Bank, which was able to quadruple its small business loan portfolio in two years while maintaining its risk exposure limits. In addition, it was able to improve its net interest margin through better risk pricing and by pricing in the value of convenience. At the beginning of the process, direct marketing, it can increase open rates six fold.

On the horizon, Numerated is planning on applying its expertise to bring rapid decision to the underwriting of larger, more complex loans (e.g., larger small business loans and commercial real estate loans) by undertaking upstream analysis (e.g., debt service coverage ratio) and delivering results to an underwriter. It is also examining how its platform could be applied to consumer credit (e.g., unsecured consumer loans and home equity lines of credit).

PREVISE

Previse, a London-based startup founded by a former Global Co-Head of Principal Strategic Investments at Goldman Sachs and a team of industry experts, is reimagining how supply chain financing (SCF) works. Previse is using machine learning to radically change SCF along two critical dimensions: payment timing and supplier reach. Its machine learning model predicts the likelihood that a large buyer will pay a supplier's invoice, thereby enabling payment upon invoice receipt as opposed to invoice approval, thereby reducing the time to payment often by months. It can also apply the same concept from the time of issuance of a purchase order. Its model makes predictions for all invoices and hence enables small to medium enterprises (SMEs) suppliers to receive financing — the long tail of suppliers which typically represent 80% plus of the total and are excluded from bank-run SCF programs.

Previse is generating multiple win-wins. First, it is enabling SMEs to be paid immediately at a relatively low cost, that is, typically half the cost of factoring, which is an SME's primary source of working capital financing. This benefits both suppliers and large buyers, in that if suppliers' cost of borrowing is lower, they could pass part of the savings to buyers in the form of lower prices. Large buyers benefit in four areas. First, they can improve their reputation and supplier relationships as well as respond to criticism that they pay small suppliers late (e.g., in Europe, they are under government pressure to reduce payment delays). Second, they receive a revenue share for partnering with Previse (a data fee). Third, buyers have additional data analytics they can leverage to, for example, further optimize their source-to-pay processes. Fourth, they can achieve operational cost savings from improved back office efficiency, including an algorithmically prioritized work queue and robo-approvals. Moreover, unlike with traditional SCF programs, buyers and suppliers do not have to do anything differently to accelerate the payment cycle, and suppliers do not need to change their AR process or sign any new contracts. As a result, Previse has removed two points of friction: e-invoicing and contract negotiation.

By opening up the roughly 80% of the $16 trillion in trade not financed, it is generating a new asset class ("pre-approved invoices") that banks, hedge funds, private equity, and

institutional investors are eager to invest in given the low rate environment. Previse allows funders to adjust risk exposure based on payment probability thresholds (e.g., 99.5%).

Previse built its machine learning model with an exceptionally rich data set: $1 trillion in spend data. After successfully testing its model through six proofs of concept (i.e., simulation basis) with six multinational buyers and hundreds of thousands of suppliers, it raised seed funding in July 2017. To train its model for each buyer, it will take in three years of a buyer's spending data and analyze 20 to 30 features of each invoice (e.g., spend category, supplier location, currency, and amount). To add a new supplier to its program, it needs a year's worth of invoice payment data. It has signed up its first multinational clients and funders and expects to go live in the first half of 2018.

WORKFUSION

WorkFusion stands out on several fronts. First, it has a record-breaking five use cases in production (out of the eight covered in this report). Second, it is unique in combining workforce orchestration, OCR, and RPA with machine learning to automate end-to-end processes. It has battle-tested its platform in production at some of the largest global banks across continents (e.g., Citibank, Standard Bank, Axis Bank). Its workforce orchestration enables a company to effectively collect robust data to train machine learning models while managing a global workforce and controlling the work quality. Its machine learning models are trained by humans, who extract and interpret unstructured data through graphical user interfaces (GUIs). Once trained, these models automate judgment work. If the model cannot determine the appropriate action, the task is routed as an exception to the appropriate analyst, through the exact same GUI used for training. This enables the machine learning model to keep learning from what humans do and increase both automation and accuracy, and thereby productivity.

WorkFusion further differentiates itself with its Virtual Data Scientist capability which enables machine learning models to be automatically created by the software. In other words, a bank using WorkFusion does not need a data scientist to create an machine learning model and hence can scale advanced AI-driven automation relatively faster.

In corporate banking, WorkFusion has been used in a myriad of use cases, ranging from account opening for a new borrower, loan booking, compliance error identification post-booking, corporate loan notices to trade finance processing, sanctions screening, payments processing, and fraud detection (the latter two are discussed in the next section).

For loan booking for a US bank, WorkFusion's software developed seven machine learning models (six classification models and one data extraction model). Evidence of its effectiveness is seen in the results achieved in production: 79% average automation rate with 92% accuracy for new loans, and 99% and 98% respectively for renewals (all statistics after the initial training). The models will continually learn from the exceptions processed by humans. Hence both automation and accuracy rates will go up.

For notices, evidence of the effectiveness of WorkFusion's approach is seen in the results achieved a top US investment bank: a 80% reduction in processing time per

notice (from 5 minutes to 59 seconds), a 90%+ accuracy across all data fields, work load balancing, and an estimated labor cost savings of $750,000.

In trade finance, WorkFusion has automated numerous end-to-end back office processes including sanctions checks and issuance of a letter of credit, eliminating the human "maker" and turbocharging the human "checker." For example, it automates import payment processing with a combination of advanced AI and robotics.

WorkFusion has potential to move to customer-facing activities in commercial banking. In 2016, WorkFusion added chatbots and began applying NLU and NLG to automate transaction and status requests which consume back office bandwidth without adding value to the client relationship. It has applied its chatbots in retail banking to automate requests related to real estate loan origination (e.g., "How much can I borrow?" "What rate can you offer me?").

COMPLIANCE AND FRAUD

One of the highest AI summits has been compliance, which is not surprising given that compliance failures have cost banks billions in fines despite their employing, in some cases, tens of thousands of dedicated staff and costly software solutions. Within corporate banking, the focus herein is on KYC, AML, and sanctions screening, which are ripe for reimagining. Growing transaction volumes, inefficient manual processes and paper trails, and tightening regulatory requirements are driving unsustainable increases in labor costs. For example, banks report AML compliance costs increasing exponentially and false positives (an alert or investigation that does not result in a suspicious activity report, or SAR) as high as 98%, which implies that a majority of investigators are not needed. At the same time the burden of KYC has grown as tax-related regulation requirements rise (e.g., Foreign Account Tax Compliance Act in the US and its global equivalent, Common Reporting Standard).

While RPA can effectively be used to undertake basic, repetitive processes (e.g., KYC remediation, pulling financials and filing with regulators; cross-checking with OFAC list), advanced AI is needed to process unstructured data, which is an important source of information supporting compliance. The results of AI-powered processes have been positive: dramatically lowering false positives and occasionally catching false negatives (examples are included below).

In addition to compliance, payments fraud mitigation and cybersecurity are critical areas being addressed by AI. AI-based solutions are well suited to tackle the challenges of business-to-business (B2B) payments and real-time payments fraud. For both, false positives have a relatively much greater adverse impact than in cards and hence must be minimized. B2B fraud mitigation requires a relatively sophisticated approach because it is much harder to detect than consumer-to-business payments fraud. It tends to be episodic, and incidents tend to have unique traits. Faster payments demand real-time authorization for all payment types. Moreover, to encourage adoption for B2B, transaction limits have to be set higher than $25,000, which requires a highly effective fraud mitigation system. In the area of cybersecurity, advanced AI can be used to tackle one of the most challenging frauds to detect: imposter fraud (e.g., business email compromise), which has become rampant.

Figure 26, 27, and 28 outline the leaders in compliance and fraud detection-related AI applications (Figure 26 lists providers in order of maturity level, and Figures 27 and 28 list them alphabetically within maturity level). These two broad use cases have the biggest competitive field and range of initiatives with a variety of providers from multi-industry to bank-focused making the ascent.

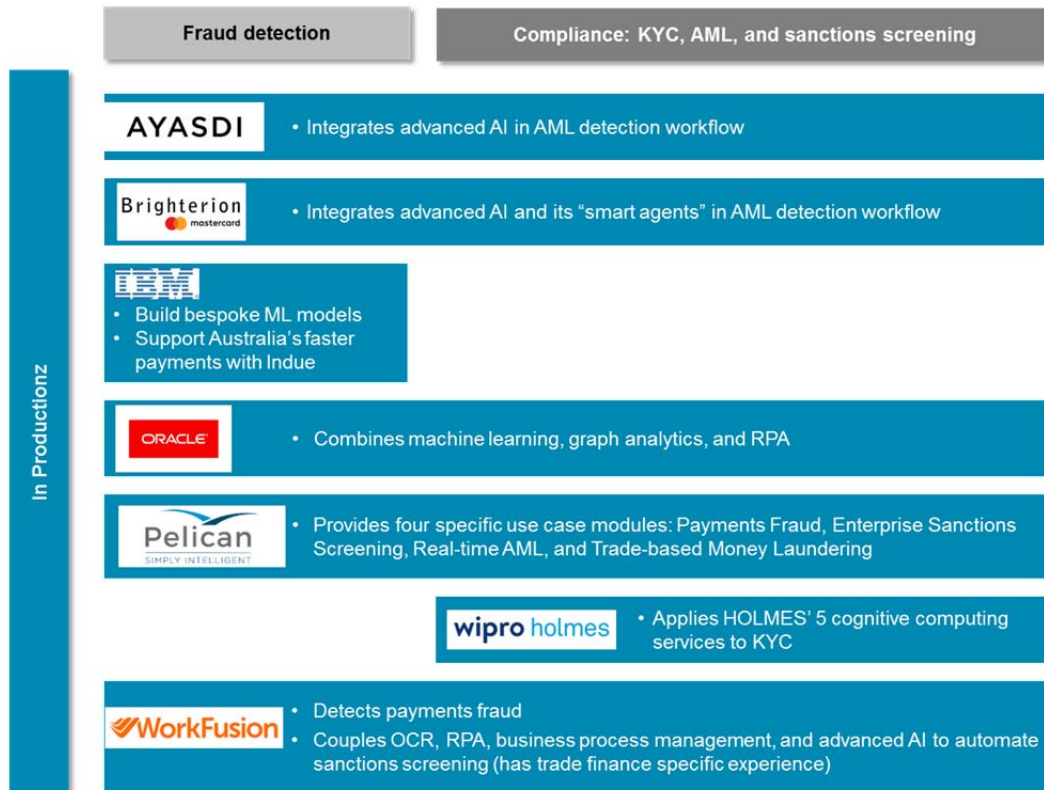Figure 26: Mountaineers in Fraud and Compliance — Overall

Company	Fraud Detection	Compliance (KYC, AML, Sanctions Screening)
AYASDI	✓✓✓	✓✓✓
Brighterion (mastercard)	✓✓✓	✓✓✓
ORACLE	✓✓✓	✓✓✓
Pelican SIMPLY INTELLIGENT	✓✓✓	✓✓✓
WorkFusion	✓✓✓	✓✓✓
IBM	✓✓✓	✓[1]
wipro holmes		✓✓✓
(intel) Saffron		✓✓
Mphasis Unleash the Next		✓

Legend: ✓✓✓ = in production
✓✓ = in implementation or released but not installed
✓ = pilot, proof of concept or scheduled for release in 2018

Source: Celent interviews and analysis
1. Includes Financial Crimes Insight for Watson—Due Diligence and Financial Crimes Alerts Insights with Watson. Outside of the areas covered herein, IBM Watson is in production in Regulatory Compliance Management and Analytics.

Figure 27: Mountaineers in Fraud and Compliance — In Production

Fraud detection	Compliance: KYC, AML, and sanctions screening

In Productionz

AYASDI
• Integrates advanced AI in AML detection workflow

Brighterion mastercard
• Integrates advanced AI and its "smart agents" in AML detection workflow

IBM
• Build bespoke ML models
• Support Australia's faster payments with Indue

ORACLE
• Combines machine learning, graph analytics, and RPA

Pelican SIMPLY INTELLIGENT
• Provides four specific use case modules: Payments Fraud, Enterprise Sanctions Screening, Real-time AML, and Trade-based Money Laundering

wipro holmes
• Applies HOLMES' 5 cognitive computing services to KYC

WorkFusion
• Detects payments fraud
• Couples OCR, RPA, business process management, and advanced AI to automate sanctions screening (has trade finance specific experience)

Source: Celent interviews and analysis

Chapter: Compliance and Fraud

AYASDI

AI is Ayasdi's domain expertise. For financial institutions, it is focused on delivering AI solutions for regulatory and risk management use cases and providing tools for data science teams and their partners to build their own applications. It has developed six application areas: financial crime, conduct risk (employee conduct), model development and risk management, regulatory risk, cyber risk, and customer intelligence (e.g., attrition risk prediction). To date its experience in corporate banking has been in financial crime, specifically AML and payments fraud detection. Citibank is both an investor and a customer, using Ayasdi's Model Accelerator to build its pre-provision net revenue models.

Its platform distinguishes itself on several fronts (Figure 29). First, it intakes data from a range of sources and puts it in a structured format (i.e., rows and columns) using Hadoop. Second, Ayasdi is the only provider to have an enterprise platform powered by topological data analysis (TDA), which is described by Ayasdi "as the most principled way [i.e., one supported by a body of mathematical knowledge and technique] ever developed to understand the shape of a dataset — allowing for a superior and unbiased method of discovering hidden patterns and relationships within a dataset." TDA combines machine learning, statistical, and geometric algorithms. Depending upon the nature of the business problem, Ayasdi's machine learning models can range from unsupervised to partially or completely supervised, with all results explainable.

Figure 29: Outline of Ayasdi's Machine Intelligence Platform
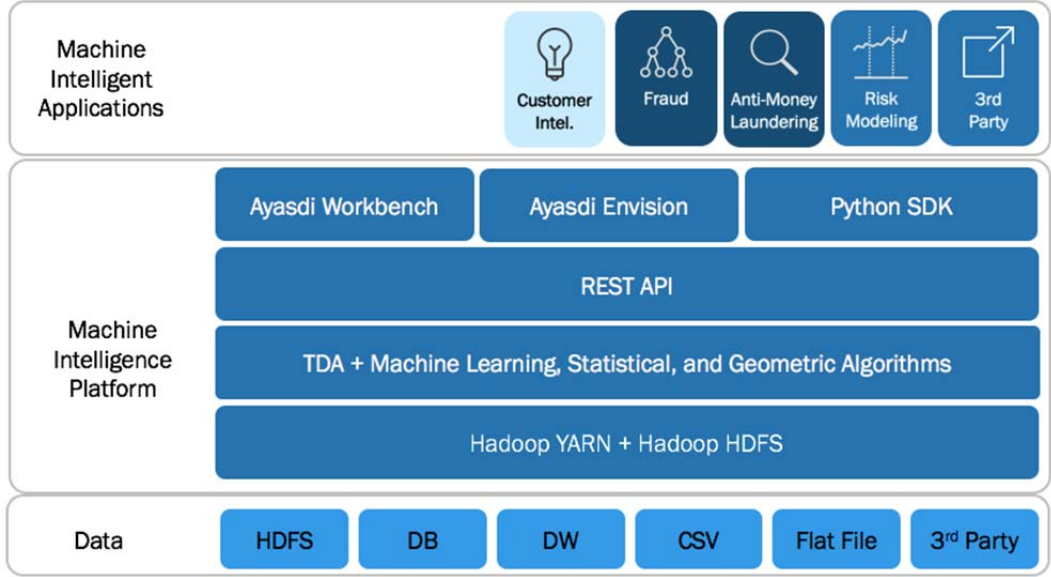
Source: Ayasdi.

Figure 30 shows how Ayasdi augments a bank's AML workflow in three steps. First, Intelligent Segmentation hones customer, transaction, and alert segmentation, improves scenarios, and optimizes thresholds with the goal to increase thresholds without increasing risk exposure (Ayasdi can take data from the customer onboarding process and use the data for threshold tuning). With this step alone, HSBC reduced false positives 20%. Second, Intelligent Alerts triages alerts based on priority level while providing the rationale in order to increase the efficiency of AML alert investigators. Third, Intelligent Typologies creates customer profiles and identifies anomalous behavior, which in turn is used to add typologies to spot new criminal behaviors and better identify true positives.
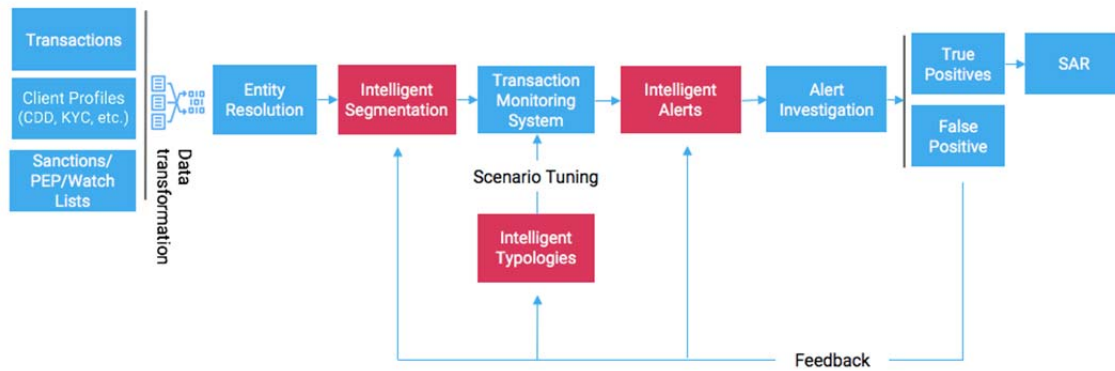
For 2018–2019, Ayasdi plans on streamlining the ability for data scientists and developers to leverage the enterprise platform to quickly perform key data science tasks, such as segmentation, anomaly detection, and hotspot analysis. Ayasdi will also create a prebuilt AML-specific AI Workbench. Its current AI Workbench is a graphical modeling environment, essentially a compressed visual summary of all of a company's data sets, which enables analysts to relatively rapidly detect risk indicators (e.g., data anomalies).

BRIGHTERION

Brighterion — acquired by MasterCard in 2017 after having a partnership for a decade — stands out on several fronts. First, it has built a highly flexible general purpose AI platform with several unique features: "smart agent" technology (adaptive unsupervised learning to update models), segment of one behavior analysis (i.e., a smart agent is built for every entity), and AI-based distributed data storage. Second, it can be deployed relatively fast. Third, its platform has been battle-tested and proven that it can handle 14 billion transactions per month at milliseconds per transaction (e.g., fastest response time is 2 milliseconds), which bodes well for its use in real time payments fraud detection. It counts six of the top 10 global banks, a leading payment network (MasterCard), and a leading merchant acquirer (WorldPay) as customers.

Brighterion's product suite includes: iComply (KYC and AML), iDetect (data breaches), iMonitor (model performance and data trends, e.g., fraud trends), iPredict (credit risk and delinquency), iPrevent (cross-channel fraud prevention), and iSupervise (case management). On the commercial side of banking, Brighterion's iComply has been used to automate KYC, AML, and sanctions screening processes. For example, WorldPay uses it to improve its merchant KYC process.

Brighterion applies 10 core advanced AI technologies as outlined in Figure 31 to power its smart agents. Its smart agents approach underlies its strong performance and long model lifespan. Brighterion's platform builds a smart agent (aka profile) for every bank customer/account (e.g., if a bank has 1 million customers, there are a million smart agents). As a result, it is able to undertake behavioral analysis across channels, which is particularly useful in detecting payments fraud and money laundering. On average, it has delivered two to three times the detection rates while reducing false positives 90 to 95%. Importantly, its models avoid being a black box by providing a rationale and reason codes

for all auto-decisions. In addition, Brighterion's smart agents technology continually monitors the unique behavior of each entity to detect new trends and behaviors, enabling the models to learn and improve over time. Brighterion's technology also results in relatively lower operational and modeling costs.

Figure 31: AI Technologies that Power Brighterion's Smart Agents
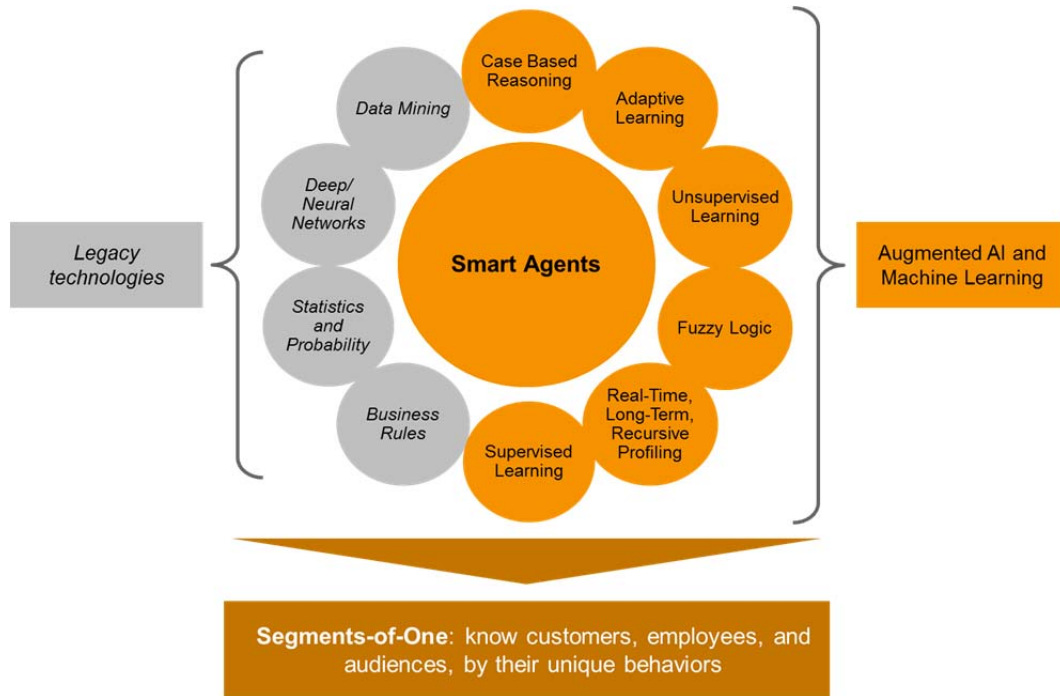
Source: Brighterion

Brighterion's patented AI-based distributed data storage enables customers to deploy relatively quickly at low cost using any data, format, volume, and source. It eliminates the need for any database or data warehousing requirements. By being able to compress data (e.g., it can reduce storage requirements 99%) and achieve very fast processing speeds, it is highly scalable. Customers often deploy within six weeks using off-the-shelf hardware. Brighterion typically requires at least 13 months of data and can build models based on small datasets (e.g., 500 rows). Brighterion builds bespoke models for each of its customers. Evidence of Brighterion's fast deployment capability is seen in WorldPay's AML deployment, which was achieved in less than six weeks.

Evidence of its technology's effectiveness is the award it received from Morgan Stanley in November 2017: the Morgan Stanley Fintech Award, which recognizes the fintech company that had the most significant impact on its mission to continuously innovate for its clients and business. Morgan Stanley's Managing Director and Global Head of Fraud spotlighted: "The open architecture of the platform offers a remarkable level of transparency, providing our teams with a faster and easier way to build, enhance, and adapt models."

During 2018, Brighterion plans on releasing version 9 of its AI platform which introduces additional supervised and unsupervised learning technologies.

IBM WATSON

IBM has coupled Watson with training by its bank compliance domain expert, Promontory Financial Group (acquired in 2016) to deliver IBM Financial Crimes Insight (FCI) with Watson with an initial focus on KYC (FCI - Due Diligence) and AML (Financial Crimes Alerts Insight with Watson). Watson comprises numerous types of machine learning models and comes with cognitive accelerators, that is, a set of tools and analytics to expedite model building, select the optimal model, and train it based on a bank's data sets, which could be relatively small.

The goal of FCI with Watson is to improve KYC and compliance analysts' productivity while reducing false positives. Watson can evaluate structured and unstructured data from a variety of data sources, including not only traditional sources (e.g., case notes, LexisNexis, customer account data) but also nontraditional, previously unconnected data (e.g., Google, Facebook, Yelp). It then generates insights (e.g., a custom narrative) and identifies the most relevant content to deliver to the analyst for case investigation. Insights include the results from its network analytics, that is, "who knows who" analysis. To facilitate and expedite decisioning, the analyst can examine graph analytics which visualize entities (person or business) associated with an investigation and relationships to other entities, such as fraud rings or collusive networks. Moreover, the analyst can determine if an entity is using variations of an identity to disguise itself and/or is on any watch list. Importantly, in training Watson, analysts review alerts and the reasons that Watson believes an alert is a false positive (e.g., alerts with same reason code were dismissed nine times before based on evidence found by Watson that recipient is a legitimate business) and then make the final decision.

Figure 32 outlines IBM's initial and anticipated applications. For KYC, it is focusing on supporting enhanced due diligence analysts. For AML, it is prioritizing and triaging alerts for case investigators and then helping them make better and faster decisions. The results have been compelling. In one proof of concept, case investigators using FCI with Watson were able to process 25 case dossiers nearly 75% faster than those not using FCI (5 minutes compared to 19 minutes).
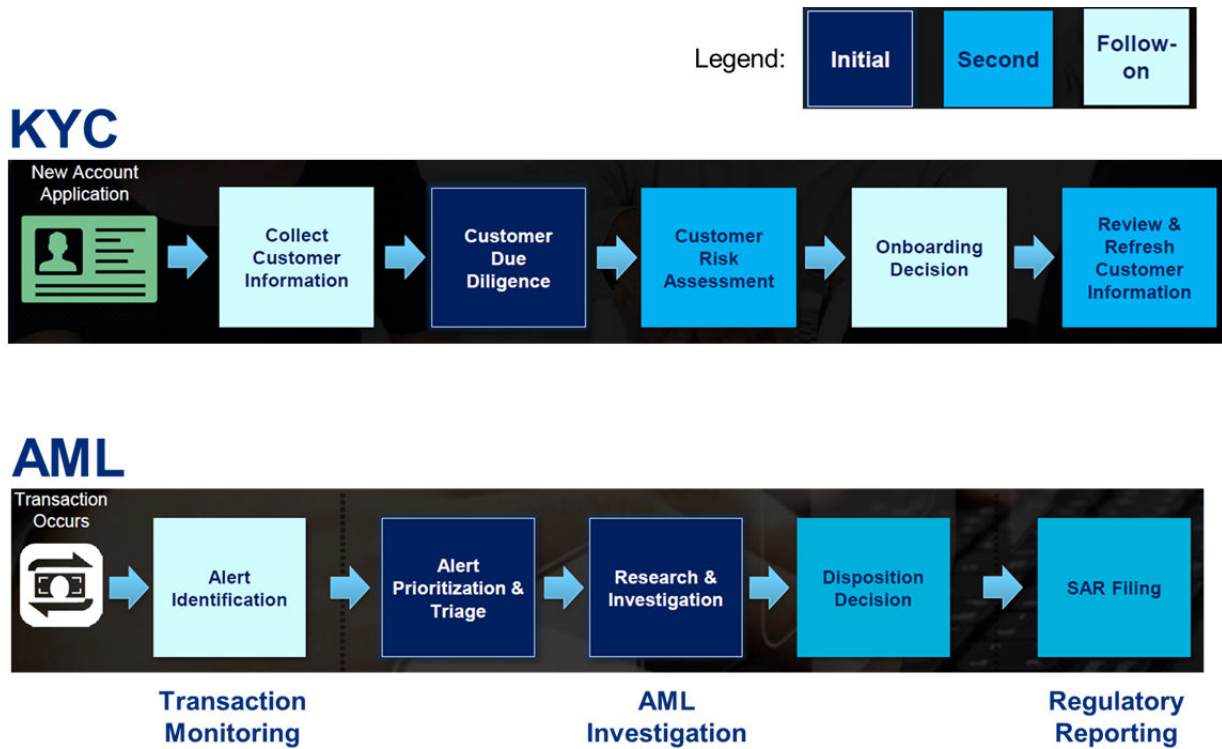
Source: IBM

IBM FCI with Watson is currently in the proof of concept stage at five global banks. Several of these POCs were completed in 2017, and these banks will begin production rollouts in Q2 2018. As part of its development roadmap, IBM will apply Watson to other parts of the workflow as outlined in Figure 32. The challenge in moving up and downstream is that an open architecture and APIs are required.

In the area of payments fraud mitigation, IBM Safer Payments uses Watson to perform granular anomaly detection on each payment as outlined in Figure 33. Hence Watson's approach differs significantly from traditional neural networks used to detect credit card fraud, which are based on the assumption that fraud exposure is the same for all banks and which need huge data sets (e.g., banks often share data) and deliver a model based on spreadsheets and fixed calculations. Watson ingests real-time payments data and compares it to historical data (6 to 12 months), seeking patterns, constantly simulating and updating the model. Watson combines machine learning hindsight with human foresight by generating rules for a human to evaluate (i.e., explainable AI). The more rules that are tested with humans, the smarter the model becomes, and the greater its ability to adapt and address new and evolving episodic fraud threats.

Source: IBM

It has recently gone into production in Australia with its partner Indue to support the New Payments Platform, Australia's faster payments network (Indue Orion Financial Crimes will offer a 24×7 fraud and sanctions monitoring service for the NPP).

INTEL SAFFRON

Intel acquired Saffron, a leader in associative memory AI, in 2015. Saffron's patented AI technology identifies connections between data points (structured and unstructured, e.g., websites and email) at the entity level (person, place, or thing) and stores the connections in an associative memory, which mimics the natural ability of the human brain to identify similarities and anomalies in a wide variety of dynamic, heterogeneous data and situations at a far greater scale than a human could. Critical for compliance applications, Saffron's AI is a white box, that is, it can explain how it arrived at a conclusion (e.g., how connections were identified, and how it makes decisions to support human analysis in complex compliance environments such as financial crimes identification and investigation).

Saffron's brand of AI has several unique aspects. First, Saffron is not a "deep learning" solution. It does not employ models requiring training against extensive data. Instead, it uses algorithmic approaches that identify unique connections and associations, often on very sparse data; following a "one-shot learning" approach rather than deep learning's many iterations against large data sets. Second, Saffron utilizes a proprietary "loss less compression" approach that creates extremely compact storage of relevant data arrays, minimizing storage and CPU hardware requirements. In addition, it has proven to be highly scalable and efficient in cases of extremely large datasets. For example, working with one of the largest global financial institutions, in fewer than eight weeks Saffron was able to establish a recall rate of over 90% of SAR filings from a truth set consisting of six months of transactional data and customer information from more than 90 million customers. From the same data set, Saffron also identified "false negatives" or activity

that was potentially suspicious but had not previously been flagged by traditional systems. Upon subsequent investigation, approximately 25% of these false negatives were deemed suspicious enough to warrant SAR filing.

Saffron's ability to illuminate multidimensional patterns has been applied to a variety of use cases across industries as well as government agencies. For example, the Department of Defense was an early Saffron client. The Intel acquisition has enabled Saffron to productize its AI capabilities with Intel Saffron's AML product being recently announced (October 2017). The solution leverages Saffron's ability to score risk and detect anomalies to identify true suspicious activity positives and increase the productivity of AML investigators. The application also offers the ability for investigators to visualize emerging activity pattern clusters or trends. Figure 34 shows how Intel Saffron supports AML investigators.

Figure 34: Intel Saffron expedites AML Investigations
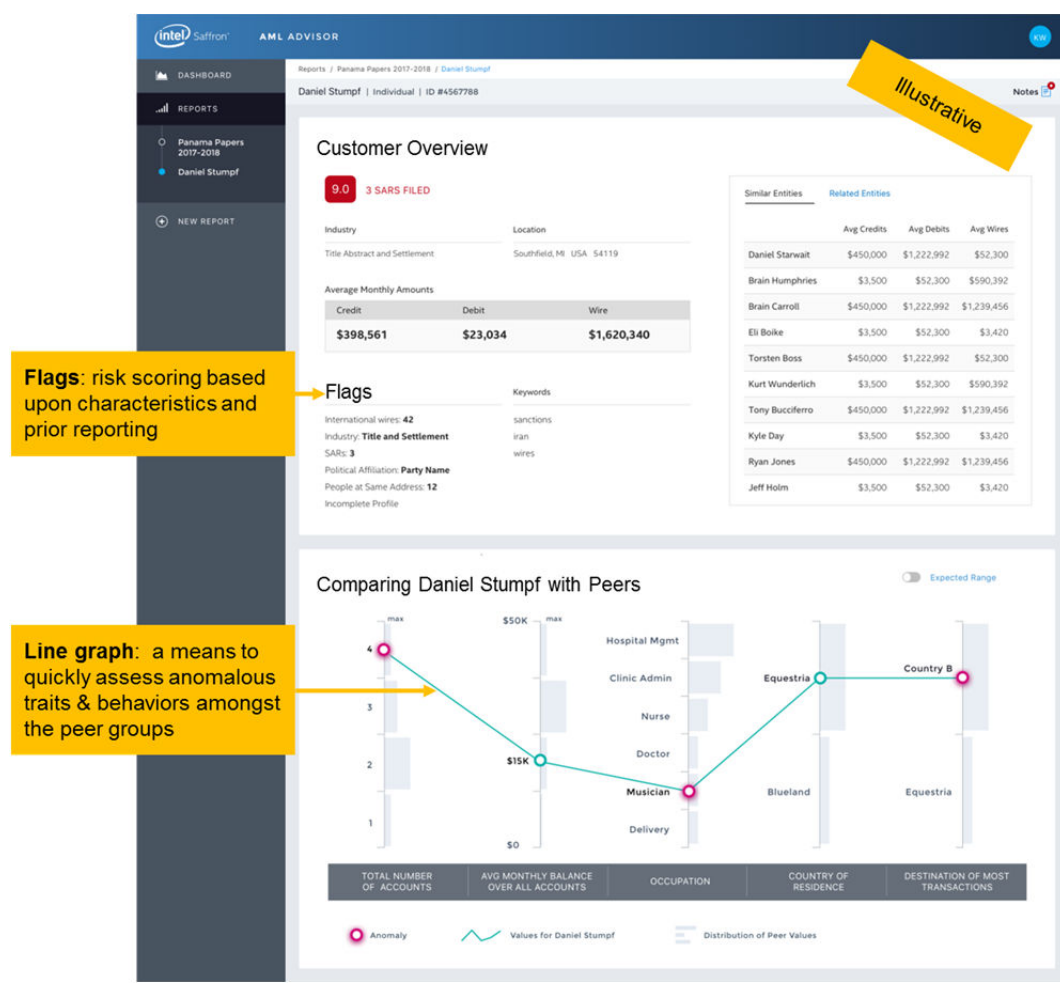
Source: Intel

Intel Saffron has also launched an Early Adopter Program to partner with a maximum of five institutions to expand its use in financial services. Bank of New Zealand is a Saffron early adopter in the customer optimization arena. Saffron's product roadmap includes eventual expansion from financial crimes compliance to the adjacent areas of fraud and

customer personalization. In 2018, the focus will be primarily on financial crimes compliance in English-speaking markets with broader expansion in both global markets and product suite development in 2019 and beyond.

ORACLE

Oracle has been incorporating advanced AI capabilities within its Financial Services Financial Crime and Compliance Management (FCCM), its detection and investigation system (over a hundred top banks worldwide use its detection system). It is one of the few vendors whose detection engines can run rules, machine learning models, and graph analytics or a combination of these techniques as part of detection logic. It began tackling the challenge of lowering false positives in AML and sanctions screening with AI in early 2016 with the objective to not only lower false positives but also to encourage banks to share information regarding detection patterns.

FCCM's goal is to serve the four key stakeholders in compliance operations efficiency: auditors, operations team (e.g., those that write scenarios and generate alerts), case investigators (i.e., the end users who should not need to understand the technology), and the data scientists who improve the models. As outlined in Figure 35, its components correspond to specific activities and analytics which support these stakeholders.

Figure 35: FCCM Core Vision
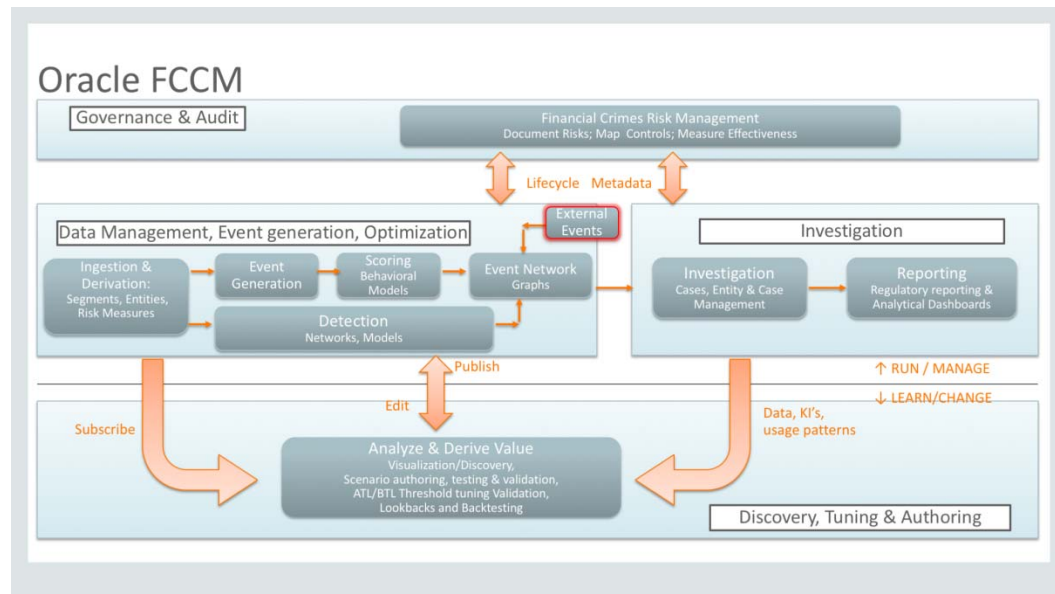
Source: Oracle

In the data management, event generation, optimization component, it is applying machine learning and graph analytics:

- Customer segmentation — machine learning used to discover customer clusters. The machine learning results are combined with expert knowledge.

- Event scoring — machine learning used to predict event importance based on the results of historical case investigations.

- Event graph creation and scoring — graph analytics and machine learning used to score graphs of events with high scoring graphs being passed on to case investigators.

- Pattern detection — machine learning applied to transaction data to detect networks of suspicious activity and AML patterns.

In the investigation component: machine learning will be applied to predict next best action (based on historical case adjudications), and NLP as well as RPA and text analytics will be used to generate suspicious activity reports.

As part of its May 2018 release, Oracle will add scoring event (alert) graphs using machine learning to prioritize an entire set of related alerts and use graph analytics in case creation.

PELICAN

Pelican, founded by an AI engineer in 1993, has been at the forefront of AI applications in sanctions compliance, fraud, and payments processing (the latter is covered in the next section). Its solutions have been battle-tested in high volume environments (e.g., TD Bank screening 65,000 transactions per day, and ICICI screening 60,000 transactions per day). Pelican's compliance and fraud solution, PelicanSecure, comprises four modules: Enterprise Sanctions Screening, Payments Fraud, Real-time AML, and Trade-based Money Laundering (Figure 36). Revenues are roughly split between PelicanSecure and PelicanPayments (see next section). PelicanSecure has produced significant improvements in sanctions screening: it has reduced the false positive rate by as much as 72% and has cut 80% of the time to undertake a sanctions review.

Pelican has been pushing the envelope in NLP and machine learning to develop its unique "semantic understanding" across all the text it analyzes. It begins with "lexical analysis," which splits text into "tokens" of words, numbers, and other symbols and then does "conceptual parsing" to associate meaning to the tokens. The result is what Pelican describes as a "precooked database" comprising action codes used by the system to take appropriate action (e.g., route a transaction to a fraud or sanctions analyst; repair a payment message for straight-through processing). As a result, its platform is a white box, providing full explanations and audit trails for each action. This explainable AI combined with investigation dashboards drive material productivity gains for case investigators.
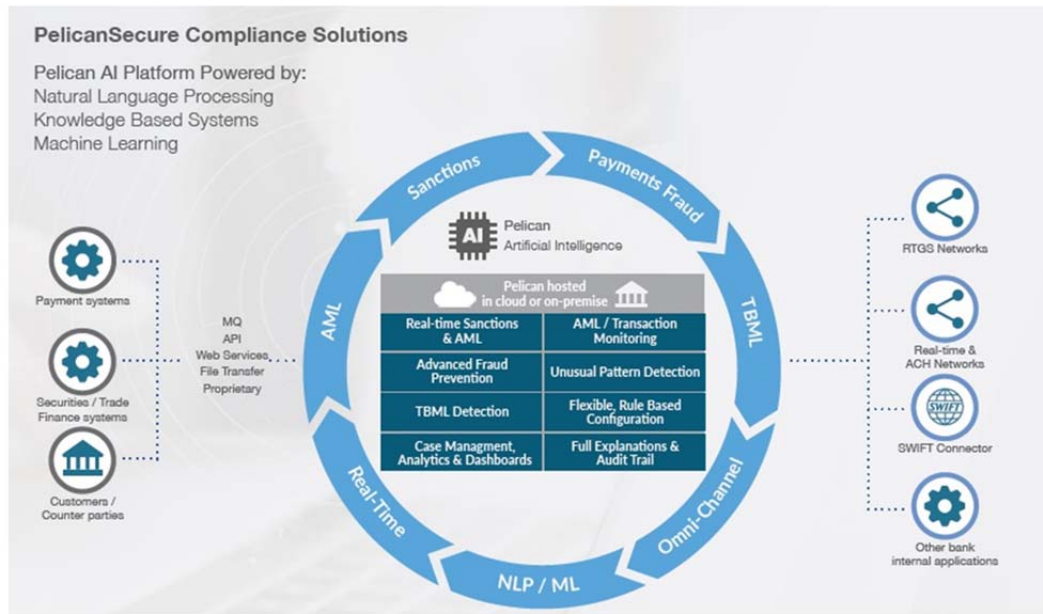
PelicanSecure combines the power of rules-based engines, NLP, and machine learning to undertake real-time detection. For Enterprise Sanctions Screening, Pelican has built over 40 matching algorithms, which generate detailed alert information, thereby reducing review times up to 80%. For Payments Fraud, it is using supervised machine learning to enhance the rules engine through self-calibration of results from analysts' fraud review and reinforced learning to continuously incorporate the actions of analysts processing alerts. In addition, it uses unsupervised machine learning with cluster and anomaly detection to identify new fraud techniques in real time.

With Trade-Based Money Laundering (TBML), Pelican is the only vendor developing an out of the box application for trade finance, addressing its particular issues (e.g., dual-use goods, ghost shipping, and commodity pricing anomalies). It is applying the NLP it has developed to contextualize the unstructured data in trade finance documents, undertake intelligent document matching and consistency checks, and detect "red flag indicators" as specified by various regulators (e.g., Singapore's MAS and US's OFAC).

In 2018, Pelican is investing heavily in advanced AI with the goal to not only improve its middle and back office applications but extend to the front office as well. It will integrate new data feeds into TBML, such as, vessel tracking and specific commodities price lists. It is enhancing PelicanSecure for corporate use cases. To push its AI to new heights, Pelican is launching a global innovation hub to explore new technology and use cases, working in partnership with clients on R&D and proofs of concept.

In addition, it will expand its payments-related AI beyond correspondent banking to real-time payments whose fraud detection poses significant challenges for banks running batch processes. It has undertaken testing with the Euro Banking Association for the SEPA Instant Credit Transfer scheme in Europe, and is working with The Clearing House in the US to facilitate its vendor integration process.

WIPRO HOLMES

Over the past four years, Wipro has developed an AI platform, HOLMES (heuristics and ontology-based learning machines and experiential systems). Wipro's approach to developing AI capabilities reflects its core capabilities as an IT services provider: target use cases and then build a holistic solution. First, it identified manually intensive use cases involving varying degrees of commodity skills and quickly arrived at KYC in the banking world. Second, it determined how a use case could be reimagined not only through advanced AI but also by using mainstream automation technology. Through partnerships, Wipro also offers other cognitive technologies (Watson, Amazon, or Microsoft) and RPA.

HOLMES comprises five cognitive computing services (Figure 37):

- Ingest data via natural interaction and interface (speech recognition, natural language interaction and generation).

- Digest the data via a knowledge engine which includes knowledge extraction from unstructured and structured data and semantic mapping.

- Analyze the data via an algorithmic intelligence engine using a combination of calculations including stochastic computations, statistics, and machine learning.

- Learn continuously via a machine learning engine.

- Process like a human through a reasoning engine which is reasoning rule and ontology-based (i.e., conceptualization of the data, which involves naming and defining types, properties, and interrelationships of entities).

Source: Wipro

As Wipro hunted for processes that could be reimagined via an AI platform, KYC not surprisingly made the short list, and by 2016, it launched E-KYC (enterprise-KYC). With the goal to digitize the entire process and automate the work "KYC makers" do, it examined the end-to-end process from data gathering to verification and audit trail generation and the work of thousands of FTEs (Figure 38). It applied AI to tackle the parts most challenging to automate, such as the gathering of external data (e.g., financials from annual reports, including footnotes, controller information) and determining the ultimate beneficiary owner.

Figure 38: E-KYC Powered by HOLMES

Source: Wipro
Note: modules include information extraction, hierarchy creation, cognitive document download, OCR, case report generation, evidence and annotation, and connector APIs.

A US-based global bank has implemented Wipro's E-KYC in limited scale production and has been able to reduce cycle time by 60% (from 6–7 hours to 1–2 hours), processing effort between 35% to 40%, and improved user experience by 30% (i.e., reduced client back and forth for data collection reduced by 30%). In addition, Wipro sees potential to significantly reduce client's paid data costs.

Wipro has onboarded the US, UK, Australia, and India, and plans to onboard Canada, a few Asian countries, Latin America, continental Europe, and the Middle East in the next few quarters. As part of these rollouts it will add select languages. On its 2018–2019 roadmap is developing transaction monitoring and expanding in AML, including screening for negative news. These developments will include a new workflow and reporting module and leverage of technologies for faster extraction and continuous learning.

Outside of KYC, Wipro is planning to apply HOLMES to documentary trade finance and is developing an intelligent solution for letters of credit validation. The solution will provide the banks with the ability to validate letters of credit against the documents provided per the ICC guidelines and to prepare reports with supporting evidence for audit purposes and any needed manual validation.

WORKFUSION

WorkFusion (described on page 37) is also tackling automation of sanctions screening (including geocoding), file build-up for KYC analysts, and payments fraud. As in all use cases, WorkFusion combines a variety of technologies to automate a process. For example, it is automating sanctions screening at a global bank, which was handling close to 2 million payments per month flagged for investigation. WorkFusion applied its SmartApp, which orchestrates the optimal tool for specific parts of the process. The results are shown Figure 39, which contrasts the manual process (current state) with WorkFusion's solution. The scope was "like for unlike" false alerts for location to company, individual to company, and company to vessel. WorkFusion is on track to reduce costs by $400,000 for this bank within six months of implementation.

Figure 39: WorkFusion Payments Sanctions Screening
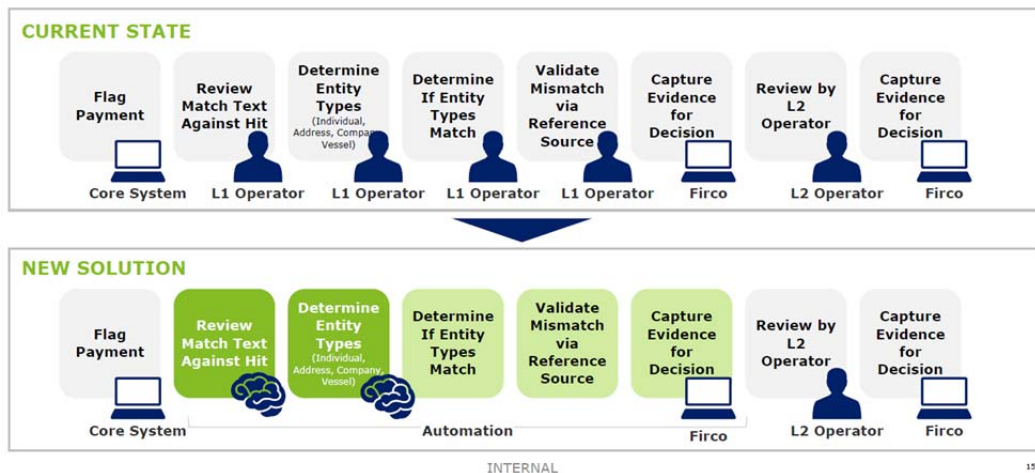
Source: WorkFusion.

On its 2018–2019 roadmap is increasing its experience across the board from onboarding KYC to AML and sanctions screening.

PAYMENTS PROCESSING

AI has been enabling critical improvements in payments processing prone to manual intervention, in particular complex cross-border payments. Key areas include repairs, correspondent bank selection, and pricing. There is ripe opportunity in payment repairs and investigations, which typically account for 75% to 80% of the labor-intensive part of a payment processing operation. Banks routinely deal with a variety of errors made by customers when they initiate payments from typos in routing numbers to the selection of an incorrect payment rail.

Two providers stood out for applying AI to payment processing: Pelican, an early mover over a decade ago, and WorkFusion.

PELICAN

Over the past decade, Pelican has been a pioneer in pushing cross-border payments straight-through processing (STP) rates past 95% by applying AI to automate manual processes. For a tier 1 bank, it increased the bank's STP rate from 85% to 97%. Its PelicanPayments STP applies machine learning and NLP to deliver repair logic, customized routing logic, and intelligent reporting, including dashboards on countries, currencies, and correspondents, and full details of all repair activity the system has undertaken. It can make repairs across any field in a message or file and can normalize unstructured data (including correction of typos) and translate these into machine-readable text. It learns from the errors operators correct, identifying patterns of repetitive behavior (e.g., a common repair of a payment instruction) and then developing auto repair tools, which can be applied at the transaction level or down to individual pieces of information (e.g., country of beneficiary). For payment routing, the platform routes based on business objectives, typically least cost or fastest route. It provides an audit trail and log to make the routing choices transparent.

In addition to automating the processing, PelicanPayments STP provides a revenue generation platform powered by NLP, which enables a bank to customize pricing for its correspondent bank clients as well as enable a client to charge for payment repairs or alternatively waive a fee (e.g., for disaster relief payments). Traditionally, a fixed, standard fee is deducted from the payment. Instead, a bank can enable its correspondent bank clients to charge a customer-specific fee based on the payment information (e.g., purpose, geography, and beneficiary). To enable customized pricing, Pelican uses AI to integrate the payments system to the bank's billing system. Figure 40 outlines the steps in setting up the revenue generation system from selecting the product, fee codes (which could be customized for individual clients), and the key words that identify the product code to setting up payment parameters and assigning the profile to the correspondent bank's customer.
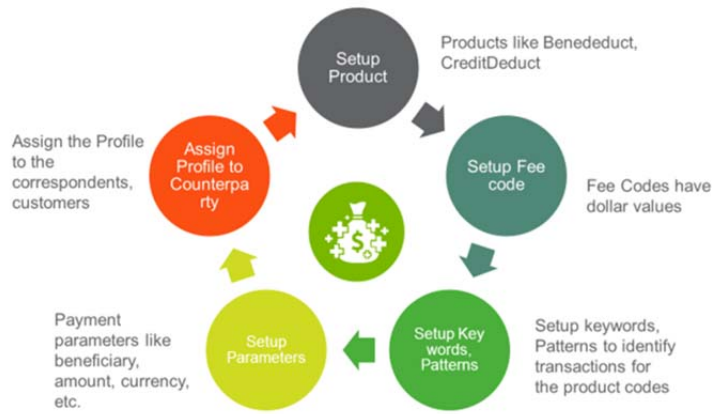
Source: Pelican.

During 2018–2019, Pelican will provide additional integration between the repair function and the API management layer. By offering payment repair service as an API, banks and other players can offer repair-as-a-service.

WORKFUSION

WorkFusion (described on page 37) has automated payments processing and payments repairs. Figure 41 outlines how WorkFusion applies RPA, AI, and BPM to automate bank name repair. For a global bank, it automated 95% of a manual process with a 96% accuracy rate, resulting in an 86% automation rate.

Figure 41: WorkFusion Payment Repairs for STP
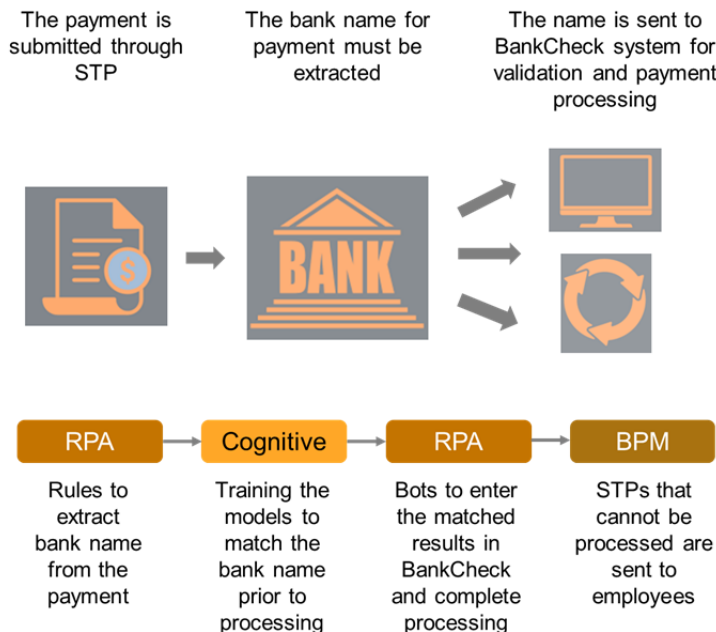
Source: WorkFusion.

PATH FORWARD

Over the next five years, more AI summits will come into view, and a few will be reached. Celent expects that AI will increasingly become a key enabler in the realization of banks' digital strategy. Bank-customer engagement will be transformed as virtual assistants prove themselves and become more commonplace. Relationship managers will increasingly embrace AI as a "partner," increasing their productivity while making their jobs more interesting. AI will stave off the inflation in operations and compliance costs.

These summits will only be reached if critical prerequisites are in place. Successful adoption of AI rests on much more than the technology. Banks which want to scale the advanced AI summits must have in place:

- Beginning at the top of the house, transformation-driven leadership with a digital strategy and willingness to embrace new paradigms and businesses based on assumptions about the future rather than the facts of the past.

- A collaborative alchemy across the business side, data team, IT, and compliance.

- An enterprise-wide AI initiatives team which ensures that all types of AI use cases can be examined, learnings shared, and best practices emulated.

- Strong data governance to reduce data preparation time and maximize the value-add of data scientists and system specialists.

- Strong model governance and explainable AI to ensure ethical behavior, privacy, and robust permissioning.

Aspiring AI summiteers must also have patience and perseverance. As Reinhold Messner, one of the few people to have climbed all fourteen peaks above 8,000 meters, said: "It's always further than it looks. It's always taller than it looks. And it's always harder than it looks." And he kept climbing.

For a further discussion of the prerequisites and examples of successful AI implementations, please see the upcoming Celent report, *Vision to Reality: AI in Action*.

Was this report useful to you? Please send any comments, questions, or suggestions for upcoming research topics to info@celent.com.

LEVERAGING CELENT'S EXPERTISE

If you found this report valuable, you might consider engaging with Celent for custom analysis and research. Our collective experience and the knowledge we gained while working on this report can help you streamline the creation, refinement, or execution of your strategies.

SUPPORT FOR FINANCIAL INSTITUTIONS

Typical projects we support related to small business, middle market, and corporate banking include:

Vendor short listing and selection. We perform discovery specific to you and your business to better understand your unique needs. We then create and administer a custom RFI to selected vendors to assist you in making rapid and accurate vendor choices.

Business practice evaluations. We spend time evaluating your business processes, particularly in implementation and product rollout. Based on our knowledge of the market, we identify potential process or technology constraints and provide clear insights that will help you implement industry best practices.

IT and business strategy creation. We collect perspectives from your executive team, your front line business and IT staff, and your customers. We then analyze your current position, institutional capabilities, and technology against your goals. If necessary, we help you reformulate your technology and business plans to address short-term and long-term needs.

SUPPORT FOR VENDORS

We provide services that help you refine your product and service offerings. Examples include:

Product and service strategy evaluation. We help you assess your market position in terms of functionality, technology, and services. Our strategy workshops will help you target the right customers and map your offerings to their needs.

Market messaging and collateral review. Based on our extensive experience with your potential clients, we assess your marketing and sales materials — including your website and any collateral.

RELATED CELENT RESEARCH

Understanding the Investment in AI in Banking
July 2017

Trends in Corporate Banking: From Disruption to Transformation
June 2017

Strategies for Enhancing Corporate Client Experience: The Future of Attended Channels
January 2017

Challenging the Status Quo: External Forces Supporting the Rise of Non-Bank Financial Services
September 2016

Faster Than a Speeding Payment: The Race to Real-Time Is Here
June 2016

Top Trends in Corporate Banking 2016: Turning Innovation into Results
March 2016

Breaking the Payments Dam: External Forces Transforming the Payments Ecosystem
November 2015

Top Trends in Corporate Banking 2015: The Business of Managing Uncertainty
April 2015

Eastern Bank: Celent Model Bank of the Year 2016
April 2016

Copyright Notice

Prepared by

Celent, a division of Oliver Wyman, Inc.

For more information please contact info@celent.com or:

Alenka Grealish   agrealish@celent.com

AMERICAS

USA

200 Clarendon Street, 12th Floor
Boston, MA 02116

Tel.: +1.617.262.3120
Fax: +1.617.262.3121

USA

1166 Avenue of the Americas
New York, NY 10036

Tel.: +1.212.541.8100
Fax: +1.212.541.8957

USA

Four Embarcadero Center, Suite 1100
San Francisco, CA 94111

Tel.: +1.415.743.7900
Fax: +1.415.743.7950

Brazil

Av. Doutor Chucri Zaidan, 920 –
4º andar
Market Place Tower I
São Paulo SP 04578-903

Tel.: +55.11.5501.1100
Fax: +55.11.5501.1110

Canada

1981 McGill College Avenue
Montréal, Québec H3A 3T5

Tel.: +1.514.499.0461

EUROPE

France

28, avenue Victor Hugo
Paris Cedex 16
75783

Tel.: +33.1.73.04.46.20
Fax: +33.1.45.02.30.01

United Kingdom

55 Baker Street
London W1U 8EW

Tel.: +44.20.7333.8333
Fax: +44.20.7333.8334

Italy

Galleria San Babila 4B
Milan 20122

Tel.: +39.02.305.771
Fax: +39.02.303.040.44

Spain

Paseo de la Castellana 216
Pl. 13
Madrid 28046

Tel.: +34.91.531.79.00
Fax: +34.91.531.79.09

Switzerland

Tessinerplatz 5
Zurich 8027

Tel.: +41.44.5533.333

ASIA

Japan

The Imperial Hotel Tower, 13th Floor
1-1-1 Uchisaiwai-cho
Chiyoda-ku, Tokyo 100-0011

Tel: +81.3.3500.3023
Fax: +81.3.3500.3059

China

Beijing Kerry Centre
South Tower, 15th Floor
1 Guanghua Road
Chaoyang, Beijing 100022

Tel: +86.10.8520.0350
Fax: +86.10.8520.0349

Singapore

8 Marina View #09-07
Asia Square Tower 1
Singapore 018960

Tel.: +65.9168.3998
Fax: +65.6327.5406

South Korea

Youngpoong Building, 22nd Floor
33 Seorin-dong, Jongno-gu
Seoul 110-752

Tel.: +82.10.3019.1417
Fax: +82.2.399.5534