World Academy of Science, Engineering and Technology
International Journal of Computer and Information Engineering
Vol:15, No:7, 2021

Reference Architecture for Intelligent Enterprise Solutions

Shankar Kambhampaty, Harish Rohan Kambhampaty

***Abstract*—**Data in IT systems in enterprises have been growing at phenomenal pace. This has provided opportunities to run analytics to gather intelligence on key business parameters that enable them to provide better products and services to customers. While there are several Artificial Intelligence/Machine Learning (AI/ML) and Business Intelligence (BI) tools and technologies available in marketplace to run analytics, there is a need for an integrated view when developing intelligent solutions in enterprises. This paper progressively elaborates a reference model for enterprise solutions, builds an integrated view of data, information and intelligence components and presents a reference architecture for intelligent enterprise solutions. Finally, it applies the reference architecture to an insurance organization. The reference architecture is the outcome of experience and insights gathered from developing intelligent solutions for several organizations.

***Keywords*—**Architecture, model, intelligence, artificial intelligence, business intelligence, AI, BI, ML, analytics, enterprise

I. Introduction

CONSUMERS are demanding information *relevant to their context, anywhere, anytime,* on any device [1]. To support this demand, enterprise IT systems store and process increasingly large amounts of data, extract intelligence from it and deliver the right business information at the right time to consumers to provide differentiated products and services. With the rise of Artificial Intelligence (AI/ML) and BI tools and technologies in recent years, analytics solutions are being developed to extract intelligence from data. The problem, however, is that point solutions for analytics are being developed in several enterprises without taking a holistic and integrated view of data and information in enterprise solutions in the context of current and future needs of the business. This has led to "patch work" of analytics systems that do not support reuse and increase cost in the long run.

The purpose of this paper is to provide reference architecture to address the above problem and enable intelligent enterprise solutions. To illustrate its use by enterprises, it has also been applied to an insurance organization. The reference architecture has evolved from the experience and insights gathered from developing intelligent solutions for several enterprises.

Shankar Kambhampaty is a DXC Distinguished Architect and Chief Technology Officer (CTO) for a major account at DXC Technology, Hyderabad, India (e-mail: skambhampaty@gmail.com).

Harish Rohan Kambhampaty is a student of Bachelor of Technology (B.Tech.) programme in Computer Science & Engineering at Gokaraju Rangaraju Institute of Engineering & Technology (GRIET), Hyderabad, India.

II. Reference Model for Enterprise Solutions to Enable Digital Transformation

Data are comprised of text, video, images that may be structured (e.g., data in relational databases), semi-structured (e.g., XML, JSON) or unstructured (e.g., social media data) that originates from various business applications, infrastructure systems and IoT enabled devices. When the raw data (e.g., stock price) are made useable with additional attributes (e.g., price going up/down) then it becomes *information*. Information in a context (e.g., stock price going up in NYSE) is *intelligence* [2].

With continuous technology and business-related disruptions in the marketplace across vertical market segments, enterprises are meeting the consumer demand of anywhere, anytime information related to their products and services by taking digital transformation initiatives and architecting their IT systems to include the principles of Bi-model IT [3]. The architecture of enterprise solutions, is therefore, gravitating towards having *systems of record*, that adopt one mode to develop and manage solutions (traditional and sequential manner, emphasizing safety and accuracy) and *systems of engagement* that adopt another mode (exploratory and nonlinear, emphasizing agility and speed) [4]. The *systems of record* expose functionality through APIs that are invoked by *consumer apps* that represent *systems of engagement*.

Based on the experience from architecting enterprise solutions for several clients, a reference model for enterprise solutions, depicted in Fig. 1, has been developed that can serve as a starting point for development of enterprise-level architecture for intelligent solutions.

Systems of record are commercial or home-grown products, that address core needs of the business (e.g., policy administration, claims and billing in insurance vertical) and are typically deployed on-premise, while *systems of engagement* are solutions that are best deployed on cloud platforms to meet consumer demand of business information anywhere, anytime. *Systems of engagement* include *consumer apps* with mobile or web frontends that may have solution components on the device that the consumers use, as well as, solution components on cloud/on-premise environments to ensure security, scalability and availability. These components also invoke the Application-Program Interfaces (APIs) in the *integration layer* exposed by enterprise business applications that constitute the *systems of record*.

Note: The direction of arrows in the figures in this paper represents the direction of invocation of the component and not data flow.

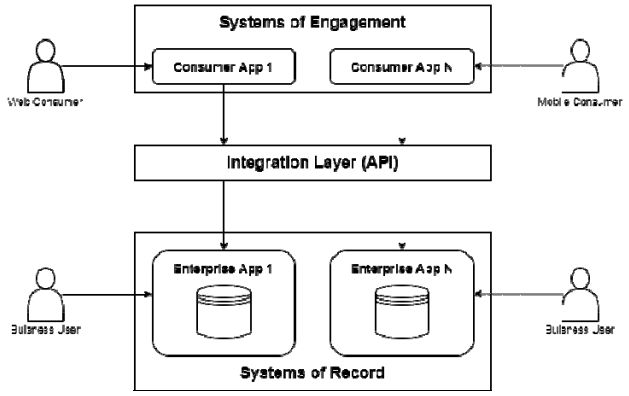

Fig. 1 Reference Model for Enterprise Solutions

The data and information related solution components are foundational to fulfilling the consumer demands and meeting the goals of the business to provide competitive products and services to consumers. The next section fleshes out the data and information components in the reference model in Fig. 1 before introducing the intelligence components into the architecture in subsequent sections.

III. Data and Information Solutions in an Enterprise

Data for *systems of record* are first stored in Online Transaction Processing (OLTP) data stores (both SQL and NoSQL). These data are then aggregated and transformed into information through a variety of solutions including operational data stores (typically snapshots from multiple transactional systems for ad-hoc reporting purpose). The data and information components for *systems of record* may be abstracted into a *data layer*. This layer includes the back-end data stores of enterprise applications. The enhanced reference model for enterprise solutions with data and information solution components is shown in Fig. 2.

The enhanced reference model elaborates the *integration layer* and depicts integration middleware component and API Gateway components. The API Gateway exposes APIs that are consumed by the *consumer app* solution components. The *integration layer* serves as a glue for all components. A *data visualization layer* is introduced in the architecture that enables enterprise applications and *systems of engagement* to get views of data in the right format. An Operational Data Store (ODS) may also be used to create a snapshot of data of interest from OLTP data stores for ad-hoc reporting purpose, which may further be exposed through integration middleware to consumer applications or enterprise applications.

IV. Need and Opportunity for Intelligent Solutions

Data in IT systems of enterprises have been growing at a tremendous pace which can provide significant value to their business growth. It presents an opportunity for enterprises to analyze data and develop valuable insights about the consumers. When consumers demand information related to their interest, anytime, anywhere and on any device, they are essentially asking providers of products and services to understand who they are, in terms of their demographic profile, what products/services they have already availed, what their preferences are, and use this knowledge to lead them towards right products/services that they may be interested in (or solve a problem they may be experiencing). This calls for a level of response often associated with *intelligence*. Therefore, enterprises need to architect and implement intelligent solutions.

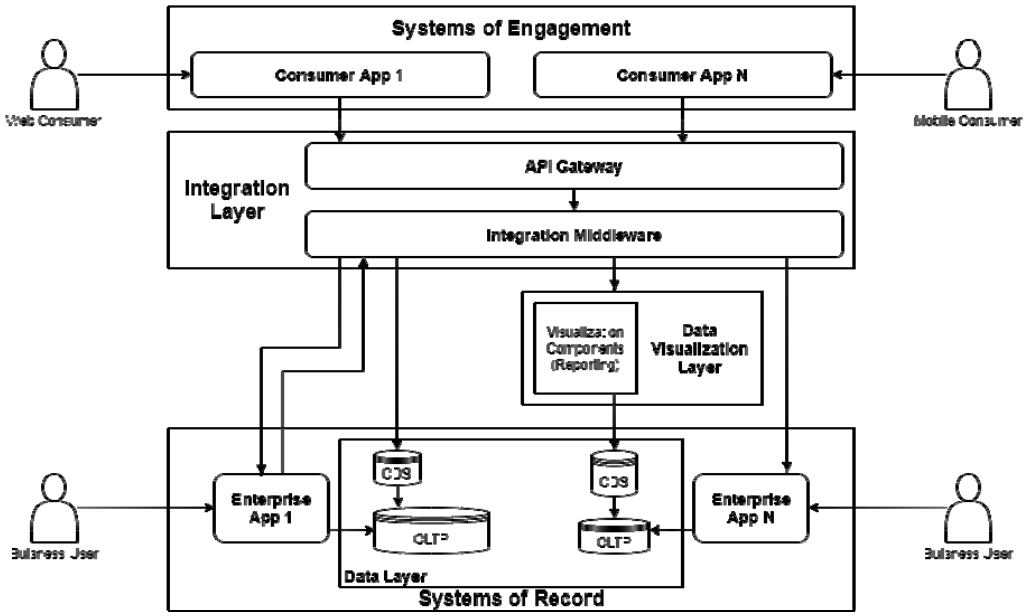

Fig. 2 Reference Model with Data and Information solution components

For enterprises to build intelligent solutions, it becomes necessary to pull together relevant data related to the consumers, determine attributes of interest to both consumers and enterprise, and establish relevant context for the

information. *Intelligence* is information in a *context*. This capability of establishing the right context and providing information to consumers makes the enterprise solution *intelligent*.

Intelligent solutions are developed by determining several attributes relevant to products/services that enterprises offer and establishing a right context for consumers. Key attributes for enterprise applications are as follows [5]:

1. Location – From what location is the consumer making the request? For instance, if the consumer is looking for an insurance quote from a specific country, the quote provided must take into account the local laws.
2. Device – What device is the consumer making the request from? If the consumer is requesting information from an iPhone, the response provided must be formatted to support the device.
3. Security – Is the consumer entitled to the requested information? Consumers need privileges to access certain information. The profile of the consumer needs to be determined and seamless authentication and authorization performed to provide them access only to the information to which they are entitled to.
4. Product/Service – What product/service is of interest to the consumer? Consumers repeatedly engage with systems and look for information that meets their goals. The solution will need to learn from user interactions and provide responses that address their implicit needs.

The data required to build intelligent solutions for enterprises are primarily derived from three sources:

1. Business Applications – Business data (e.g., policy or claims in insurance) from OLTP data stores are extracted, transformed and loaded into Data lakes, Data Marts and Data warehouses to extract intelligence.
2. Infrastructure Components – Operational data from infrastructure components (e.g., using monitoring components like Dynatrace) are pulled into tools (e.g., Splunk) to extract intelligence to improve the efficiency of data center infrastructure.
3. Devices – Streaming data from machines and vehicles (e.g., telemetry device fitted to a car) is used to extract intelligence and send relevant notifications to consumers.

V. Intelligent Solutions for the Enterprise

Investopedia defines data analytics as "science of analyzing raw data in order to make conclusions about that information" [6]. Data analytics techniques have been used to extract intelligence. The intelligence so extracted may be used to refine business decisions in an enterprise.

Traditionally, there have been four types of data analytics applied in enterprises to extract intelligence. More recently, there has been an emergence of the fifth type of data analytics. The five types of data analytics applied in enterprises are shown in Table I [7].

The five types of data analytics techniques to extract intelligence find wide application in enterprises across vertical market segments. Table II presents a summary of intelligent enterprise solutions in key market segments [8].

TABLE I
Types of Data Analytics

Data Analytics	Purpose	Technologies/Tools
Descriptive	To determine what has happened (e.g., Year-on-Year product sales by product category/ region/customer segment)	Traditionally BI (e.g., MS SSAS, IBM Cognos®, MicroStrategy, SAP Business Objects)
Diagnostic	To determine why it happened (e.g., Why sales have decreased or increased for a specific year)	Traditionally BI (e.g., Qlik Sense, Tableau Desktop, Chartio, InsightSquared)
Predictive	To determine what will happen next (e.g., What would sales be next quarter in a specific area and who will buy)	Big Data (e.g., Hadoop, R), AI/ML (e.g., Python with SciKit/TensorFlow, AWS Sagemaker, Azure ML)
Prescriptive	To determine what should be done (e.g., What types of products pair well together and how to price products)	AI/ML (e.g., Python with SciKit/TensorFlow, AWS Sagemaker, Azure ML)
Cognitive	To determine what to do, why and how (e.g., What types of products pair well together, how to price products and how to course correct based on customer buying behavior)	AI/ML (e.g., Python with SciKit/TensorFlow), AWS Sagemaker, Azure ML, Azure Cognitive Services, AWS Cognitive services

TABLE II
Intelligent Solutions across Vertical Segments

Vertical Market Segment	Intelligent Solutions
Banking, Financial Services & Insurance	Trading Analytics, Fraud Analytics, Customer retention Analytics, Loss ratio Analytics, Anti-Money laundering, Customer 360, Customer financial/insurance risk analytics,
Healthcare, Life Sciences	Predictive intervention, Fraud & Abuse detection, Disease Risk intervention, Patient 360
Manufacturing, Automotive	Market Analytics, Location Analytics, Predictive Maintenance, Supply Chain Optimization
Travel, Transportation, Logistics	Passenger Analytics, Revenue Integrity Analytics, Intelligent Freight, Location Intelligence
Retail	Retail Analytics, Predictive Supply Chain, Customer Engagement, Loss Prevention
Energy	Smart Meters, Predictive Maintenance, Location Intelligence, Performance Optimization

The technologies that enable intelligence in solutions in the enterprise, broadly, fall under two categories, BI and AI/ML, which are explained in next couple of sections.

VI. Intelligent Solutions with BI Technologies

As shown in Table I, BI technologies are well suited to perform *descriptive* and *diagnostic* analytics to determine what happened and why it happened for the topic of interest to the business user. These technologies provide the capabilities of Online Analytical processing (OLAP) that have been around more than two decades, and are mature with significant skill base in enterprises. The BI technologies may be grouped under two categories:

1. Multi-dimensional Stores
2. Visualization tools

The data from various OLTP transactional data stores are typically maintained in an intermediary area, called ODS that facilitates ad-hoc reporting for operational users in near real-time and also helps as staging area for subsequent data

transformation for building Data Marts and Datawarehouses. The Data Mart or Datawarehouses are designed using dimensional modeling principles consisting of a set of *dimension* and *fact* tables. The processed data in Data Mart/Datawarehouse are further aggregated in the form of multi-dimensional structures, called *cubes* in OLAP servers (e.g., Microsoft SSAS). Reporting/visualization tools that use MDX queries and other techniques make sense of the data and enable extracting intelligence on what happened and why it happened for the topic of interest to the business user. The commonly used tools are Microsoft SQL Server Reporting Services (SSRS) and Microsoft PowerBI, along with Qlik, Google Charts, FusionCharts and Tableau.

While the Data Marts/Datawarehouses are meant to store structured relational data, the data from various transactional systems can also be maintained in Data Lake which has become very popular with the advent of Big Data. The main characteristic feature of Data Lake is maintenance of structured data and unstructured data in its native form. The data schema is determined for Data Lake at the time of reading the data (schema-on-read) while the data schema for Data Marts/Data warehouses is determined at the time of loading the data (schema-on-write).

The first step towards developing an integrated view of data, information and intelligence in an enterprise is to study the use of ODS, Data Mart, Datawarehouse, and Data Lake in business units and IT groups and make right architecture decisions, technology choices and investments.

VII. Intelligent Solutions with AI/ML Technologies

The advances made in AI/ML technologies in the past decade are revolutionizing, not only how intelligence is extracted from data, but also the nature of products and services offered to consumers. Information relevant to consumer context, anywhere, anytime, on any device has not only become a reality but has raised consumer demands even more for enhanced capabilities in products and services.

The AI techniques before the advent of this millennium were built on *rule-based* and *case-based reasoning* approaches. The emergence of cloud, mobile and the internet of things (IoT) resulted in data being generated at web-scale that accelerated the adoption of ML techniques to extract intelligence. The ML techniques are a set of algorithms implemented in programs that can learn from data and improve from experience, without human intervention [9]. The learning techniques result in creation of a ML *model.*

Fig. 3 depicts the approach to build a ML model. At a high-level, there are three steps:

1. Capture and Cleanse data
2. Build ML Model
3. Visualize

Step1. Capture and Cleanse data: The historical data for key variables of a given domain are captured. For instance, for automobile insurance domain, historical data related to key variables such as policies, claims, customers and vehicles are captured, and issues related to incorrect/ missing data are fixed.

Step2. Build ML Model: In this step, feature variables and target variables are identified. The historical data collected is split into training and test data. The model is trained using the training data and the accuracy of the prediction is tested with test data and scored by computing performance metrics. Should the accuracy not be at the required level, the model is re-trained and if that does not work, feature variables and algorithms used to train the model or changed until the accuracy of prediction is satisfactory.

Step3. Visualize: Once the model is trained for a set of feature variables with training/test data, it forms the basis to perform *predictive, prescriptive* and *cognitive* data analytics and extracting intelligence.

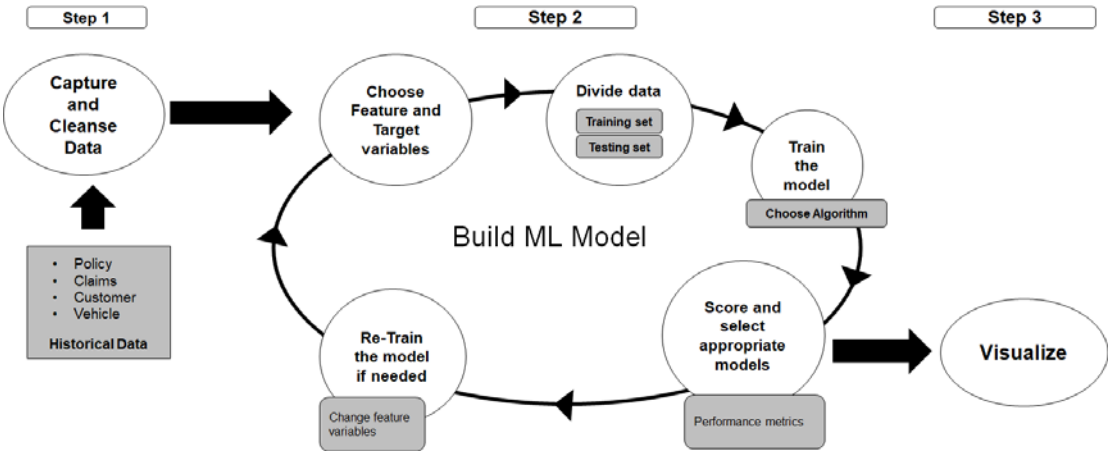

Fig. 3 Approach to build a ML Model

The learning techniques broadly fall under four categories as depicted in Table III [10].

There are several AI/ML technologies, tools and frameworks that are available as commercial products from

software vendors, as services on cloud platforms, and as open-source resources. Intelligence solutions in an enterprise may be built with them. Table IV provides the key products for AI/ML available from these sources [11].

TABLE III
ML TECHNIQUES

Learning Technique	Type of Data	Algorithms	Application
Supervised	Structured data	k-Nearest Neighbor	Risk Analytics
		Naive Bayes	Fraud Analytics
		Decision Trees	
		Linear Regression	
		Support Vector Machines (SVM)	
		Neural Networks	
Unsupervised	Unstructured data (Text)	k-means clustering	Sentiment Analysis
		Association Rules	
Reinforcement Learning	Unstructured data (Text)	Q-Learning	Computer played board games (Chess, Go)
		Temporal Difference (TD)	Robotic hands
		Deep Adversarial Networks	Self-driving cars
Deep Learning	Unstructured data (Text, speech, image)	Deep Neural Networks (DNN)	Image Recognition Voice Recognition

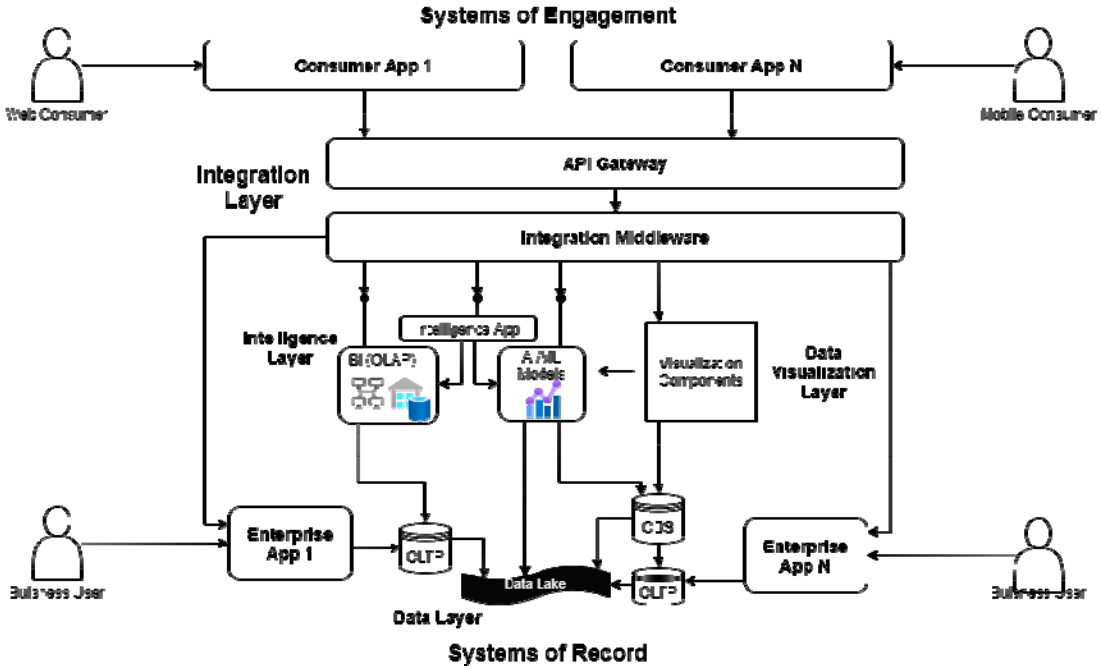

Fig. 4 Reference Architecture for Intelligent Enterprise solutions

TABLE IV
AI/ML TECHNOLOGIES

Source	Product
Open Source	Compose, Core ML Tools, Cortex, Featuretools, GoLearn, Gradio, H2O, Oryx, PyTorch Lightning, TensorFlow, MxNet, Scikit-learn, Shogun, Spark MLlib, Weka, Caffe, Anaconda (Individual Edition)
Commercial Cloud Services	Anaconda (Team Edition & Enterprise Edition) AWS AI/ML Services • ML Services – Amazon SageMaker • AI Services – Amazon Poly, Lex, Rekognition, Comprehend Azure AI Services • Azure ML • Cognitive services (search, Vision, Speech, Text) • Azure Bot Framework • Azure Databricks Google AI Services • Cloud ML Engine • BigQueryML

VIII. Reference Architecture for Intelligent Enterprise Solutions

Putting it all together, Fig. 4 depicts the reference architecture for intelligent enterprise solutions.

The reference architecture for intelligent enterprise solutions has *systems of record* that include enterprise applications with core business functions that change less frequently and *systems of engagement* that interact with consumers and change more frequently. The *systems of record* expose functionality through APIs that are invoked by *consumer apps* that constitute *systems of engagement*.

The *systems of record* are typically deployed on-premise, while *systems of engagement* are deployed on highly scalable and available infrastructure such as cloud platforms. The *systems of engagement* support mobile and web front ends and devices that consumers use to interact with the enterprise (e.g.,

Smart speakers such as Amazon Alexa, Google Home or the Apple HomePod) to meet consumer demand of business information anywhere, anytime. These components invoke the APIs in the *integration layer* exposed by enterprise business applications that constitute the *systems of record*.

The *integration layer* consists of integration middleware component and API Gateway components. The API Gateway exposes APIs that are consumed by the *consumer app* solution components. The *integration layer* serves as a glue for all components. A *data visualization layer* enables enterprise applications and *systems of engagement* to get views of data in the right format.

Intelligence in solutions in an enterprise is extracted by applying AI/ML and BI techniques on the data and information stores. The solution components that implement these techniques may be abstracted into a separate layer, referred to as *intelligence layer*. It contains both AI/ML and BI components whose functionality is exposed as services and consumed through the integration middleware both by enterprise applications and consumer apps. A Data Lake has also been included in the *data layer*. AI/ML solution components access these data for training data, test data and for performing analytics to extract intelligence and formulate the right response to requests both from consumers and enterprise applications. Introduction of the intelligence components (AI/ML & BI) enables development of new apps to extract intelligence and provide new capabilities to meet consumer demands. The *intelligence app* [12] shown in Fig. 4 performs that function.

IX. Applying Reference Architecture for Intelligent Enterprise Solutions to Insurance Organization

This section illustrates how the reference architecture defined in the earlier section may be applied to an enterprise. An insurance organization has been chosen for this purpose. Microsoft technologies have been given as examples for different components. Fig. 5 depicts the reference architecture at a high-level for a generic insurance organization.

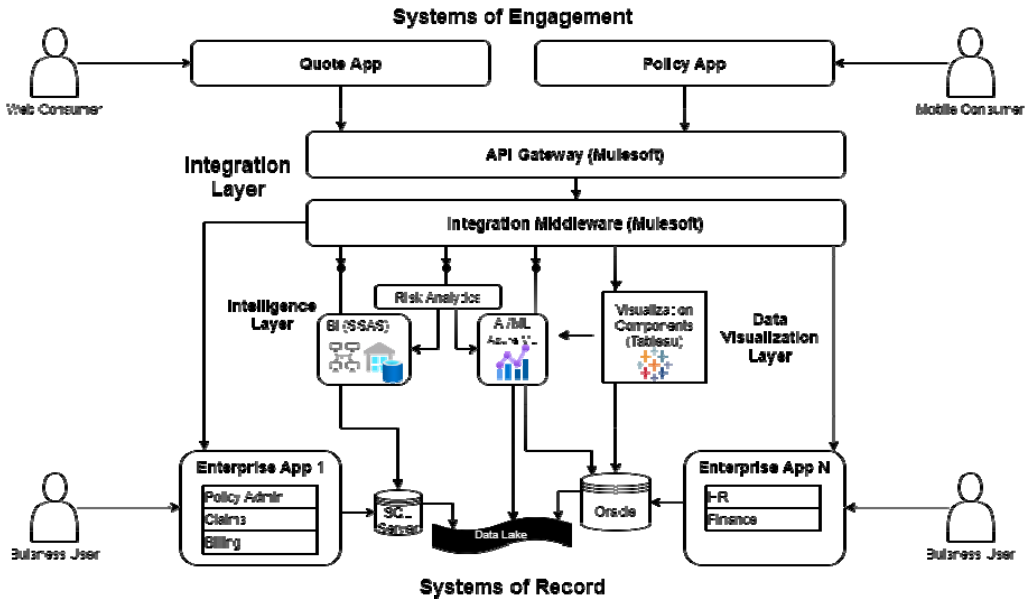

Fig. 5 Reference Architecture for Intelligent Enterprise solutions for Insurance organization

In the insurance domain, due to digital disruption, there is a clear shift to some portion of agent led business to digital platforms (e.g., Compare in US). The consumer compares premium to be paid for the required coverage among the participating insurance companies and decides on the insurance organization to go with. In addition, there is a clear trend for consumers to pay renewal premium online than go through an agent/broker. Also, in some cases consumers use devices such as smart speakers (like Amazon Alexa, Google Home or the Apple HomePod) and so the response provided by enterprise solution components must match the request from consumers both in format and content.

In the past, insurance companies would use a set of rules to offer quotes or specify renewal premiums. But given that customer acquisition and retention is key for insurance business, insurance companies are building increasing levels of intelligence to their enterprise solutions to provide highly customized quotes leveraging the intelligence to be competitive in the marketplace. The reference architecture addresses these requirements through an *integration layer* that exposes APIs. The APIs are invoked by the mobile, web or smart speaker consumer applications in *systems of engagement* as well as third party digital platforms to provide quotes to potential customers.

The API invocation results in AI/ML models and BI components being called to extract the intelligence on the fly

to service the consumers.

Enterprise Applications in an insurance organization have policy admin, claims and billing modules that implement the lead to quote, quote to policy, FNOL to claims and billing workflows. They have been shown to have SQL Server as back-end. Likewise, ERP systems implement the HR and Finance functionality. They have been shown to have Oracle Database server as back-end. These constitute the *systems of record*. SSAS server performs OLAP functions and Azure ML services build models for predictive analytics. Risk analytics is a good candidate for an *intelligence app* of an insurance organization. The intelligence so gathered may be used by the policy admin module of the enterprise application to provide a more intelligent and competitive quote to consumers. In this way, the reference architecture with suitable modifications and enhancements and use of appropriate technologies may be applied to any enterprise.

X. Implementation of Reference Architecture

The reference architecture for intelligent enterprise solutions is best suited to be implemented at enterprise-level. To that end, the architecture groups within enterprises may consider its adoption as part of their enterprise architecture initiatives.

The following steps are recommended in implementing the Reference Architecture:

1) Review of the IT Roadmap of the enterprise.
2) Determination of strategic importance of intelligence technologies in meeting the organization goals.
3) Assessment of maturity of implementation of data and information solutions within the enterprise.
4) Identification of use cases to implement intelligence solutions.
5) Development of business case to justify the investment required for implementation.
6) Execution of pilot to provide the technology choices and ROI.

XI. Conclusion

Consumers increasingly demand information relevant to their context, anywhere, anytime, on any device. A structured well-thought-out architecture that takes an integrated view of data, information and intelligence components is key to be able to support the demand. Intelligence components include both AI/ML and BI and both are needed to support the intelligence needs of the enterprise. The reference architecture presented in this paper serves as a good starting point for an organization to define or refine its enterprise-level architecture and plan steps for implementation.

Disclaimer

This disclaimer informs readers that the views, thoughts, and opinions expressed in this paper belong solely to the authors, and not necessarily to the author's employer, institution or other groups or individuals.

References

[1] S. Kambhampaty, "Reference Architecture for SMAC solutions", Proceedings of the Sixth International Conference on Computer Science and Information Technology (CCSIT - 2016), 2016. https://arxiv.org/abs/1601.06389

[2] S. Kambhampaty, "It's All About Data", Forbes article, 2018. https://www.forbes.com/sites/forbestechcouncil/2018/07/23/its-all-about-data/?sh=f41d58e68556

[3] D. Scott, S. Mingay "Scaling Bimodal — Fusing IT With the Business", A Gartner Trend Insight Report. ID: G00331680, 2017. https://www.gartner.com/en/doc/3772092-scaling-bimodal-fusing-it-with-the-business-a-gartner-trend-insight-report

[4] MPulse, "System of Record vs. System of Engagement, Part 1: What Maintenance Managers Need to Know about ERP Systems", 2016. https://mpulsesoftware.com/blog/maintenance-management/system-record-vs-system-engagement-part-1/#:~:text=A%20System%20of%20Record%20is,data%20repository%20of%20an%20organization.&text=A%20System%20of%20Engagement%20is,usable%20tool%20for%20capturing%20data

[5] S. Kambhampaty, "'Beam Me Up, Scotty:' The Next-Generation User Experience", Forbes article, 2019. https://www.forbes.com/sites/forbestechcouncil/2020/04/20/beam-me-up-scotty-the-next-generation-user-experience/?sh=3b7efac41b71

[6] J. Frankenfield, "Data Analytics", 2020. https://www.investopedia.com/terms/d/data-analytics.asp

[7] E. Wilson, "The differences between descriptive, diagnostic, predictive & Cognitive analytics", 2020. https://demand-planning.com/2020/01/20/the-differences-between-descriptive-diagnostic-predictive-cognitive-analytics

[8] D. Gupta, *Data Analytics and Its Application in Various Industries*, Apress, Berkeley, CA, ISBN: 978-1-4842-3524-9, 2018. https://link.springer.com/chapter/10.1007/978-1-4842-3525-6_1

[9] R. Shaw, "The 10 Best Machine Learning Algorithms for Data Science Beginners", Dataquest, 2019. https://www.dataquest.io/blog/top-10-machine-learning-algorithms-for-beginners/

[10] D. Fumo, "Types of Machine Learning Algorithms You Should Know", towards data science, 2017. https://towardsdatascience.com/types-of-machine-learning-algorithms-you-should-know-953a08248861

[11] S. Yegulalp, "14 open source tools to make the most of machine learning", InfoWorld, 2020. https://www.infoworld.com/article/3575420/14-open-source-tools-to-make-the-most-of-machine-learning.html

[12] delaware. "Intelligent apps: the next generation of applications", 2020. https://www.delawareconsulting.com/en-us/solutions/intelligent-apps#:~:text=Intelligent%20apps%20are%20applications%20that,personalized%20and%20adaptive%20user%20experiences

Shankar Kambhampaty is a DXC Distinguished Architect and Chief Technology Officer (CTO) for a major account at DXC Technology, Hyderabad, India. He is author of several international conference papers including the book titled, Service Oriented Architecture and Microservice Architecture for Enterprise, Cloud, Big Data and Mobile, Wiley India publication.

Harish Rohan Kambhampaty is a student of Bachelor of Technology (B.Tech.) programme in Computer Science & Engineering at Gokaraju Rangaraju Institute of Engineering & Technology (GRIET), Hyderabad, India. He is a member of the Advanced Academic Center (AAC) and an active contributor to AI/ML, web development and mobile solutions.