Mindtree

TECH BEACON

THE TECHNOLOGY HARBINGER

Artificial Intelligence

Internet of Things

Conversational Apps

Cloud Computing

Blockchain

Application Development

eXtended Reality

ABOUT

TECH BEACON

THE TECHNOLOGY HARBINGER

Technology is core to Mindtree. It is extremely important in building capability and capacity for us to keep a sharp focus on emerging technologies. We also make it a point to understand the application of these technologies in the context of our customers. In this effort, we constantly explore and experiment with emerging technologies and gauge their maturity levels for consumption by enterprises. We do this by comparing and contrasting them with similar tools already in use.

Tech Beacon is an initiative to compile outcomes of our experimentation and technology adoption guidance for enterprises. We group technologies under three categories: Invest, Experiment and Watch. We explain these categories in detail in this report.

Direction

The Mindtree Tech Beacon provides direction to enterprises on various technologies by categorizing them as Invest, Experiment or Watch technologies.

Invest

These technologies have a high potential for generating business value in the near future. Mindtree intends to build capability internally or identify a partner who can help our project teams learn to bootstrap these technologies. We arrive at these technologies by observing customer demands, industry trends, opportunities and open source community support. Most of these technologies were either implemented in production-grade applications for customers, or in internal reference implementations.

Experiment

These technologies have not yet seen mainstream adoption but show huge promise for the near future. We arrive at these technologies by listening to our customers' interests, following industry buzz, and noting indications of increased activity in various forums and analyst reports. We need to experiment and build capability in these technologies to have an early adopter edge.

Watch

The technologies in this category show promise through their architecture, the business cases they support, and a community push for adoption. These technologies are not yet adopted by enterprises, but are being evaluated by enterprises for their adoption benefits. While we do not have to invest much in them yet, we need to keep a close watch on these technologies and move them to the 'Experiment' category as they mature or become more relevant to enterprises.

Table of
Contents

Architecture Patterns

Business needs are rapidly evolving and are driving architecture evolution. In this section, we present certain key business needs, and the architecture paradigms that have come about to support these business needs. We've grouped these patterns into 3 categories:

Patterns of interaction within/between systems

Distributed Architecture and Service Mesh

Increased adoption of microservices at scale necessitated common cross-service orthogonal needs like resiliency, service discovery, location transparency, security, etc to be handled in a standardized, scalable and infrastructure independent manner. Service Mesh has evolved as the answer - an abstraction layer for providing these aspects integrated with services in the compute cluster, itself configurable through APIs. By abstracting away common cross-service needs into a containerized, configurable layer, Service Mesh reduces the operational complexity that would otherwise result when using a large, complex, interconnected ecosystem of micro-services. Common implementations are Envoy, Linkerd and Istio - with the open-source Istio project getting the support of industry giants like Google, IBM and Lyft.

Reactive Autonomic Systems/Architecture

Enterprises in the last few years have become more and more technology focussed and have adopted a plethora of technologies to meet their customers direct and indirect demands. However, as architectural complexity increases, so does the chance of things not going as planned - and as architects, we are ultimately responsible for how a system behaves when faced with certain scenarios, both expected and unexpected. Autonomic computing is about engineering self-healing, self-optimizing, self-configuring and self-protecting systems to increase the robustness of software systems and minimize impact of malfunctioning systems on customers.

Reactive Architectural principles (codified in the reactive manifesto - reactivemanifesto.org), aims to deliver the how of autonomic computing - software systems that are responsive, resilient, elastic and overall, 'aware of and reactive to their context'. Reactive architecture achieves this by providing a set of guiding principles (responsiveness, resiliency, elasticity and message-driven/non-blocking style), that ensure that problems or unexpected scenarios with a single component do not compromise the integrity of the system as a whole - ensuring a better experience for the user. Given the trend of software architecture evolution that has happened in the last decade, reactive autonomic computing/architectures should be a top priority item for the CTO of an enterprise.

Architecture as a reactive activity – an anti-pattern

While the design of a solution can follow a reactive style, architecture as an activity cannot - unfortunately, the need for speed in enabling a new capability, combined with a plethora of products offering ready-to-use, cloud-native, SaaS solutions - when put together - have increased the number of 'stick-it-together, it-should-work' approaches to solution design. While time to market is certainly achieved in this model, risks/challenges like multiple copies of data, unclear systems of record, process duplication, etc. have started proliferating.

The success of this model for a target context has now become highly dependent on a combination of 2 things:
- The extent to which the enterprise's processes are realized by the selected SaaS solutions, and
- The least number of products needed to implement the complete process set.

Businesses for whom the value for the first variable (extent of fulfillment) is low or the value for the second factor (number of solutions needed) is high - cannot afford to apply architecture as a reactive activity - but should instead set aside planned time/effort for architectural design and governance to ensure success

Patterns for Information-to-Insight

Converting data to insight

While availability of processing power and adoption of machine learning algorithms have increased, an ML model is still very very specific and tuned to a certain enterprise's context - and tasks like data cleanup, outlier removal, selection of models/factors are all still very much first done by a human before the solution becomes mainly automated. This has resulted in the need for integrating exploratory data analysis tasks as part of a processing pipeline, with certain demands on the designed solution as a result. Harvesting insight from data requires synthesizing data from many sources, integrating multiple constructs like big data storage, data analysis, text mining and neural networks together before services can deliver this to consuming channels/downstream systems. Essential components for a solution include:

- A language *(R, Python or Julia)*
- Associated libraries/components *(Numpy, Matplotlib, Pandas)*
- The data store or information layer *(Spark, Hadoop, Cassandra)*
- Integration components that tie these all together *(SparkML, PySpark, Sparklyr, Python Cassandra driver)*, etc.

Lambda and Kappa Architectural Styles

As companies become more and more Digital, the availability of data pertaining to every aspect of an enterprise's journey and its surrounding context (what its partners/suppliers are doing, what its customers are doing/thinking), has grown in leaps and bounds. So has the demand to understand all of this available information and generate useful insight from them. Understanding the customer or generating insights requires two distinct stages:

- Processing a large volume of information
- Delivering the generated insight to consuming channels/downstream systems.

Architectural needs for both these stages are fairly discrete - and Lambda architecture evolved to provide a consistent pattern by separating the solution into a high-in-volume-extremely-fluid batch layer and a speed layer that provides a real-time view into the generated insight. The main challenge with Lambda was the need to maintain two branches/versions of code that dealt (in very different ways) with the same data - and Kappa is the next iteration in this evolution. While Lambda stores the data and then generates the insight in the batch layer, Kappa generates insight as-the-data-arrives within its 'streaming layer' - the only downside being that any re-generation requires a complete playback. Kappa doesnt replace Lambda - the selected pattern very much depends on the context of the solution, and exactly when the insight-generation can happen in the target context.

Distributed Ledger

As part of optimising value streams, enterprises have begun collaborating with each other and sharing data with a view to improving consistency and accuracy across their processes. The approach to enable this sharing originally evolved from file-interface-upload model to using APIs, the newer generation of businesses are turning to blockchain as a way to ensure consistent, distributed transactions - in real-time. The goal of a distributed ledger provided by Blockchain is to have a consensus of shared and replicated digital data across the global network of participating providers. This is beneficial for any distributed transaction where accuracy and consistency are paramount. This framework gives organizations a secure, digital alternative compared to manual processes built on interface/API styles of integration which needed expensive, regular, manual reconciliation tasks.

Patterns for Interaction with systems

Conversational Systems

Human machine interaction has come a long way from the days of punch cards. From spending months of training for a human to understand how to use a system, today machines are going through months of training to understand how to interact with humans, understand what humans say, what they intend and how to map the specified intent to an available business process/API call. Architecturally, this translates to a conversational layer as a channel *(with context/NLP and integration capabilities),* in addition to traditional web and mobile. While NLP capabilities are growing in leaps and bounds, aspects like multiple languages, semantic meaning, dialect variations, and lack of standardisation across components add to the complexity of the designed solution. Libraries like Rasa, Wit.ai, Microsoft's LUIS provide core language processing capabilities, which when combined with an integration/abstraction layer that maps to existing APIs/processes, can enable conversational interaction with existing systems really quickly.

Visual apps, AR/VR, Deep Learning

Machines understanding the world we humans live in - the cognitive domain of computer vision - has also been growing exponentially. Components/libraries like PyTorch for vision, OpenCV for object recognition, Tesseract for character recognition/OCR, and deep neural networks built on libraries like TensorFlow enable easy adoption and rollout of capabilities that combine/match/analyze images or video with product details, item location, stock available, etc as well as security related use cases like facial recognition and biometrics.

Delivering information back to the user becomes the domain of AR/VR which can now superimpose objects with 3D transformations on to live video, enabling systems to interact with the user seamlessly. Solutions that need to deliver these capabilities still require an architecture thats based on the principles outlined above, with the libraries mentioned in this section integrated with the rest of the ecosystem.
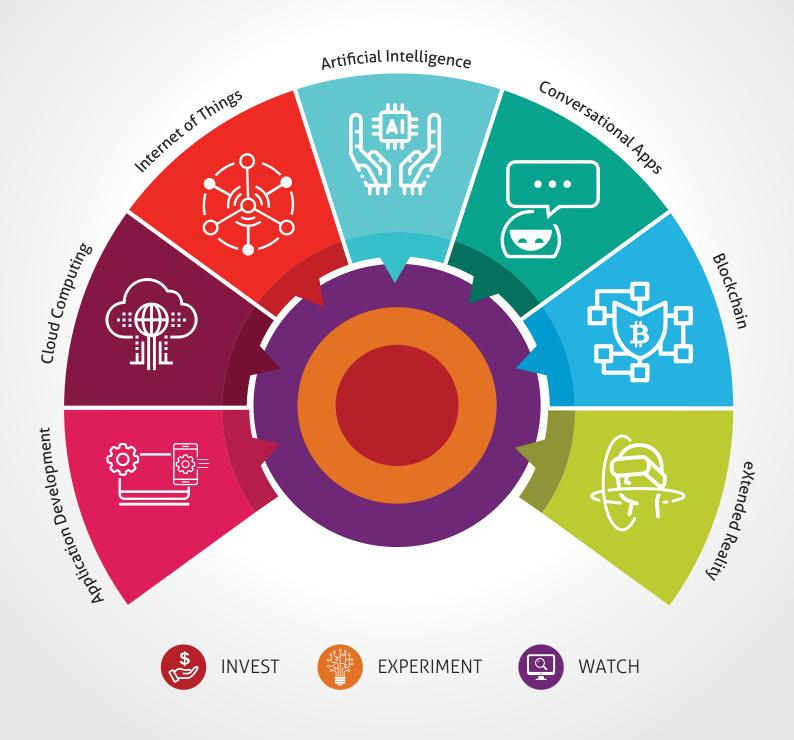
Humans interacting with Artificial Intelligence – Roles, anti-patterns?

While intelligent machines integrated with processes can provide seamless and always-on access, businesses are realizing that there are quite a few scenarios where a human is still the better judge. Components/patterns/styles listed above can enable access to a set of capabilities, however, the ultimate choice of where to draw the boundary with regard to the role of humans vs the role of machine 'intelligence', the interaction between humans and machine 'intelligence', and the level of decision making in an application system is the responsibility of us as architects of the solution. Models trained by data with little understanding/corrections will reflect the bias inherent in the data, the real risk is when humans stop paying attention to the output from these models.

While explainable AI aims to provide a reasoning behind a choice, that may not always be the right path to take – and CTOs of today, while adopting machine learning and AI in their processes should also carefully draw in boundaries of interaction and responsibility between man and machine. It looks like a solution with thresholds, human approvals past certain ranges and manual steps *(as outlined in the ethical AI principles https://ethical.institute/#principles*) may actually be a better end state than a fully automated solution where we're at the mercy of the machines.

Focus Areas

Mindtree presents a point of view on each of the following broad areas of focus in the technology domain.

Artificial Intelligence

Conversational Apps

Internet of Things

Blockchain

Cloud Computing

eXtended Reality

Application Development

INVEST   EXPERIMENT   WATCH

Application Development

The open-source market has never been more vibrant. Today, open-source technologies drive entire technology domains such as Web application and application programming interface (API) development frameworks, UI framework and database engines. Commercial software vendors are also making available a plethora of new technology products.

Cloud Computing

With a wide variety of cloud services right from compute and storage to databases and application development tools, today organizations can create applications entirely based on the products from a cloud provider. This saves the precious time of development and operations teams in installation, and management of hardware and software required to develop an application. With a clear enterprise policy to move existing on-premise applications to cloud, today it is possible for enterprises to leverage cloud platforms to save time and cost.

Internet of Things

The Internet of Things continues to emerge as a major focus across most industry sectors. With the huge market potential of IoT, the need for standardization in the near term is real. Major IT solution providers have jumped in and started providing a variety of tools and technologies.

Artificial Intelligence

Artificial intelligence enables natural and contextual interaction with tools that augment users' experiences, using the power of machine-based intelligence. Tap into an ever-growing collection of powerful artificial intelligence algorithms for vision, speech, language and knowledge.

Conversational Apps

Conversational apps provide an expressive and meaningful combination of human interactions with technology to give best customer service available in the market. It provides an opportunity for the user to communicate with machine in their natural language rather than in a syntax specific commands, along with omnipresence and help increase the productivity. Conversational apps help business to reduce spending on administrative and customer services.
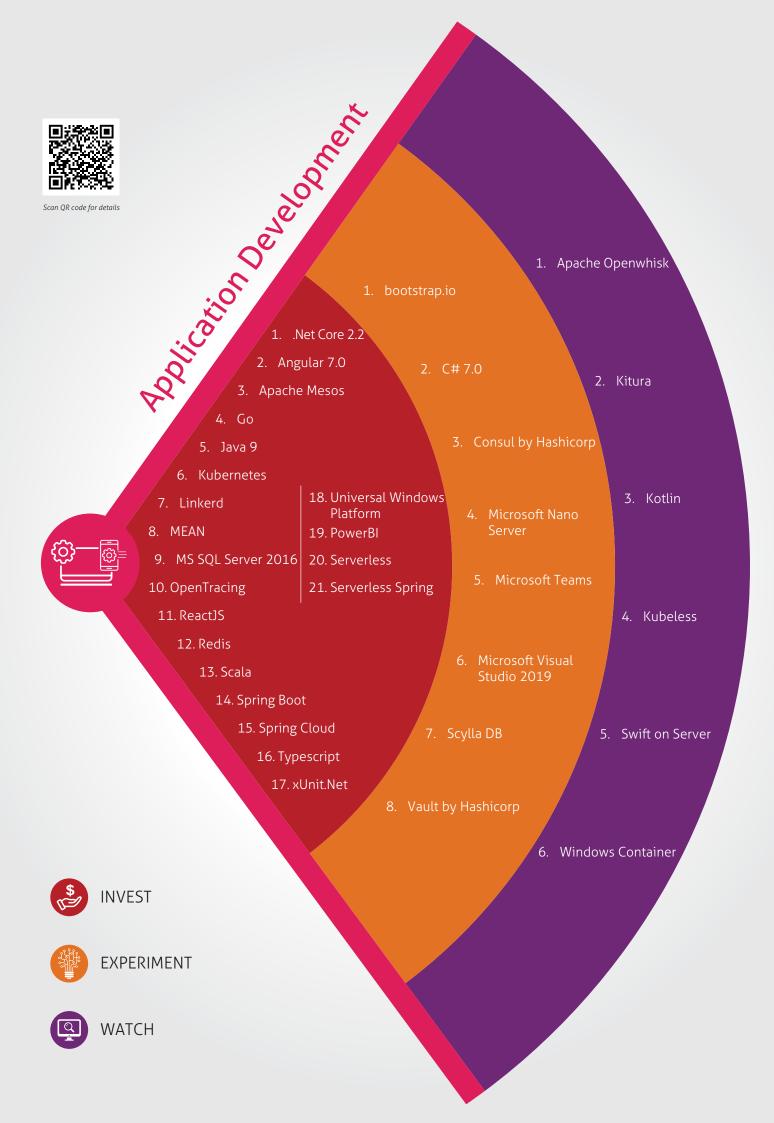
Blockchain

Blockchain is emerging as a fundamental technology that can potentially have a disruptive impact on traditional business models. It provides an incorruptible digital ledger of economic transactions that can be programmed to record not just financial transactions but virtually everything of value.

eXtended Reality

eXtended Reality is growing swiftly within the industry with huge number of wearables, devices and SDKs powering it up. eXtended Reality is a large term which covers Virtual, Augmented, Mixed/Hybrid and many more to follow. All of them have abilities to alter the way we interact with computers, which in turn opens up large number of opportunities. Though the hardware market is still evolving towards future, currently there is no stoppage in putting up mobile based realities.

Application Development

Watch (purple):

1. Apache Openwhisk
2. Kitura
3. Kotlin
4. Kubeless
5. Swift on Server
6. Windows Container

Experiment (orange):

1. bootstrap.io
2. C# 7.0
3. Consul by Hashicorp
4. Microsoft Nano Server
5. Microsoft Teams
6. Microsoft Visual Studio 2019
7. Scylla DB
8. Vault by Hashicorp

Invest (red):

1. .Net Core 2.2
2. Angular 7.0
3. Apache Mesos
4. Go
5. Java 9
6. Kubernetes
7. Linkerd
8. MEAN
9. MS SQL Server 2016
10. OpenTracing
11. ReactJS
12. Redis
13. Scala
14. Spring Boot
15. Spring Cloud
16. Typescript
17. xUnit.Net
18. Universal Windows Platform
19. PowerBI
20. Serverless
21. Serverless Spring

INVEST

EXPERIMENT

WATCH

Invest

.Net Core 2.2

.NET Core is a general purpose and cross-platform development platform for building applications that need to run on Windows, Linux and MacOS. Essential for applications that require high performance and scalability due to the high performant .Net Core Runtime. The platform facilitates mixing microservices developed with Ruby, Java or .Net Framework and deployment image to the container is of smaller size.

Angular 7.0

Angular 7.0 has many improvements over its previous versions. It supports Node 10, Typescript 3.1 and RxJS 6.3 with more support for routing, themes, gestures and animations along with virtual scrolling for very large scrollable lists and automatic rendering for drag and drop of content. Performance wise, it automatically adds or removes the metadata polyfill, which makes applications smaller in production.

Apache Mesos

Apache Mesos is a cluster manager that provides efficient resource isolation and sharing across distributed applications or frameworks. It abstracts CPU, memory, storage, and other compute resources away from machines (physical or virtual), enabling fault-tolerant and elastic distributed systems to easily be built and run effectively. Mesos is open-source software originally developed at the University of California at Berkeley.

Go

Go is an expressive, concise, clean, and efficient general purpose programming language. It compiles quickly to machine code, yet has the convenience of garbage collection and the power of runtime reflection. It is a fast, statically typed, compiled language that feels like a dynamically typed, interpreted language. It has of late gained significant prominence in server software design and construction.

Java 9

Java 9 is one of the most-awaited Java releases that will introduce several new features to an already rich programming ecosystem. The most significant of them are the introduction of jShell, Java Microbenchmarking Harness (JMH), HTTP 2.0 support, Java Platform Module System (JPMS) and several language and API improvements that should benefit the general Java developer.

Kubernetes

Kubernetes is an open-source container-orchestration platform. It helps manage containers, and scale & automate the deployment. With applications moving to cloud and using containers as a basic deployment unit, Kubernetes comes in handy to manage the deployment of containers and leverage cloud resources in an efficient manner.

Linkerd

Linkerd with microservices becoming the de facto architecture style for modern web applications, it brings certain operational challenges as the number of microservices grow in an organization. Typical challenges include scaling, dynamic routing, service discovery, failure handling, etc. Linkerd, a service-mesh framework brings in these features which can be applied on a service transparently without impacting the business logic. Linkerd can be applied to cloud as well as on-premise applications seamlessly.

MEAN

MEAN is a framework for building MongoDB, Node.js, Express and AngularJS-based applications. It is designed to provide a quick and organized way to start developing MEAN-based web apps with useful modules like Mongoose and Passport. These and other such modules come pre-bundled and configured along with this framework. It mainly tries to address the connection points between existing popular frameworks and solve common integration problems.

MS SQL Server 2016

MS SQL Server 2016 is the foundation for Microsoft's data strategy. With this new release, Microsoft delivers an end-to-end data management and business analytics solution, for demanding business applications as well as insights into business data on any device.

OpenTracing

OpenTracing is a modern web and enterprise applications are based on distributed architectures where multiple instances of micro-services serve the users from multiple data centers and cloud infrastructure across the world. To manage such complex software installation, enterprises need to have more emphasis on monitoring. OpenTracing is one such framework which helps developers trace calls across services in a distributed environment and make sense of operations to address issues quickly.

OpenTracing is a framework that is trying to bring in standardization on which various tools can be plugged in to make monitoring more effective.

ReactJS

ReactJS is an open-source JavaScript library that provides a view for data rendered as HTML. React views are typically rendered using components that contain additional components specified as custom HTML tags.

Redis

Redis is an open-source, in-memory data structure store which can be used as a database, cache or message broker. It supports data structures such as strings, hashes, lists, sets, sorted sets with range queries, bitmaps, hyperlog logs and geospatial indexes with radius queries. Redis has built-in replication, Lua scripting, LRU eviction, transactions, different levels of on-disk persistence and provides high availability via Redis Sentinel and automatic partitioning with Redis Cluster.

Scala

Scala is a popular object-oriented and functional programming language that has seen good adoption, especially for developing reactive applications. Also, Scala tends to be the programming language of choice for analytics applications that leverage the Apache Spark ecosystem.

Spring Boot

Spring Boot enables developers to build Spring applications quickly by taking a convention over configuration approach. The Spring Boot ecosystem is very vibrant and can be used to help realize a micro-services-based architecture.

Spring Cloud

Spring Cloud is a collection of tools from Pivotal that provides solutions for quickly building some of the patterns commonly encountered in distributed systems. Spring Cloud builds upon some of the common building blocks of the Spring ecosystem and the Netflix Common Runtime Services and Libraries.

TypeScript

TypeScript is a language for application-scale JavaScript. It adds optional types, classes and modules to JavaScript and supports tools for large-scale JavaScript applications for any browser, host or OS. TypeScript compiles to readable, standards-based JavaScript.

xUnit.net

xUnit.net is a free, open-source, community-focused unit testing tool for the .NET Framework. Written by the original inventor of NUnit, xUnit.net is the latest technology for unit testing C#, F#, VB.NET and other .NET languages. xUnit.net works with ReSharper, CodeRush, TestDriven.NET and Xamarin. It is also a part of the .NET Foundation, and operates under their code of conduct.

Universal Windows Platform

Windows 10 introduced the Universal Windows Platform (UWP), which enables the Windows Runtime model evolve and bring it into the Windows 10 unified core. As part of the core, the UWP now provides a common app platform that is available for every device that runs Windows 10. This package can be installed on a wide range of devices.

Power BI

Microsoft Power BI is a collection of online services and features that enable you to find and visualize data, share discoveries and collaborate in intuitive new ways.

Serverless

Serverless is your toolkit for deploying and operating server-less architectures. Functions can be written in a cloud-agnostic way and deployed on any cloud using the Serverless toolkit. This will increase productivity and allows developers to focus on logic and not worry about the cloud provider.

Serverless Spring

Serverless architecture has gained a lot of traction over the last two years. However, until now enterprises had to depend on companies like AWS for their AWS lambda service. Now with Spring framework taking the jump into server-less architecture model, we expect more organizations to experiment with this style. Given Spring framework's strong position in the enterprise application development space, Serverless Spring is a framework to look out for.

Experiment

bootstrap.io

bootstrap.io is a developing simple web applications especially single-page applications should not need a developer to write a lot of code. Frameworks like bootstrap.io are helping developers to move towards code-less development. This helps speed up development time and enable better quality code.

C# 7.0

C# 7.0 adds a number of new features and brings a focus on data consumption, code simplification and performance. A few of the newly added features are: Tuples, Pattern Matching, Local Functions, Out Parameters and Literal Improvements. Perhaps the most useful features with significant benefits would be Tuples and Pattern Matching.

Consul by Hashicorp

Consul is a versatile toolkit for discovering and configuring services within your infrastructure. It provides several key capabilities like service-discovery, health-checks and keyvalue store. Consul is designed to be friendly to both traditional application developers and the DevOps community, making it perfect for modern, elastic and highly available applications.

Microsoft Nano Server

Nano Server is a lightweight operating system, optimized for running 'cloud-native' applications based on containers and micro-services. It can also be used as an agile and cost- effective data-center host with a dramatically smaller footprint, supporting remotely managed installation and optimizations for the cloud and DevOps workflows.

Microsoft Teams

Microsoft Teams is a new chat-based workspace in the Office 365 ecosystem. It is an entirely new experience that brings together people, conversations and content along with tools that teams need to easily collaborate and achieve more. The service is a cloud-based app that integrates well with familiar Office applications and is available to paying Office 365 customers. The Office 365 component adds a group chat tool to the company's office suite, and supports both persistent and threaded chats to keep everyone engaged. It also acts as a hub for teamwork, which brings multiple good features of Office 365 together in one place. It is customizable at team level to enforce two-factor authentication, single sign-on through Active Directory and encryption of data in transit and at rest.

Microsoft Visual Studio 2019

Enhancements to code navigation, IntelliSense, refactoring, code fixes, and debugging, save time and effort on everyday tasks regardless of language or platform. For teams embracing DevOps, Visual Studio 2019 streamlines the inner loop and speeds up code flow with brand new real-time features such as live unit testing and real-time architectural dependency validation.

Scylla DB

Scylla DB takes a new approach to implement the NoSQL DB. Unlike typical NoSQL DBs, Scylla uses multiple engines to run the queries in a 'share-nothing' approach. This means each CPU core can have Scylla engine with its own CPU and memory. This helps scale up and scale out to increase by many folds compared to traditional Nosql DBs.

Vault by Hashicorp

Vault is a tool for securely accessing secrets. A secret is anything that you want to tightly control access to, such as API keys, passwords, certificates, and more. Vault provides a unified interface to any secret, while providing tight access control and recording a detailed audit log.

Watch

Apache Openwhisk

Apache Openwhisk for enterprises with a strong affinity for developing applications using open-source frameworks, Apache OpenWhisk is a good option to develop server-less apps on the cloud. Apache OpenWhisk offers many features, but is still in an incubating stage. But, it is worth to keep an eye on as there are many server-less frameworks becoming available in the market.

Kitura

Kitura is a new, modular, package-based web framework written in the Swift language. It is a high-performance and simple-to-use web framework for building modern Swift applications.

Kotlin

Kotlin is a functional programming language, statically typed and is 100% interoperable with Java and Android. Since Android and Java platforms are widely used, Kotlin is gaining a lot of traction and is worth to keep a watch on.

Kubeless

Kubeless is a kubernetes native server-less framework to build Function as a Service applications. With kubernetes having established itself as standard for container orchestration for web-scale applications, kubeless is the framework to look out for, as server-less applications are gaining traction.

Swift on Server

Swift is a high-performance system programming language. It has a clean and modern syntax, offers seamless access to existing C and Objective-C code and frameworks, and is memory-safe by default. Although inspired by Objective-C and many other languages, Swift is not itself a C-derived language. As a complete and independent language, Swift packages core features like flow control, data structures, and functions, with high-level constructs like objects, protocols, closures and generics. Swift embraces modules, eliminating the need for headers and the code duplication they entail.

Windows Container

Windows Container is an isolated place where an application can run without affecting the rest of the system, and vice versa. Containers are the next evolution in virtualization. If you were inside a container, it would look very much like you were inside a newly installed physical computer or a virtual machine. And, to Docker, a Windows Container can be managed in the same way as any other container.

Windows Container Types:

- *Windows Server Containers:* provide application isolation through process and namespace isolation technology. A Windows Server container shares a kernel with the container host and all containers running on the host.

- *Hyper-V Containers:* expand on the isolation provided by Windows Server Containers by running each container in a highly optimized virtual machine. In this configuration the kernel of the container host is not shared with other Hyper-V Containers.

Cloud Computing

1. Compute
2. Data Pipelines
3. Database
4. Cloud Datawarehouse
5. API Gateway
6. Migration Services

1. Compute - Multi-Cloud Migration
2. Data Pipelines
3. Azure Data Lake Gen 2 & AWS Lake Formation
4. Databases

1. AWS Firecracker
2. AWS Outpost

INVEST

EXPERIMENT

WATCH

Invest

Compute

In Cloud Compute ecosystems, serverless cloud services, containers and container orchestration services are witnessing increasing adoption.

Container orchestration engines handle the lifecycle, scaling, auto-healing including failover, updates, monitoring of container-based systems. Leading services of investment are AWS ECS for Kubernetes, AWS Fargate, Azure Kubernetes Service, Google Kubernetes Engine.

Serverless computing platforms drive event driven programming Leading services of investment are AWS Lambda, AWS Step Functions, Azure Functions and Google Cloud Functions.

Data Pipelines

Data Pipelines is almost every new Big Data engagement starts on Cloud owing to the separation of compute and storage coupled with price per insight commercial model. Data Pipelines are key achieve consistency in the data ingestion process.

Key areas of investment include Step Functions coupled with Lambda functions for serverless data pipelines with low volume and AWS Glue for more continuous workloads. These pair well with AWS Kinesis.

From an Azure perspective Azure Data Factory Pipelines continue to be an area of investment. Azure Data Factory is a fully managed ETL service that can be used to compose and orchestrate data integration workflows spanning on-premises and cloud. These pair well Azure Event Hubs for continuous data.

From a Google Cloud perspective Cloud Dataproc, Dataflow and Pub/Sub are areas of investment.

Database

Managed database services continue to increase. All relational and NoSQL database services across cloud providers continue to see investment. Multi-master globally redundant databases like AWS Aurora, AWS DynamoDB, Azure Cosmos DB, Google Spanner see more investment.

Cloud Datawarehouse

Managed Datawarehouse services which make use of object storage and scalable capacity continue to see high adoption. Key areas of investment are AWS Redshift, Azure SQL Datawarehouse, Snowflake.

API Gateway

Pace of cloud adoption has increased. Key customer requests are to migrate Cloud Native migration services and third AWS API Gateway makes it possible to develop, publish APIs for all types of applications in a secure and scalable manner. API Gateway is a core component for serverless functions and microservices deployed in the cloud.

AWS API Gateway and Azure API management continue to be areas of investment.

Migration Services

Getting to cloud is a repeatable and predictable manner is a key request from customers. Towards this automatic discovery, workload analysis, programmatic migrations including databases is key to success. In addition to third party services cloud provider specific services are very important.

Azure Server Migration service and AWS Database migration service are investment priorities. AWS Database Migration Service helps migrate databases to AWS quickly and securely. The source database remains fully operational during the migration, minimizing downtime to applications that rely on the database.

Experiment

Compute - Multi-Cloud Migration

Modern architectures with Multi-Cloud adoption is a key expectation from customers. This will evolve from cloud provider specific orchestration engines like EKS, AKS. Towards this Google Anthos holds a lot of promise. Anthos lets one build and manage modern hybrid applications across environments. This is powered by Kubernetes and other industry-leading open-source technologies from Google.

From a serverless programming stand-point AWS Lambda Layers promises to provide a simple way to manage common software and data across multiple functions.

Data Pipelines

Data Pipelines are evolving from notebook-based workflows towards graphical workflows. Data Flow is a new feature of Azure Data Factory (ADF) that allows development through graphical data transformation logic that can be executed as activities within ADF pipelines.

Azure Data Lake Gen 2 & AWS Lake Formation

Azure Data Lake Gen 2 is a set of capabilities dedicated to big data analytics which combines blob storage and Azure Data Lake Storage. This brings in parity to AWS S3 as a data lake storage.

Similarly, AWS Lake Formation is a service that makes it easy to setup data lakes.

Databases

New types of data processing is becoming key to adoption. Towards this databases like AWS Timestream and AWS Quantum Ledger database (QLDB) will be areas of experimentation. QLDB is a fully managed, purpose-built ledger database that provides an immutable and cryptographically verifiable history of all changes made to the application's data while Timestream is a managed time series database out. AWS Firecracker is an open-source virtualization technology that is purpose-built for creating and managing secure, multitenant containers and functions-based services.
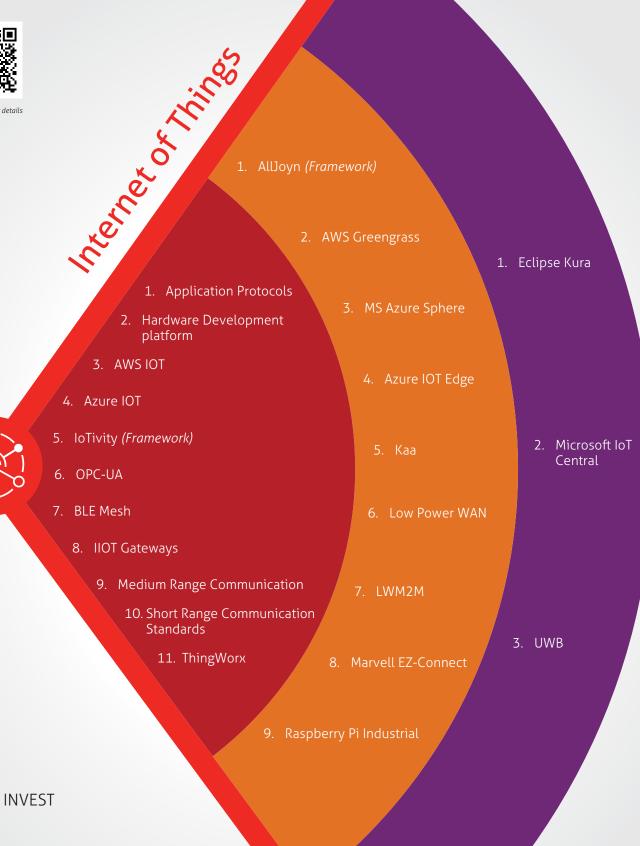
Watch

AWS Firecracker

With container technologies taking root and Docker leading the way, space for new container technologies needs to be watched out. AWS Firecracker is an open-source virtualization technology that is purpose-built for creating and managing secure, multitenant containers and functions-based services.

AWS Outpost

Hybrid cloud workloads is a key requirement from multiple enterprises. AWS Outposts bring native AWS services, infrastructure, and operating models to virtually any data center, co-location space, or on-premises facility.

Mindtree

Internet of Things

Red section:
1. Application Protocols
2. Hardware Development platform
3. AWS IOT
4. Azure IOT
5. IoTivity *(Framework)*
6. OPC-UA
7. BLE Mesh
8. IIOT Gateways
9. Medium Range Communication
10. Short Range Communication Standards
11. ThingWorx

Orange section:
1. AllJoyn *(Framework)*
2. AWS Greengrass
3. MS Azure Sphere
4. Azure IOT Edge
5. Kaa
6. Low Power WAN
7. LWM2M
8. Marvell EZ-Connect
9. Raspberry Pi Industrial

Purple section:
1. Eclipse Kura
2. Microsoft IoT Central
3. UWB

INVEST

EXPERIMENT

WATCH

Invest

Application Protocols

- **MQTT:** MQTT is a lightweight machine-to-machine communication protocol that uses a publish-subscribe messaging transport. Its small footprint is ideal for IoT devices.

- **AMQP:** Advanced Message Queuing Protocol (AMQP) is an open standard message-oriented middleware. It boasts of a low latency and high throughput queuing mechanism that is ideal for IoT workloads.

- **HTTP:** While HTTP is a rather heavy protocol for constrained devices, it continues to be used for IoT. REST over HTTP is a common technique for communication.

- **CoAP:** CoAP is a data transfer protocol designed for needs of constrained devices. It runs over UDP not on TCP and follows the client server model. CoAP is designed to interoperate with HTTP and the RESTful web at large through simple proxies.

- **6LoWPAN:** 6LoWPAN is a low-power wireless mesh network where every node has its own IPv6 address, allowing it to connect directly to the Internet using open standards.

Hardware Development Platform

- **Arduino:** Arduino is an open standards physical computing platform based on a simple microcontroller board and a development environment for writing software for the board. It is used in a variety of IoT use cases, such as taking inputs from sensors and switches, developing interactive objects and controlling a variety of subsystems. Arduino sketches enable quick development for prototype use cases. Where a Pi is suitable for prototyping gateways, Arduino is suitable for prototyping edge nodes.

- **ARM mBed:** ARM mBed accelerates product development with a complete stack that includes an open standards embedded operating system, client and transport layer security and an array of tools to develop and deploy software for ARM Cortex-M-based IoT devices. The overall ecosystem makes this an attractive choice.

- **ESP:** ESP8266 is a low-power, highly integrated Wi-Fi solution embedded with 8Mbit flash. ESP8266EX is among the most integrated Wi-Fi chips in the industry. Measuring just 5 mm x 5 mm, ESP8266EX requires minimal external circuitry and integrates a 32-bit Tensilica MCU, standard digital peripheral interfaces, antenna switches, RF balun, power amplifier, low noise receive amplifier, filters and power management modules – all in one small package. ESP32 is a 2.4 GHz Wi-Fi and Bluetooth combo chip TSMC low power 40 nm technology. Initially targeted for prototyping, now we have seen widespread adoption of this module in IOT applications.

- **Raspberry Pi:** Raspberry Pi is a low-cost, credit-card sized computer that plugs into a computer monitor or TV. It uses a standard keyboard and mouse, and is capable of doing everything you would expect a desktop computer to do. It interfaces with USB, Ethernet, Bluetooth and Wi-Fi. The presence of SPI and GPIO banks allows the Pi to also be used to connect over serial, and to a variety of other sensors and actuators. We believe that the price point and flexibility makes it a good prototyping platform

AWS IoT

The AWS IoT platform provides secure, bi-directional communication between Internet-connected things like the sensors, actuators, embedded devices or smart appliances, and the Amazon Web Services cloud. You can use it to collect, store, and analyze telemetry data from multiple devices, and create applications that enable users to control these devices from their phones or tablets. AWS' openness, richness of platform capability, and its robustness make it a good IoT platform. Integration with other upcoming AWS technologies such as Greengrass for edge, and Machine Learning, makes it a good platform to build full IoT solutions.

Azure IoT

Microsoft Azure IoT Hub allows us to connect, monitor and manage millions of IoT assets from the Azure Cloud. With the addition of features like native MQTT v3.1.1 protocol support, X.509 certificate-based device authentication and device management, Azure IoT Hub continues to grow into a stable

IoT platform. Combined with greater support for the open-source stack, and big-data capabilities around Spark and Hadoop, Azure becomes a good choice for building IoT solutions in general, and for building on the Microsoft stack in particular.

IoTivity (Framework)

IoTivity is an open-source software framework enabling seamless device-to-device connectivity to address the emerging needs of the Internet of Things. Backed by Samsung and Intel, it is designed to work on Wi-Fi Direct, BT/BTLE, ZigBee and ZWave. As AllJoyn and IoTivity are merged, current devices running either AllJoyn or IoTivity will be interoperable and backward compatible.

OPC-UA

OPC UA is a standard that ensures the open connectivity, interoperability, security, and reliability of industrial automation devices and systems. OPC UA is widely recognized as the key communication and data modeling technology for the Industry 4.0 initiative. It works with many software platforms, is completely scalable, and has been widely adopted across several industries. OPC UA is widely used connectivity in Industrial IOT use cases.

BLE Mesh

BLE Mesh is a key enabling technology for IoT, since it removes the single hop topology restrictions and enables communication from every device to every other device within the mesh network. BLE mesh is based on existing, low-power BLE wireless technology so the devices can run on low power.

Power amplifier, low noise receive amplifier, filters and power management modules – all in one small package. ESP32 is a 2.4 GHz Wi-Fi and Bluetooth combo chip TSMC low power 40 nm technology. Initially targeted for prototyping, we have seen widespread adoption of BLE mesh across the IOT applications.

Intel-based IoT gateways: Dell, Advantech

In physical settings, gateways need to have both good computing power, and also be built robustly, to withstand conditions such as heat and dust. Dell and Advantech have a series of edge gateway products based on the Intel stack, which will be necessary to provide intelligence on the edge.

These are based on Intel Quark and Atom processors. The equivalent prototype boards from Intel are Edison (based on Atom), and Galileo (based on Quark).

Medium Range Communication: Sub GHz

Sub-GHz technology is an ideal choice for wireless applications requiring long-range, low power consumption and low data rate. This can transmit data to distant hubs, often a few miles away, without hopping from node to node. This long-range transmission capability reduces the need for multiple expensive base stations or repeaters. This can be an alternative to LoRA, without requiring a tie-in into that ecosystem.

Short Range Communication protocols / Standards

Short-range communication protocols are typically used in communication between sensor nodes and gateways. Based on scenario, different protocols and standards need to be used:

- **RFID/NFC:** RFID and NFC technologies continue to see traction in areas like retail. Tags are being used extensively for low-cost identification purposes.

- **ZigBee:** ZigBee is a low-cost, low-power, wireless mesh network standard targeted at the wide development of long battery life devices in wireless control and monitoring applications.

- **BLE:** Bluetooth low energy (BLE), now marketed as Bluetooth Smart, is a wireless personal area network technology designed and marketed by the Bluetooth Special Interest Group. It targets the consumer space with novel applications for the healthcare, fitness, beacons, security and home entertainment industries. Bluetooth 5 now has mesh support which will help better serve IoT use cases.

ThingWorx

ThingWorx is a platform from PTC, well known in the world of machine-to-machine and IoT applications for the quick and intuitive development of IoT dashboards. ThingWorx accelerates IoT application development by compressing the design-develop-deploy cycles to reduce time to market and spur innovation.

Experiment

AllJoyn (Framework)

AllJoyn is an open source software framework that makes it easy for devices and apps to discover and communicate with each other. Developers can write applications for interoperability regardless of transport layer, manufacturer, and without the need for Internet access. The wide industry backing, and its expansion beyond home automation makes it a good candidate to consider for standardization.

AWS Greengrass

AWS Greengrass is an edge software component that allows local compute, data filtering and secure communications between devices and the cloud. Developers can leverage the freedom to use familiar programming languages to easily execute program logic via AWS Lambda functions on the edge. While this service is available in limited preview, we believe the service has adequate merit to be included in our 'Experiment' category.

MS Azure Sphere

Azure Sphere is a Linux based operating system created by Microsoft to provide the end to end security solution to the MCUs. Microsoft Azure Sphere solution is a three-part solution, the first part is Azure Sphere certified MCUs, the second part is Azure Sphere Operating System running on MCUs and the third part is Azure Sphere security service which ensures a secure communication between the device and cloud for over-the-air update, authentication and error reporting capabilities

Azure IOT Edge

Azure IOT Edge is a fully managed service that is designed for the use cases where you need to analyze the data on the device (edge). So IOT edge deliver the cloud intelligence locally by deploying and running artificial intelligence (AI), Azure services, and custom logic directly on cross-platform IoT devices.

Kaa

Kaa is an open-source IoT platform that promises to cater to a wide range of business domains. We found its ability to integrate with client-provided custom code as one of its strong points. It also addresses core communication concerns and lets developers focus on the business logic around their IoT use cases. Its scalable server architecture is another strong point. Strong community support and noteworthy interest across industries makes this IoT platform an interesting area of experimentation.

Low Power WAN (LoRA, Sigfox, NB-IoT)

LPWAN technologies are used to connect low-cost, low-power and low-bandwidth devices. Competing technologies continue to emerge in the LPWAN segment which uses Sub GHz frequency range, with the focus on maximizing the range and minimizing the transmission power. Many companies are building a business model around this. We believe this is an important component of long-range communication, where solutions require non-cellular connectivity. Smart cities are another candidate for LPWAN solutions. LoRA and Sigfox are more established compared to NB-IoT. LoRA has seen greater adoption, and is a candidate that we believe will move quickly from 'Experiment' to 'Invest'.

LWM2M (Specification)

LWM2M is a new open industry protocol developed by OMA to provide a means of remotely performing service enablement and application management for embedded devices and connected appliances. It is a communication protocol for use between client software on an M2M device and server software on a M2M management and service platform. The backing by OMA, and its rich specification will drive its adoption to a large extent.

Marvell EZ-Connect

The Marvell EZ-Connect platform is helping to rapidly bring innovative devices to the fast-growing IoT market. Built on Marvell's silicon and coupled with industry-standard certifications and leading Kinoma software, EZ-Connect is widely considered the most complete of industry-leading connectivity portfolios. It's hardware and software solutions are capable of enabling more efficient time-to-market IoT designs.

Raspberry Pi Industrial

The Compute Module is a Raspberry Pi in a more flexible form factor, intended for industrial application. CM3 is the Raspberry Pi 3-based revision of the original Compute Module. This comes with BCM2837 processor, 1 GB RAM and a 4 GB eMMC Flash device, all integrated on to a small 67.6 mm x 31 mm board which fits into a standard DDR2 SODIMM connector. This is a good candidate to use when we are planning to productize our own hardware platform which is prototyped using Raspberry Pi3. This compares favorably with platforms like Advantech from a cost standpoint.

Watch

Eclipse Kura

Eclipse Kura is an Eclipse IoT project that provides a platform for building IoT gateways. It is a smart application container that enables remote management of such gateways and provides a wide range of APIs for allowing you to write and deploy your own IoT application. Kura runs on top of the Java Virtual Machine (JVM) and leverages OSGi, a dynamic component system for Java, to simplify the process of writing reusable software building blocks. Kura APIs offer easy access to the underlying hardware including serial ports, GPS, watchdog, USB, GPIOs, I2C, etc. It also offers OSGi bundle to simplify the management of network configurations, the communication with IoT servers, and remote management of the gateway.

Microsoft IoT Central

IoT Central is a fully managed IoT SaaS solution which enables users to connect, monitor and manage their IoT assets at scale.

UWB

Specifically for indoor navigation, high bandwidth and extremely short-pulse waveforms help in reducing the effect of multi-path interference and facilitate determination of ToA (Time of Arrival) for burst transmission between transmitter-receiver pairs. This makes UWB a good solution for indoor positioning. UWB signals are able to effectively measure distance between two devices up to 10-cm of accuracy, compared to roughly 5-m accuracy for Wi-Fi and Bluetooth. This needs to be looked at in relation to adoption (UWB by itself has not seen significant adoption), and power consumption, and so we place it in the 'Watch' category.

Artificial Intelligence

INVEST

1. Azure ML Studio
2. Google Cognitive Services
3. MLLib Spark
4. Python
5. Scikit-Learn
6. XGBoost
7. Alexa
8. RASA NLU
9. Custom Speech Service

EXPERIMENT

1. PyTorch
2. Tensorflow 2.0
3. Caffe2
4. Keras-rl
5. PyTorch Geometric
6. TensorFlowOnSpark
7. Universe-Gym
8. Amazon Polly
9. Amazon Comprehend
10. Amazon Lex
11. Amazon Personalize
12. Amazon Rekognition
13. Amazon Sagemaker
14. Amazon Textract
15. Amazon Transcribe
16. Amazon Translate

WATCH

1. PyTorch-BigGraph
2. DGL
3. Pyro

Invest

Azure ML Studio

Azure ML allows users to create and train models, then turn them into APIs that can be consumed by other services.

Google Cognitive Services

Google Cloud Vision API has powerful machine-learning models to understand the content of an image. It exposes an easy-to-use REST API. It quickly classifies images into thousands of categories, detects individual objects and faces within images, and finds and reads printed words contained within images.

- **Google Cloud Speech API:** Google Cloud Speech API enables developers to convert audio to text by applying powerful neural network models in an easy to use API.

  The API recognizes over 80 languages and variants. It can help transcribe the text of users dictating to an application's microphone, enable command-and-control through voice, or transcribe audio files, among many other use cases.

- **Google Cloud Natural Language API:** Google Cloud Natural Language API reveals the structure and meaning of text by offering powerful machine-learning models in an easy-to-use API. It can be used to extract information about people, places, events and much more, mentioned in text documents, news articles or blog posts. It can be used to understand sentiment about products on social media or parse intent from customer conversations happening in a call center or a messaging app.

MLlib Spark

The goal of Apache Spark's machine learning library is to make practical machine learning scalable and easy. It consists of common learning algorithms and utilities, including classification, regression, clustering, collaborative filtering, dimensionality reduction, as well as lower-level optimization primitives and higher-level pipeline APIs.

Python

Python is an interpreted, high level programming language which is very convenient for developers for rapid application development. With no compilation step the edit-test-debug cycle is incredibly fast leading to increased developer productivity. It is the preferred language for machine learning researchers and engineers. It has a rich set of libraries for numerical and statistical operations. All major Machine learning and Deep learning frameworks have excellent python support

Scikit-Learn

Scikit-Learn is a well-designed, open-source tool for data analysis, data mining and machine learning. It is completely Python-based and leverages Python packages like Numpy and Scipy. The library has a consistent programming interface and is easy to get started.

XGBoost

XGBoost is a machine library using gradient-boosted decision trees designed for speed and performance. It supports regression, classification, ranking and user-defined objectives. It is scalable, supports parallel and distributed execution and provides interfaces to multiple programming languages. XGBoost is very effective for prediction tasks on tabular or structured data-sets. XGBoost library is used in many data science and machine learning challenges.

Alexa

Alexa is Amazon's voice control system with a built-in natural-language processing system. Alexa is built in the cloud, so it is always getting smarter. The more customers use Alexa, the more she adapts to speech patterns, vocabulary and personal preferences.

RASA NLU

An open source NLP tool for intent classification and entity extraction. It helps in building own language parsers using existing NLP and ML libraries.

Custom Speech Service

Custom Speech Service enables you to create customized language models and acoustic models tailored to your application and your users. By uploading your specific speech and/or text data to Custom Speech Service, you can create custom models that can be used in conjunction with Microsoft's existing state-of-the-art speech models.

Experiment

PyTorch

PyTorch is a Python-based tensor computing framework with GPU acceleration. PyTorch is a fast, minimalist network which is easy to extend with custom neural network modules written in Python. Compared to other deep learning frameworks, PyTorch is easier to debug and understanding of error messages and stack traces is relatively easier. PyTorch supports building dynamic neural networks which are built on tape-based autograd system.

TensorFlow 2.0

Tensorflow is one of the most widely adopted library for deep learning. Version 2 focuses on easy model building using Keras. Eager execution has been incorporated in Tensorflow 2.0, which is an imperative programming environment that evaluates operations immediately (without building graphs). Operations return concrete values instead of constructing a computational graph to run later. This makes it easier to understand and get started. Eager execution provides an intuitive interface because it has natural pythonic control flow Instead of graph control flow and this makes it easier to understand and make debugging easier as well. It introduced tf.function() decorator, which has a cool new feature of AutoGraph, which lets you write graph code using natural Python syntax. tf.function() allows you to transform a subset of Python syntax into portable, high-performance TensorFlow graphs. TensorFlow provides different software stacks like Tensorflow Lite, Tensorflow.js which makes tensorflow deployment more convenient. Using which the models can not only be deployed in servers but on mobile and browsers as well.

Caffe2

Caffe is a deeplearning framework implemented in C++ that has a useful Python interface. Caffe2 is a high performance and modular framework which can be deployed to servers and mobile devices. It focuses on scalable and cross-platform support. It is developer-friendly for deployment as well as in terms of performance. Caffe2 team has closely worked with Amazon Web Services, Intel Corp., Microsoft Corp., Qualcomm Inc. and Nvidia Corp. to optimize Caffe2 for the cloud and

mobile devices. For example, Caffe2 harnesses the power of Adreno graphics processing units and Hexagon digital signal processors on Qualcomm Inc.'s Snapdragon chips. Caffe2 features built-in distributed training using the NCCL multi-GPU communications library. This means that you can very quickly scale up or down without refactoring your design. Caffe2 delivers near-linear scaling of deep learning training achieving up to 7.7x speed up with 8 GPUs, compared to a single GPU training. Caffe2 was designed to support the deep learning AI workflow from end-to-end. Its modular, lightweight framework can take full advantage of large scale GPU servers as well as run efficiently on lower powered embedded edge devices.

Keras-rl

Keras-rl implements state-of-the-art deep reinforcement learning algorithms in Python and seamlessly integrates with the deep learning library Keras. Keras-rl makes it really easy to run state-of-the-art deep reinforcement learning algorithms, uses Keras and thus supports Theano or TensorFlow back-end. Hence algorithms can efficiently be trained either on CPU or GPU. It works with OpenAI Gym out of the box as well and makes evaluating and playing around with different algorithms relatively easy. Keras-rl is extendable which makes it easy to implement your own environments and even algorithms by simply extending some simple abstract classes.
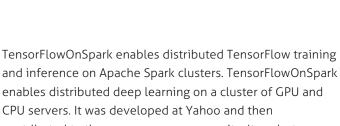
PyTorch Geometric

Most existing deep learning techniques deal with Euclidean data. However many real world data like graphs is non-Euclidean and irregular in structure. Geometric Deep Learning deals with the extension of Deep Learning techniques to graph/manifold structured data.

PyTorch Geometric consists of various methods for deep learning on graphs and other irregular structures. These methods include graph convolution, graph embeddings and graph attention based networks.

TensorFlowOnSpark

TensorFlowOnSpark enables merging of deep learning with big data platforms.

TensorFlowOnSpark enables distributed TensorFlow training and inference on Apache Spark clusters. TensorFlowOnSpark enables distributed deep learning on a cluster of GPU and CPU servers. It was developed at Yahoo and then contributed to the open-source community. It seeks to minimize the amount of code changes required to run existing TensorFlow programs on a shared grid. Its Spark-compatible API helps manage the TensorFlow cluster.

Universe-Gym

Universe allows an AI agent to use a computer like a human does, by looking at screen pixels and operating a virtual keyboard and mouse. Universe lets you train a single agent on any task a human can complete with a computer. Universe exposes a wide range of environments through a common interface: the agent operates a remote desktop by observing pixels of a screen and producing keyboard and mouse commands. Existing game environments and real-world browser tasks can be controlled from Universe-Gym environment. Gym is a toolkit for developing and comparing reinforcement learning algorithms. With Universe, any program can be run within a Gym environment.

Amazon Polly

Amazon Polly is a service that turns text into life like speech. It allows to create applications that talk and build entirely new categories of speech-enabled products. It uses advanced deep learning technologies to synthesize speech that sounds like a human voice.

Amazon Comprehend

Amazon Comprehend uses machine learning to help uncover the insights and relationships in your unstructured data. It identifies the language of the text; extracts key phrases, places, people, brands, or events. It understands how positive or negative the text is, it can analyze text using tokenization and parts of speech; and automatically organizes a collection of text files by topic.

Amazon Lex

Amazon Lex is a service for building conversational interfaces into any application using voice and text. It helps in recognizing the intent of the text, to enable to build applications with highly engaging user experiences and lifelike conversational interactions.

The same deep learning technologies that power Amazon Alexa are available to any developer, enabling them to quickly and easily build sophisticated, natural language, conversational bots.

Amazon Personalize

Amazon Personalize is used to improve customer engagement by powering personalized product and content recommendations, tailored search results, and targeted marketing promotions.

It provides an activity stream from application – page views, signups, purchases, and so forth – as well as an inventory of the items user want to recommend, such as articles, products, videos, or music.

Amazon Rekognition

The Rekognition API enables integration of image and video analysis to the existing applications. When the image or video is provided to the Rekognition API, the service identifies the objects, people, text, scenes, and activities, as well as detect any inappropriate content.

It also provides highly accurate facial analysis and facial recognition on images and video provided to it.

Amazon Sagemaker

Sagemaker provides every developer and data scientist with the ability to build, train, and deploy machine learning models quickly.

It is a fully-managed service that covers the entire machine learning workflow to label and prepare your data, choose an algorithm, train the algorithm, tune and optimize it for deployment, make predictions, and take action.

Amazon Textract

The API automatically extracts text and data from scanned documents. It goes beyond simple optical character recognition (OCR) wherein it identifies the contents of fields in forms and information stored in tables and gives a meaningful output in various formats.

It can be used to quickly automate document workflows, enabling the user to process millions of document pages.

Amazon Transcribe

Amazon Transcribe is an automatic speech recognition (ASR) service that makes it easy for developers to add speech-to-text capability to their applications. It can analyze audio files stored in Amazon S3 and have the service return a text file of the transcribed speech.

Transcribe also works on live audio stream and gives a stream of transcripts in real time.

Amazon Translate

Amazon Translate is a neural machine translation service that delivers fast, high-quality, and affordable language translation. It uses deep learning models to deliver more accurate and more natural sounding translation than traditional statistical and rule-based translation algorithms.
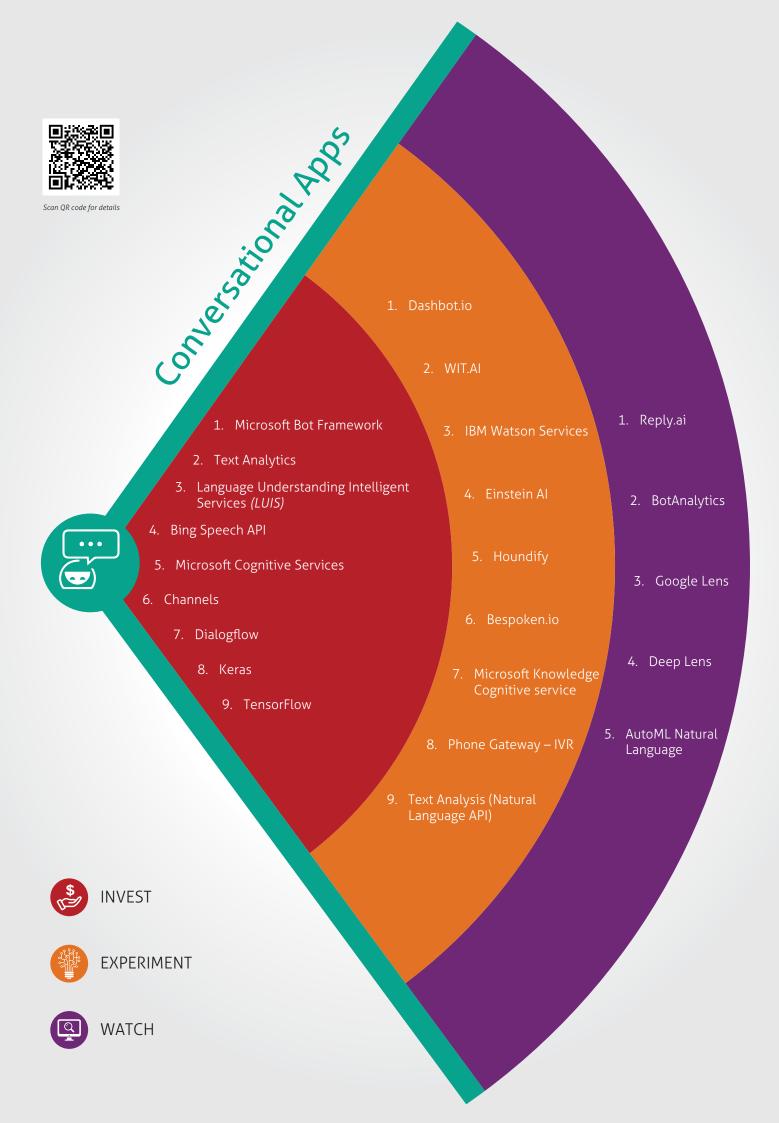
Watch

PyTorch-BigGraph

PyTorch-BigGraph is a distributed system for learning graph embeddings for large graphs, particularly big web interaction graphs with up to billions of entities and trillions of edges.
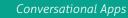
DGL

Deep Graph Library (DGL) is an implementation of graph neural network model family, on top of existing DL frameworks (e.g. PyTorch, MXNet, Gluon etc.). DGL is designed to be compatible and agnostic to the existing tensor frameworks.

Pyro

Pyro is a universal probabilistic programming language (PPL) written in Python and supported by PyTorch on the back end. Pyro enables flexible and expressive deep probabilistic modeling, unifying the best of modern deep learning and Bayesian modeling. With support of variational inference techniques, Pyro aims to address the key challenge of non-scalability of inference of probabilistic systems.

Scan QR code for details

Conversational Apps

Microsoft Bot Framework (red)
1. Microsoft Bot Framework
2. Text Analytics
3. Language Understanding Intelligent Services *(LUIS)*
4. Bing Speech API
5. Microsoft Cognitive Services
6. Channels
7. Dialogflow
8. Keras
9. TensorFlow

(orange)
1. Dashbot.io
2. WIT.AI
3. IBM Watson Services
4. Einstein AI
5. Houndify
6. Bespoken.io
7. Microsoft Knowledge Cognitive service
8. Phone Gateway – IVR
9. Text Analysis (Natural Language API)

(purple)
1. Reply.ai
2. BotAnalytics
3. Google Lens
4. Deep Lens
5. AutoML Natural Language

INVEST

EXPERIMENT

WATCH

Invest

Microsoft Bot Framework

The Microsoft Bot Framework provides just what you need to build and connect intelligent bots that interact naturally wherever your users are talking, from text/sms to Skype, Slack, Office 365 mail and other popular services. The framework provides developers with a developer portal and SDK to build Bots, a Bot Connector service to connect to social channels such as Twitter and Slack, and a Bot Directory to discover and use existing bots.

Text Analytics

Text Analytics are a set of Azure cloud based APIs that are useful for semantic analysis (analysis of raw text for positive and negative sentiment usually for customer survey), key phrase extraction, language detection in text for further translation and also Named Entity Recognition for identity terms like people, place, organization, dates, quantities, percentages and currencies.

Language Understanding Intelligent Services (LUIS)

LUIS is designed to enable developers to build smart applications that can understand human language and accordingly react to user requests. With LUIS, a developer can quickly deploy an HTTP endpoint that will take the sentences sent to it and interpret them in terms of their intents (the intentions they convey) and entities (key information relevant to the intent). LUIS supports several languages such as English, French, Italian, German, Spanish, Brazilian, Portuguese, Japanese, Korean and Chinese. Many of the common chat bots can be easily built using LUIS.

Bing Speech API

Bing Speech is a cloud-based API that provides advanced algorithms to process spoken language. With Bing Speech API, you can add speech-driven actions to your apps, including real-time interaction with the user. Bing Speech APIs and libraries enable speech capabilities on all internet-connected devices. Every major platform including Android, iOS, Windows, and 3rd-party IoT device are supported. It offers industry-leading speech-to-text, text-to-speech, and language understanding capabilities delivered through the cloud. Microsoft uses Bing Speech API for Windows applications like Cortana and Skype Translator as well as Android applications like Bing Torque for Android Wear and Android Phone.

Microsoft Cognitive Services

Microsoft Cognitive Services provide a suite of powerful machine learning APIs for computer vision, speech, language understanding and text analysis. These powerful services are exposed as easy to use REST interfaces. Computer vision APIs include image classification, recognizing celebrities, face and gender detection, detection of text in images and even classification of frames in real-time video. Microsoft offers services for content moderation including image, text and video. Microsoft speech services have very powerful speech services where the language model can be fine-tuned and customized by tailoring it to the vocabulary of the application and the speaking style of your users. It supports language detection, translation services and even custom translation system.

Channels

In the conversational world, Channels are the primary source to interact with the users. There are several channels which are categorized as Social, Enterprise and Custom. The following are in focus

- Telegram
- Twilio
- Twitter
- Whatsapp
- KiK
- Skype
- Teams
- Slack
- Facebook
- Line

Dialogflow

Dialogflow is a platform provided by Google to build engaging voice and text-based conversational interfaces powered by AI and also connect with social channels along with popular platforms and devices.

Keras

Keras is a high-level neural network API, written in Python and capable of running on top of either TensorFlow or Theano. It is open-source and has a modular design which can be easily extended with custom modules. It has a high-level interface that enables rapid prototyping of complex neural network architectures. Keras supports convolutional networks, recurrent networks and a combination of the two.

TensorFlow

TensorFlow is a popular, general, flexible, open-source framework for numerical computation using data flow graphs that can be deployed on CPUs and GPUs. It can be used to build various types of neural network architectures for deep learning. It is written in C++ and has Python bindings. The flexible architecture allows the deployment of computation to one or more CPUs or GPUs in a desktop, server, or mobile device with a single API.

TensorFlow supports multi-GPU training and the servers can be grouped into a cluster with the ability to distribute a computation task across the cluster. It integrates with a suite of visualization tools called TensorBoard which allows for better understanding, optimizing and quantitative metric plotting of data being trained by neural networks.

Experiment

Dashbot.io

Dashbot.io is a bot that can be integrated with any chat bot to monitor and generate the conversational metrics.

WIT.AI

Wit.ai makes it easy for developers to build applications and devices that you can talk or text to. It empowers developers with an open and extensible natural language platform. Wit.ai learns human language from every interaction, and leverages the community – what is learned is shared across developers. Wit.ai is free for personal and commercial usage.

IBM Watson Services

Watson and its cognitive capabilities mirror some of the key cognitive elements of human expertise. Watson provides a variety of smart services that help in building compelling cognitive applications. Watson has its strengths in natural language understanding including understanding Context around the language.

Watson provides services for conversation, dialog, document conversion, language translation, natural language classification, tone analysis, etc. Watson has Retrieve & Rank services which for example, can help an experienced technician to quickly find solutions from dense product manuals. Watson AlchemyLanguage provides a set of APIs that offers text analysis through natural language processing. These services are exposed as easy to use REST APIs.

Einstein AI

Einstein AI is a product from Salesforce that is powered by Predictive analytics, natural language processing, machine learning and smart data discovery.

Houndify

Houndify is a platform to integrate voice and conversational intelligence for customization of products

Bespoken.io

Bespoken.io is an automated testing and monitoring tool for Alexa and Google Actions.

Microsoft Knowledge Cognitive service

Maps complex information and data in order to solve tasks such as intelligent recommendations and semantic search.

Phone Gateway – IVR

Provided by Google, it is used to build conversational IVR using telephony gateway.

Text Analysis (Natural Language API)

A google cloud based Natural Language API, which supports sentiment analysis, entity analysis, entity sentiment analysis, syntactic analysis and content classification.

Watch

Reply.ai

Reply.ai is a platform to build and manage the conversational strategy focused on customer services and have human agents take over.

BotAnalytics

BotAnalytics is a conversational analytics tool focused on analyzing engagement and retention measurement for chatbots.

Google Lens

Google Lens Translates text in real time, look up words, add events to your calendar, call a number, and more to use in conversational apps.
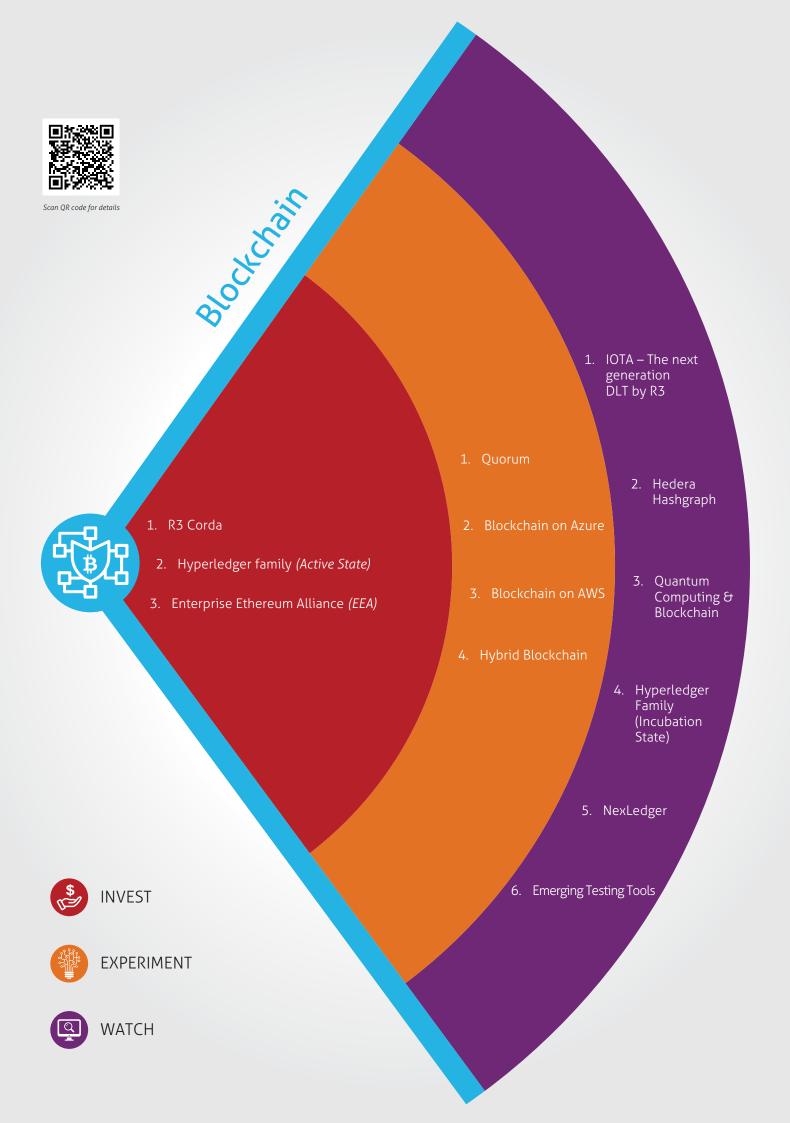
Deep Lens

AWS DeepLens is easy to customize and is fully programmable using AWS Lambda. The deep learning models in DeepLens even run as part of an AWS Lambda function, providing a familiar programming environment to experiment with and integrate with conversational ai.

AutoML Natural Language

Google provides AutoML, a cloud based service that can extract, classify and detect sentiment from the uploaded and labeled text document for a domain specific keywords and phrases to get the insights.

Blockchain

Scan QR code for details

1. R3 Corda

2. Hyperledger family *(Active State)*

3. Enterprise Ethereum Alliance *(EEA)*

1. Quorum

2. Blockchain on Azure

3. Blockchain on AWS

4. Hybrid Blockchain

1. IOTA – The next generation DLT by R3

2. Hedera Hashgraph

3. Quantum Computing & Blockchain

4. Hyperledger Family (Incubation State)

5. NexLedger

6. Emerging Testing Tools

INVEST

EXPERIMENT

WATCH

Invest

R3 Corda

Corda is a framework specifically used by the Banking and Financial institutions to explore various use cases. Corda also has a unique consensus protocol called Notary which is suitable for a lot of financial transactions. Corda is recommended for any financial institution looking to implement a Blockchain use case.

Hyperledger family (Active State)

Hyperledger has production ready products like Composer, Fabric, Iroha & Sawtooth. 'Hyperledger Composer' is used for building business networks easily. 'Hyperledger Fabric' is a platform for distributed ledger solutions, underpinned by a modular architecture delivering high degrees of confidentiality, resiliency, flexibility and scalability. 'Hyperledger Iroha' in C++, with an emphasis on supporting mobile applications with essential functionality for asset, information or identity management. 'Hyperledger Sawtooth' is the distributed Ledger supporting Multi-Language support.

Enterprise Ethereum Alliance (EEA)

EEA is a member-driven standard organizations provides a decentralized ecosystem which is open-source, public, blockchain-based distributed computing platform and operating system featuring smart contract (scripting) functionality. It supports a modified version of Nakamoto consensus via transaction-based state transitions.

Experiment

Quorum

Quorum is an enterprise-ready distributed ledger and smart contract platform similar to Ethereum, founded by JP Morgan Chase. This platform is ideal for any application requiring high speed and high throughput processing of private transactions within a permissioned group of known participants. Quorum addresses specific challenges to blockchain technology adoption within the financial industry, and beyond.

Blockchain on Azure

With Blockchain technology gaining traction across industries, many enterprises want to explore use cases. Microsoft Azure gives a jump-start with pre-built templates for Ethereum, Hyperledger Fabric, Corda and Chain platforms. For enterprises starting new, instead of setting up infrastructure, this is a very good option to look for.

Blockchain on AWS

AWS provides a fully managed, scalable blockchain service. AWS' blockchain service makes it easy to setup, deploy, and manage scalable blockchain networks, eliminating the need for you to rely on expensive consulting implementations.

Hybrid Blockchain

Hybrid Blockchain, by design, is a combination of different characteristics of public and private blockchains. It allows to determine what information stays private and what information is made public. Thus, any system can reap the benefits of both the approaches.

Watch

IOTA – The next generation DLT by R3

This is the first open-source permissionless distributed ledger that is being built to power the future of Internet of Things with feeless micro-transactions and data integrity for machines. It is based on a new distributed ledger technology, the Tangle, which overcomes the inefficiencies of current Blockchain designs and introduces a new way of reaching consensus in a decentralized peer-to-peer system.

Hedera Hashgraph

Hashgraph is a distributed ledger technology envisaged on asynchronous Byzantine Fault Tolerance (aBFT) consensus algorithm amalgamated with the concept of gossip protocol, capable of securing the platform against attacks. Hedera Hashgraph is a public distributed ledger based on the Hashgraph algorithm which uses DAG for time-sequencing.

Quantum Computing & Blockchain

In case of Quantum computing, data analyzation is done using quantum-mechanical components. It will be a threat to Blockchain as quantum computing will calculate the complex mathematical equations which will compromise Blockchain public key cyptography, thus making Blockchain as highly vulnerable technology.

Hyperledger Family (Incubation State)

Hyperledger family has some of its project in incubation state and yet to go live. These projects has some unique features which will lift blockchain technology to another level altogether. For eg, 'Hyperledger Burrow' is a permissioned ethereum smart contract blockchain. 'Hyperledger Explorer' will provide an WebUI facility. 'Hyperledger Quilt' is for interoperability solution for blockchain, DLTs and other type of ledgers.

NexLedger

Nexledger is a blockchain platform developed by Samsung SDS. It is flexible and scalable with the ability to deliver wide spectrum of customizable use case application services based on enterprise specific requirements.

Emerging Testing Tools

- **Ethereum Tester:** It is an open source testing library. The setup is pretty easy with a manageable API support for various Testing requirements.

- **BitcoinJ:** It is a Java-based framework built for Bitcoin-based apps that enables you to interact with the real BTC network and various testing activities. It is an open network available for assistance.

- **Truffle:** It is a commonly referred tool for Ethereum developers, which brings in good testing features, such as automated contract testing. The framework holds capabilities beyond just testing functionality within the blockchain application.

- **Ganache (previously TestRPC):** This tool is solely built to test Ethereum contracts locally. It creates a simulated blockchain that allows anyone to use various accounts for testing.

- **Hyperledger Composer:** Unlike the other tools mentioned in this article, Hyperledger Composer lacks most of blockchain testing features. Therefore, it's only useful for testing apps before they are launched. The three main areas it excels in are interactive, automated system and automated unit testing.

eXtended Reality

INVEST

1. Unity
2. ARKit
3. ARCore
4. Blender
5. Autodesk Products

EXPERIMENT

1. Unreal Engine
2. OpenCV
3. Vuforia
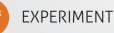4. Wikitude
5. Light variants of ML

WATCH

1. WebXR
2. 6D.ai
3. 8th Wall
4. Service Offerings
5. Amazon Sumerian

Invest

Unity

Unity is a cross-platform game engine developed by Unity Technologies which is used for creating 2D/3D applications, video games, animations for websites, gaming consoles, mobile devices, etc. Unity is notable for its ability to target games to multiple platforms. Within a project, developers have control over delivery to mobile devices, web browsers, desktops, and consoles. This tool is also used for creating content for eXtended Reality devices – Google VR, Microsoft HoloLens, HTC Vive, Oculus Rift, smartphones – Android and iOS etc.

ARKit

Augmented reality (AR) describes user experiences that add 2D or 3D elements to the live view from a device's camera in a way that makes those elements appear to inhabit the real world. ARKit combines device motion tracking, camera scene capture, advanced scene processing, and display conveniences to simplify the task of building an AR experience. It integrates iOS device camera and motion features to produce augmented reality experiences in the application.

ARCore

ARCore is a platform by Google to build augmented reality applications to integrate virtual content with the use of Motion Tracking, Light Estimation and Estimation understanding techniques. ARCore enables your smartphone to understand the world and interact with information. Some of the features of ARCore are open for both Android and iOS to create shared AR experiences.

Blender

Blender is a free, open source 3D creation suite for powering up the eXtended Reality applications. This application can handle full 3D pipeline – modeling, rigging, animation, simulation, rendering, compositing and motion tracking etc. It even supports python scripting for basic animation and rendering.

Autodesk Products

Autodesk has been a leader in the area of 3D content creation for years. Some of the popular products are Maya, AutoCAD, 3DS Max, Revit; which are used by many industrial designers, architects, real estate designers, game artists, many animation studios and more. Due to vast availability of the pre-curated content and artists, this is one of the popular choice to enable 3D content on eXtended Reality applications.

Experiment

Unreal Engine

Unreal Engine is a complete suite of creation tools which are designed to meet ambitious artistic visions, by Epic Games. Bringing photorealistic renders and content is a jewel for its crown. This can be used to develop eXtended Reality applications for HTC Vive, mobiles, PC etc.

OpenCV

OpenCV is a library of functions targeting to resolve real-time computer vision problems. It is supported in many platforms such as Windows, Linux, MacOS, smartphones – iOS and Android. Being a BSD-licensed product, OpenCV makes it easy for businesses to utilize and modify the code. It has over 2500 optimized algorithms which can be used to detect and recognize faces, identify objects, classify human actions in videos, track camera movements, track moving

Mindtree

objects, extract 3D models of objects, produce 3D point clouds from stereo cameras, stitch images together to produce a high resolution image of an entire scene, find similar images from an image database, remove red eyes from images taken using flash, follow eye movements, recognize scenery and establish markers to overlay it with augmented reality, etc.

Vuforia

Vuforia is an Augmented Reality SDK for devices which enables the creation of Augmented Reality applications. It uses Computer Vision technology to recognize and track Image Targets and simple 3D objects, such as boxes, in real time. This image registration capability enables developers to position and orient virtual objects, such as 3D models and other media, in relation to real-world images when these are viewed through the camera of a mobile device / HoloLens. The virtual object then tracks the position and orientation of the image in real time so that the viewer's perspective of the object corresponds with their perspective on the Image Target, so that it appears that the virtual object is a part of the real-world scene.

Wikitude

Wikitude is a mobile augmented reality based SDK, which is supported on iOS, Android , UWP. Just like Vuforia, this can also recognize and render based on markers/images. This also has a support for Geo based AR to power the applications.

Light variants of ML

Just having 3D models with basic functions aren't enough for providing great experiences. So we use AI to power up the apps, to behave naturally. Few of the frameworks involve TensorFlow, TensorFlow Lite, CoreML, ML Toolkit etc. They provide the necessary gesture/object/face/body detection and also perform multiple analysis right on the user's palm.

Watch

WebXR

We are entering a phase where JavaScript can pull out an eXtended Reality experience right on the website, where the user doesn't need to install any special apps. The computer/smartphone screens will be the portal for the end users to enter the world of eXtended Reality. Though the set of libraries in this space are not mature enough, but this will be the next platform to experience with. Currently supported on limited browsers and devices.

6D.ai

Though it's just a startup, it has made everyone to turn up towards them by their product which does faster SLAM. 6D Reality Platform AR with Depth, Occlusion and Physics, and all by using only a smartphone camera. This is currently in beta and could be a game changer once the cat is out of bag.

8th Wall

8th Wall is breaking the barriers between the digital and physical by bridging the development gap between iOS and Android, and even on web. Their platform enables the app to work seamlessly regardless of the platform the user is on.

Service offerings

There are many services provided by Microsoft Azure, Amazon Web Services and Google Cloud services which are very useful for eXtended Reality applications. These mainly constitute Cognitive services which are Vision based and Speech, remote rendering, spatial anchors etc., to enable the applications to perform the heavy work off-load to cloud, also to enable complex features to the devices which aren't capable enough.

Amazon Sumerian

Amazon Sumerian is another development platform for eXtended Reality applications, without even requiring specialized programming or 3D graphics expertise. Applications built on Sumerian can target to Oculus Rift, HTC Vive, Google Daydream, Android and iOS mobile devices and more. This mainly works as a WebVR implementation rather a native application.

This report and its content is based on the current available information in public domain about the referred technologies and our own experiences of working with them. Mindtree holds a neutral position on fitment of a particular technology/technologies for client–specific business problems and recommends evaluation on a case-by-case basis.

About Mindtree

Mindtree is a global IT consulting and services company which helps clients across 17 countries achieve business agility, competitive edge, and growth. We harness the power of Continuous Delivery, our digital expertise, industry knowledge, and research in emerging technologies to drive efficiencies and enable business innovation for over 340 clients. Mindtree is consistently regarded as one of the best places to work. This is a reflection of our entrepreneurial, collaborative and dedicated "Mindtree Minds" who embody the winning culture that defines our commitment to excellence, innovation, and co-creation. To learn more about us, visit www.mindtree.com or follow us @Mindtree_Ltd