



PERSISTENT

Predictions for 2016: Eight Game-Changing Digital Technologies for Enterprise Transformation

Office of the CTO



Contents

| | |
|--|----|
| Introduction | 2 |
| Trend 1: Mobile User Experience | 3 |
| Trend 2: Internet of Things (IoT)..... | 4 |
| Trend 3: Adaptive Security and Data Governance | 5 |
| Trend 4: Data Lakes..... | 6 |
| Trend 5: Interactive Analytics at Scale..... | 7 |
| Trend 6: Machine Intelligence | 8 |
| Trend 7: Blockchain..... | 9 |
| Trend 8: Containers..... | 10 |
| Conclusion | 11 |
| Contributors | 11 |

Introduction

2016 will see a significant increase in the volume, velocity, and variety of enterprises creating and executing on their digital strategies. This change will only serve to accelerate the shift of business spending from IT Modernization to Digital Transformation. At Persistent Systems, we see the following eight digital technology trends as game changers for software-driven enterprises in 2016 and beyond. The trends fall into three major categories.

- The first two trends (**Mobile User Experience** and **Internet of Things**) pertain to the interactions among the enterprise, people, and devices.
- The next three trends (**Adaptive Security and Data Governance**, **Data Lakes**, and **Interactive Analytics at Scale**) relate to the lifeblood of today's software-driven business – data, along with the analytics performed on this data to obtain insights and drive business actions.
- The final three trends (**Machine Intelligence**, **Blockchain**, and **Containers**) describe enabling technologies that are broadly applicable to various aspects of digital transformation across different industries.

The trends are of course interrelated: thus, machine intelligence techniques are needed for the adaptive security solutions that drive secure IoT platforms, while IoT platforms are needed to harvest the data that we rely on for insights. We have indicated these linkages wherever appropriate.

As a CEO, CIO, or Line of Business Head, you are likely most interested in understanding two things: the nature and degree of impact that each technology will have on your business; and how you can best leverage the trend to successfully execute your digital strategy. We have chosen to keep the description of the trends brief (no more than a page each), focusing on two things: the most significant developments that will emerge or strengthen in the coming year, and ways for businesses to leverage them. We also provide a set of 3-5 business questions with each technology trend to help you get started down the path of discovery.

Trend 1: Mobile User Experience

User experience (often abbreviated as UX) is the centerpiece of digital transformation. You must “wow” your users almost immediately to engage them, or you risk losing them. In 2016, we will see a significant focus on user experience as a business strategy, not only to grow top-line revenue but also to create stickier, long-term user engagement. It is also clear that much of the user experience innovation in 2016 will be mobile-centric.

One of the biggest changes we predict in 2016 is that *the app interface itself will fade into the background to make the experience more contextual, intuitive, and action-oriented*. Siri and Google Now are the most well-known exemplars of this trend, and in 2016 more enterprise apps will emulate them – popping open just at the right moment with the right information, while otherwise remaining in the background. Reducing the friction for the user, from having to actually open the app to use it, is no small advantage. *Learning the habits and preferences of the user to customize the experience and app triggers will be the new normal*. Both iOS and Android notification centers have made significant changes in recent years, opening the opportunity for a richer app interaction within the notification itself.

Simple and intuitive interfaces will dominate in 2016, with designers opting to only keep the frequently-used actions directly available on an app screen while hiding everything else. This also means the polarizing “hamburger menu” which hides all the other options is not going away anytime soon.

Users will also expect a seamless experience of the service on multiple devices and touch points as they use a single service across smartphones, desktops, tablets, and wearables. *Maintaining consistent state, interface, and capabilities across multiple devices will be the new normal*.

We will also witness *a significant increase in the offerings and maturity of cross-platform development and testing platforms* to move cross-platform apps closer to native apps in delivering a smooth and responsive experience. Cross-device testing, usability testing, and cloud-hosted in-app analytics will make application experience consistent across the growing spaghetti of devices and platform versions.

Inter-app communication, both data and control, will mature in 2016 to make it easier for apps to communicate securely. Users won't face jarring app switches when carrying out a foreign app action from within their current app, but will instead be able to complete multi-step transactions of the foreign app from within the currently open app.

As a decision maker in your enterprise, ask yourself these questions in 2016 about the experience that your customers, vendors, and partners are expecting from your enterprise systems.

- 1. Am I wowing them with an intuitive, streamlined, and frictionless user experience?**
- 2. Are my services available to them whenever and wherever they need them, even if they are not using my application?**
- 3. How can I help them through my services just when they need help?**
- 4. How can I provide them a uniform, seamless, and consistent experience across channels?**

Trend 2: Internet of Things (IoT)

Predictions for the number of smart devices (or “things”) that will be online in 2016 are in the billions, with continuing growth forecast for the foreseeable future. The scale and diversity of IoT endpoints has been growing rapidly over the last couple of years, ranging from smart wearables, to environmental beacons and sensors, to smart gadgets and drones, to appliances, buildings, and cars that interact with their users. The greatest growth has been in the consumer segment, with wearables and home automation dominating. Outside the consumer segment, the growth has been strongest in the industrial segment and assembly lines. *In 2016, we will see a strong growth of IoT devices in commerce, transportation, infrastructure, and enterprise.* From shopping malls to roads to office buildings, these devices will power the ambient user experience helping deliver the just-in-time contextual action, insight, or information to the user. We will also see growth in deployment of systems that not only sense the world around them but react autonomously. Examples include entire buildings that optimize energy consumption and monitor air quality, self-driving vehicles, and traffic light systems that react not just individually based on proximal sensors but collectively based on a network of those to regulate traffic. In the consumer space, we will see *more interoperability within devices and platforms* that will allow the lightbulb from one manufacturer to respond to inputs from the thermostat/motion sensor of another via a platform built by a third.

IoT security received some attention in 2015, with the odd appliance hack demo and the more famous smart Jeep disablement incidents. In 2016, however, IoT will play in the hacking majors, with at least one major security incident that will be real (not a controlled experiment) and with severe consequences. This may be the impetus needed for the industry to finally think security by design in this space. *From end-to-end encryption to secure local storage to over-the-air patches, we will see security baked into the IoT pipeline.*

While smart device and sensor manufacturers will benefit from the IoT boom, the true business value of IoT lies in the analysis of the data generated and gathered by these devices and the business actions that can be taken. It is an open opportunity for both manufacturers and platform players to seize. Intelligent systems – ones that learn from user behavior, environmental data gathered from their sensors, and inputs received from other devices around them – will be the biggest winners.

Use the following questions to determine how your enterprise should focus on IoT in 2016.

- 1. How can I leverage IoT to create a low friction proactive engagement with my customer?**
- 2. Can I reduce errors and speed up my business processes by deploying sensors and devices that capture the data currently gathered manually?**
- 3. How can I transform data collection from my assembly line or manufacturing processes to achieve a more continuous, faster, and comprehensive data pipeline complete with real-time insights?**
- 4. What are the opportunities for automation in my operations and supply chain where I can leverage IoT to reduce cost?**

Trend 3: Adaptive Security and Data Governance

From data breaches to ransomware to state-sponsored hacking, security threats have taken a whole new serious dimension in recent years. Additionally, the network perimeter keeps expanding due to digital strategies that involve Bring Your Own Device (BYOD) and the opening up of controlled enterprise access to vendors, partners, and customers. This puts security squarely as one of the most important strategic pillars of a successful digital plan.

Fixed enterprise security policies have come up short against the multi-pronged attack vectors such as zero-day exploits and Advanced Persistent Threats (APTs). As digital strategies open the enterprise further and drive up customer data collection, protection against these threats require that the security policies are enforced using a dynamic risk based model rather than with static rules. Enterprise security tools of 2016 will ingest data and signals from endpoint security agents, application logs, external threat feeds, access control systems, and Data Loss Prevention (DLP) platforms to continually score users and their activity on the network, and to adopt the security posture needed to suit the complete context rather than isolated events. For example, an authorized user transferring large amount of data at odd hours may raise a flag, even if the access is authorized. These systems will also correlate data across input sources to detect suspicious activity. The emphasis will be to raise fewer but highly accurate alerts that require immediate attention and action. Such systems will leverage machine learning algorithms to adapt the security posture appropriate for the networks being monitored.

Data breaches rocked our world in 2014 and 2015, starting with the shopping majors in the United States to the federal government's own comprehensive background check database of its most sensitive information on its employees. Stringent data privacy laws led by Europe as well as *a heightened sense of awareness and concern for their private data in the minds of the consumers* is the second driving force why data governance will be a major issue in 2016. Given the dynamic network boundary, a one-size-fits-all approach does not work. Enterprises will adopt tools that let them prioritize and designate the governance of their data and resources at the most granular level. Compliance tools that regulate cross-border data exchange will be an integral part of the data-driven digital ecosystem. In domains such as healthcare, genomics research, and personalized medicine, the need for *the ability to share data while maintaining strict confidentiality standards* will be paramount. In the IoT segment, the computationally challenged end-devices means deeper data analysis has to happen in the cloud; given the highly personal nature of IoT data, this will require granular governance and data retention policies.

Ask the following questions to determine how your enterprise fares on the security dimension.

- 1. How ready is my network to handle new threats that have unknown patterns, can be executed by insiders, and are designed to be stealthy?**
- 2. Do I have a Data Breach Prevention and Handling strategy in place along with the right set of tools?**
- 3. How compliant is my enterprise with the data collection, retention, and security regulations of all the countries in which it operates?**

Trend 4: Data Lakes

A data lake, as originally described by James Dixon, is a large storage repository of raw data, in contrast to a data mart which is a smaller repository to store only subset of attributes and aggregated data. Over the years, the term data lake has come to describe a platform that can ingest multi-structured data from various internal and external sources and make it available for applications to consume however they want. Data lakes offer organizations not only a cost-effective way to store a lot of data, but also a way to optimize existing businesses and address new business opportunities, such as cross-channel analysis of customer interactions and behaviors to offer personalized services; storage and analysis of various kinds of patients' data (physician's notes, radiology images, prescriptions, reports, etc.) to predict the likelihood of re-admission; and storage and analysis of Web data (server logs, clickstream data, etc.) to improve advertisement selection and placement.

Most data lake deployments today are based on Apache HDFS (Hadoop Distributed File System), which is a low-cost solution to store big data. Although HDFS is core to data lakes, organizations need tools around data quality, data governance, metadata discovery, etc. to make this data ready for businesses to consume. Without the tools to govern and manage your data lake, it effectively turn into a data graveyard. While some organizations have successful data lake implementations, the technology is still not mature. There are still limited options available for data ingestion frameworks, data quality, metadata discovery, etc. This is a key reason why the adoption of data lakes in enterprises has been slow.

Organizations like Cloudera, Hortonworks (Apache Hadoop 2.0, Apache Falcon, Apache Atlas), Waterline Data (Automated Data Inventory), StreamSets (Data Collector), LinkedIn (Gobblin – Data Ingestion framework) and many others are working on building and enhancing products to simplify and streamline data lake implementations. *In 2016, we will see the emergence of more out-of-the-box technology options for data lakes, along with best practices and guidelines that make it easier for enterprises to implement data lakes, both on-premise as well cloud-based.*

Answer the following questions to help you decide whether your organization can benefit from investing in a data lake architecture.

- 1. Do we have access to large amounts of multi-structured data coming in, which we are not storing or processing today? Are we generating large amounts of un/multi-structured, multi-format data that we are throwing away after on-the-fly decisions?**
- 2. Are our current analytics and reports based only on limited or aggregated data, and could we do better with the underlying raw data?**
- 3. What data silos do we currently have, i.e., data fragmented across the organization with no effective way to analyze all the data together?**

Trend 5: Interactive Analytics at Scale

“Life is a banquet and most poor suckers are starving to death”, said Auntie Mame in the eponymous musical. That can be applied to a lot of organizations today when it comes to data – too much can be, well, too much to take in. Today, it is common for enterprises to try and analyze terabytes to petabytes of multi-structured data, where text, audio, and visual data are analyzed side-by-side with structured tabular data. At the same time, users are making increasing demands on the data visualization component of data analysis: visualization and analysis at multiple resolutions, as well as a higher bar for aesthetics and ease of use for data visualization and intuitive ways of interacting with them. While big data technologies allow analyzing large data sets, they have also brought the additional complexity of multiple technologies and different programming paradigms. Additionally, integration challenges remain at every interface in the data analysis pipeline. These challenges have made the task of data analysis increasingly difficult and time consuming.

In 2016, emerging data platforms will tackle this challenge by abstracting the complexity of this pipeline for enterprise IT and business analysts and offer a “single stack” interface that allows the analysts to simply focus on the business problem and define their own analytics in a declarative form which the system then translates to a pipeline of processing from loading to cleanup to analysis and visualization. These platforms will be collaborative by design and not as an afterthought.

Enterprise business analysts and decision makers will also expect to be able to interactively analyze this data and be able to drill down into details or explore alternative hypotheses. Our diminishing ability to anticipate the query load a priori (coupled with schema-on-use) in a big data ecosystem has led to a renewed interest in *online aggregation techniques*. Online aggregation attempts to handle ad hoc queries quickly without pre-computation – by incrementally refining approximate answers. This approach requires a fundamental change in the end-user mindset, where complete accuracy in large-scale analytic queries is unimportant, and that it is more important to be able to balance accuracy and running time in a flexible way. Techniques such as on-the-fly sampling, pre-computed samples on base tables, query resolution reduction to match consumer display constraints, as well as a synthesis of these approaches to build end-to-end systems – where all of these techniques can be leveraged in tandem – will increasingly become more mainstream in 2016.

To understand if this trend is something your enterprise should capitalize on in 2016, ask yourself:

- 1. Could my Business Analysts derive new insights if they had the ability to interact with data via intuitive interfaces, or would they be overwhelmed by the data deluge?**
- 2. What are my potential cost and time savings and ROI on a system that lets my team customize, modify, and personalize their views and data explorations in minutes?**
- 3. What critical business decisions would improve if we had the ability to explore large-scale raw or historical data on the fly rather than just relying on pre-canned dashboards?**

Trend 6: Machine Intelligence

"I'm sorry Dave, I'm afraid I can't do that." That line, spoken by HAL 9000 in *2001: A Space Odyssey*, offered a 1960s glimpse of what we imagined machine intelligence would be like in the future, the future in this case being 15 years ago! But while we may not have flying cars or robots that control us, machine intelligence is in fact here. From optimizing an Uber driver's route (and maximizing revenue) to recognizing faces in photographs, software is taking on the responsibility for complex processes that require "learning" of new information from multiple data sources and past outcomes, often needing to perform this learning in real-time.

Until now, advanced machine intelligence and deep learning techniques were largely limited to the realm of scientists with advanced degrees and enterprises with vast resources. With the availability of platforms such as Watson (IBM), Azure ML (Microsoft), TensorFlow (Google), and various open-source options, these techniques are now becoming available to every enterprise. Of course, enterprises will still need to model their problems, come up with the right training scenarios for the algorithms, and train their workforce in using these tools. However, harnessing the power of advanced learning algorithms at scale has become incredibly accessible with these platforms. *2016 will see major adoption of machine learning algorithms and platforms in driving operational efficiencies, predicting failures, customizing user experience and in discovering new revenue channels.*

2015 was a breakthrough year in terms of machine intelligence advances both in terms of democratizing it through the aforementioned cloud services and strides made in vertical solutions towards machine autonomy. In 2016 we expect to see a much stronger impetus towards unsupervised and reinforcement learning techniques to build on that.

Watch out in many more places for automated customer service chatbots that can understand human questions and provide the right answer. Autonomous vehicles are one of the biggest success stories of this technology domain, and we will see early adoption in the industrial applications for these in 2016. One of the biggest impacts deep learning technologies will have is in the cybersecurity domain. It is the next big hope to be able to detect previously unseen malware or stealthy and persistent threats by connecting the dots on seemingly unconnected network or system events and user behaviors. Machine intelligence will also power major advances in genomics, personalized medicine, and drug discovery.

Ask the following questions to determine if your enterprise can benefit from machine intelligence.

- 1. Which parts of my operations and customer engagements processes can I improve by leveraging machine learning?**
- 2. Do I have sufficient and high-quality data coming into the enterprise that allows me to employ algorithms for analyzing the past and predicting the best move forward?**
- 3. How do I leverage machine intelligence for cost savings through increased automation in operations?**

Trend 7: Blockchain

Blockchain, the underlying technology that powers Bitcoin, has been independently recognized for its broader potential for some time now. Blockchain is a distributed, transparent, and auditable ledger that does not depend on any central trusted entity but instead establishes trust, immutability, and integrity through cryptographically powered distributed consensus models.

Blockchain offers a distributed trust and/or consensus model that eliminates dependencies on central entities. Multiple players have adopted and extended the blockchain platform for uses beyond Bitcoin. Key constraints that plagued Bitcoin – such as anonymity, proof of work (mining) overheads, and permission-less participation – are eliminated by some of these platforms. These emerging platforms are geared for more nuanced use cases, including: support for a closed, permissioned ecosystem; trusted special nodes; and alternative consensus models that do not require mining as a proof of work.

The second half of 2015 saw some breakout platforms and alliances emerge to take this forward in a wide range of domains including finance, IoT, and supply chain. Examples include Ripple (international payments), Counterparty (financial contracts), and Ethereum (smart contracts). We will see a sharp rise in companies building domain-specific solutions using blockchain in 2016, and some of the early players will cement their position in the various domains. In late 2015, the Linux Foundation partnered with a large group of technology (IBM, Intel, etc.) and financial services (J. P. Morgan, Wells Fargo, etc.) players to build an enterprise-grade blockchain that will be open source. This is a very positive development towards establishing a large-scale developer community as well as enterprise adoption in the future. We will see early results of that in 2016.

In 2016 we expect to see the emerging platforms mature and sign on early adopters. Processes that requires central trusted third parties for trust and clearance and that want to eliminate cross border delays in processing or multiple intermediaries will be redesigned using blockchain technology. We expect to see significant disruptive innovation, especially but not limited to the Financial, Government, and Commerce domains.

To determine if blockchain can help your enterprise, consider the following questions. Keep in mind that this technology trend is at an earlier point in its lifecycle than the other seven.

- 1. Do we have multi-party settlement processes that are slowed down by the intermediary trusted third party mediating the process?**
- 2. Is our supply chain process too cumbersome due to long drawn transactional verifications at intermediate points?**
- 3. Are we or our customers conducting payment processes across more than one country and would like to optimize both the cost and time they take?**

Trend 8: Containers

In today's cloud infrastructures, a "container" is a packaged runtime environment for an application along with its dependencies and configuration files. Containers enable us to move applications from one environment to another without breaking, and in turn allow greater agility between the development and operations teams. The past two years have seen an incredible growth in the container ecosystem, from core technology innovation to rich suites of tools for orchestration, monitoring, and management. Such advances will be a key driver for widespread enterprise adoption of containers in 2016.

While there has been a great interest in containers among enterprises, and many tests of the technology, their actual adoption in production environments has been quite low. Security concerns have been the biggest reason for this reluctance. Several recent developments have addressed those issues, leading us to predict a much higher adoption rate for containers in 2016.

Since multiple containers run atop the same OS (either directly or inside a VM) at user level, the isolation between applications is not sufficient for creating a secure multi-tenant environment. Support for user namespaces in Docker and other containers has alleviated this issue substantially. Trustworthiness of images has been another concern, which should go away with systems like CoreOS introducing signed images that can be verified before execution.

Another major concern is that IT security has no visibility inside the container itself. New startups like TwistLock are addressing that problem and also enabling the enterprise IT security team to extend their security and access control postures to the containers and provide the much needed audit trails.

Another reason creating a drag on operations is that persistent data storage has to be external to the container, which in turns implies that it is disconnected during migration. New tools like Flocker have come up to address this and enable stateful services inside containers.

Enterprises need to start evaluating and testing containers now, as adopting this technology will require a significant change in the skills, operations, and workflows that are historically slow to change due to enterprise complexity. Containers and the ecosystem around them should be considered as a catalyst for change in the enterprise, to improve processes, reduce infrastructure complexity, and speed up innovation.

To decide whether containers are an appropriate technology choice for you, ask yourself these questions.

- 1. Are my engineering and ops teams managing a large set of applications distributed across multiple data centers and environments?**
- 2. Is my cloud adoption suffering due to the complexity of managing and migrating applications across various on-premise and cloud environments?**
- 3. Are my development and ops teams struggling to collaborate and spending time in solving environment-related issues?**
- 4. How can I introduce a devops culture in my engineering organization?**

Conclusion

So there you have it: eight technology trends that we think are going to make the biggest impact in 2016. Every enterprise, regardless of scale or domain, should look at leveraging one or more of these in driving its digital transformation strategy. The best way to proceed is pick the biggest opportunities you see to apply digital transformation, and then ask the kind of questions we have included here to determine which digital technologies will help you address those opportunities. Then start from the end-user perspective and solve for creating the ideal digital experience for them. Iterate in small intervals, get feedback from end users, and repeat. These are exciting times for innovation, and we hope this document guides your enterprise digital transformation journey in 2016 and beyond.

Contributors

- Chandraprakash Jain (Senior Architect – CTO Office)
- Chandrakant Shinde (Senior Architect – CTO Office)
- Dhruva Ray (Principal Architect – CTO Office)
- Dr. Pandurang Kamat (Chief Architect – CTO Office)
- Dr. Siddhartha Chatterjee (Chief Technology Officer)



PERSISTENT

About Persistent Systems

Persistent Systems (BSE & NSE: PERSISTENT) is a global company specializing in software product and technology services. For over two decades, Persistent has consistently been selected as the trusted innovation partner for the world's largest technology brands, leading enterprises and pioneering start-ups. Persistent has a global team of more than 7,800 employees worldwide including offices and delivery centers in North America, Europe, and Asia. Persistent develops best-in-class solutions in key next-generation technology areas including Analytics, Big Data, Cloud Computing, Mobility and Social, for the independent software vendors (ISVs), telecommunications and media, life sciences and healthcare, and financial services verticals. For more information, please visit: www.persistent.com

India

Persistent Systems Limited

Bhageerath, 402,
Senapati Bapat Road
Pune 411016.
Tel: +91 (20) 2570 2000
Fax: +91 (20) 2567 8901

USA

Persistent Systems, Inc.

2055 Laurelwood Road, Suite 210
Santa Clara, CA 95054
Tel: +1 (408) 216 7010
Fax: +1 (408) 451 9177
Email: info@persistent.com

DISCLAIMER: "The trademarks or trade names mentioned in this report are property of their respective owners and are included for reference only and do not imply a connection or relationship between Persistent Systems and these companies."