Internet of Things (IoT)

Ben Brock

MSCS 8331 Advanced Data Mining

20-04-2017

* The internet of things (IoT) is a computing concept that describes a scenario where everyday physical objects are connected to the internet and can identify themselves to other devices or processes, via an IP address.
* The IoT is significant because an object that can represent itself digitally becomes something greater than the object by itself.
  * No longer does the object just relate to the process;
  * it now connects to surrounding objects and database data, permitting “big data” analytics and insights.

What is IoT?

The Internet of Things (IoT) is the network of physical objects or "things" embedded with electronics, software, sensors, and network connectivity, which enables these objects to collect and exchange data.

IoT allows objects to be sensed and controlled remotely across existing network infrastructure, creating opportunities for more direct integration between the physical world and computer-based systems, and resulting in improved efficiency, accuracy and economic benefit.

"**Things**," in the IoT sense, can refer to a wide variety of devices such as heart monitoring implants, biochip transponders on farm animals, electric clams in coastal waters, automobiles with built-in sensors, DNA analysis devices for environmental/food/pathogen monitoring or field operation devices that assist fire-fighters in search and rescue operations.

These devices collect useful data with the help of various existing technologies and then autonomously flow the data between other devices.

How IoT Works?

Internet of Things is not the result of a single novel technology; instead, several complementary technical developments provide capabilities that taken together help to bridge the gap between the virtual and physical world. These capabilities include:

- ***Communication and cooperation***
- ***Addressability***
- ***Identification***
- ***Sensing***
- ***Actuation***
- ***Embedded information processing***
- ***Localization***
- ***User interfaces***

How IoT Works?

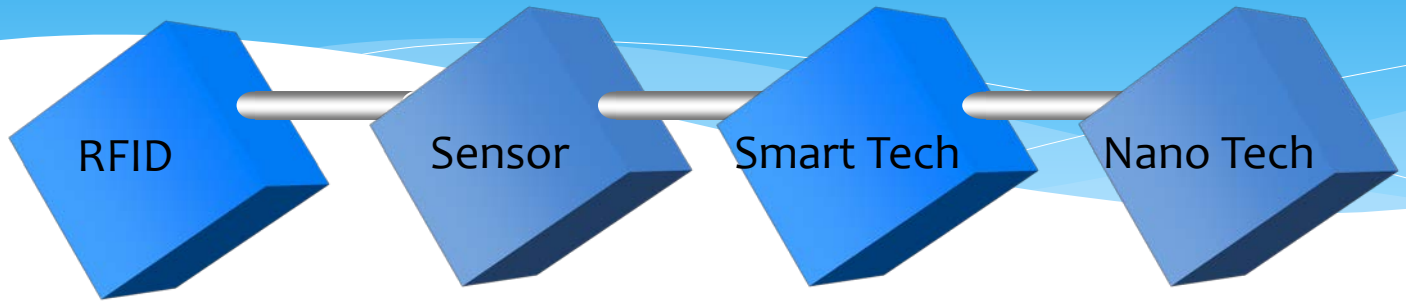

To identify and track the data of things	To collect and process the data to detect the changes in the physical status of things	To enhance the power of the network by devolving processing capabilities to different part of the network.	To make the smaller and smaller things have the ability to connect and interact.

Current Status & Future Prospect of IoT

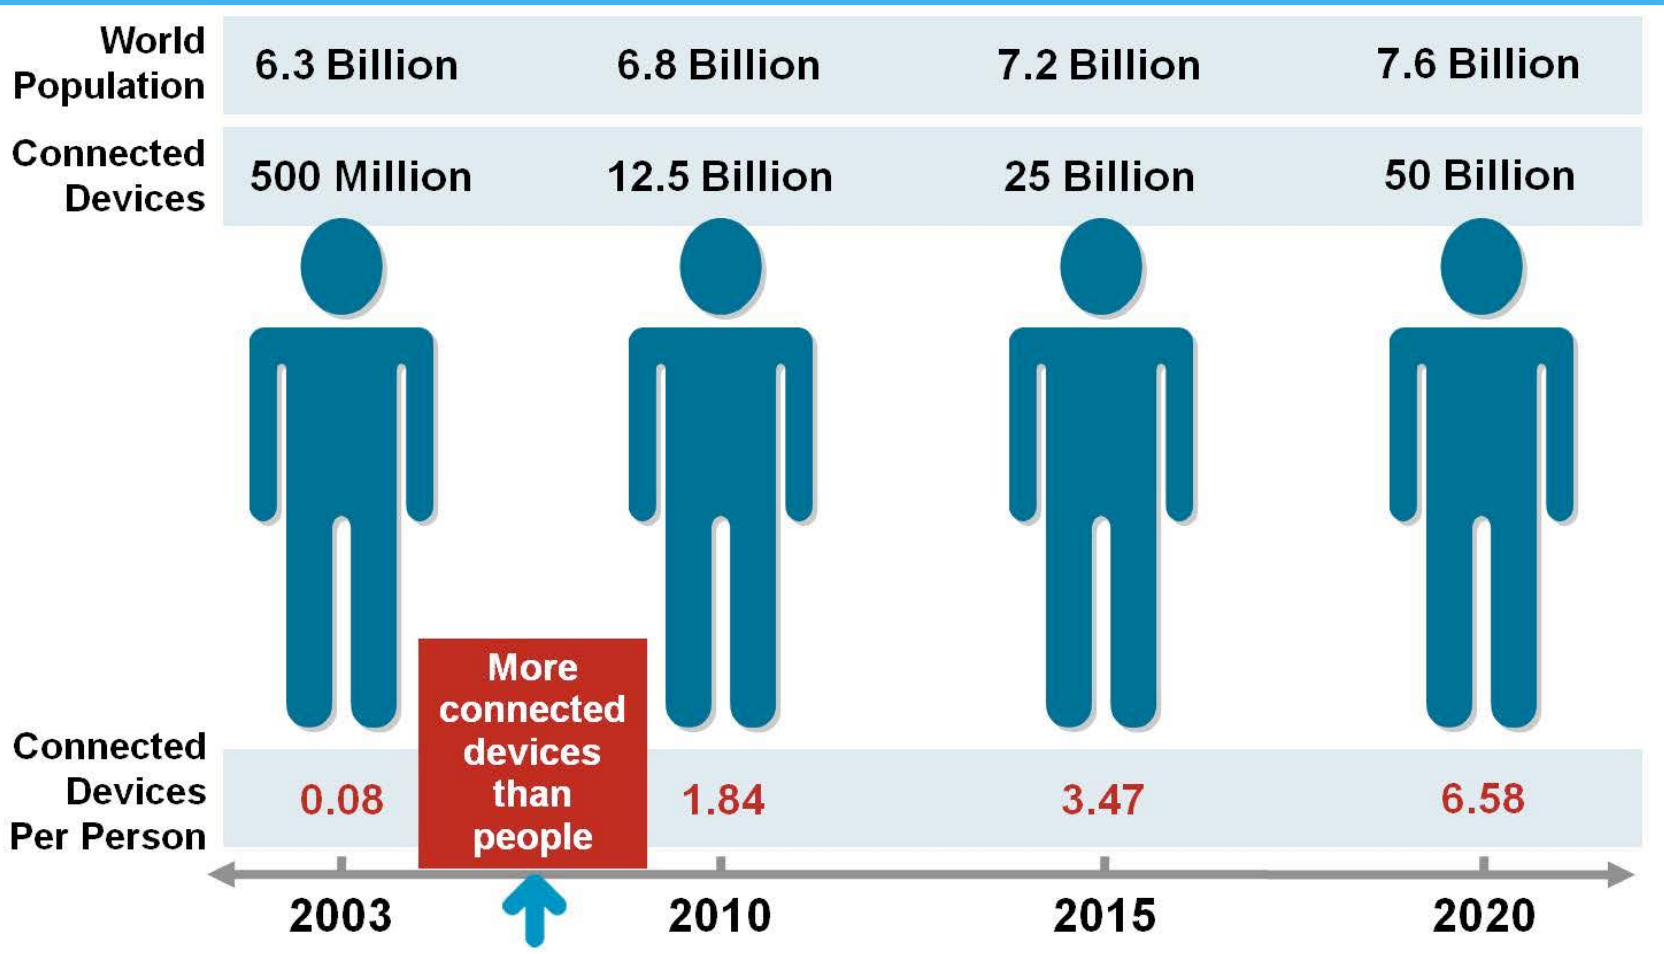

"Change is the only thing permanent in this world"

IoT as a Network of Networks:

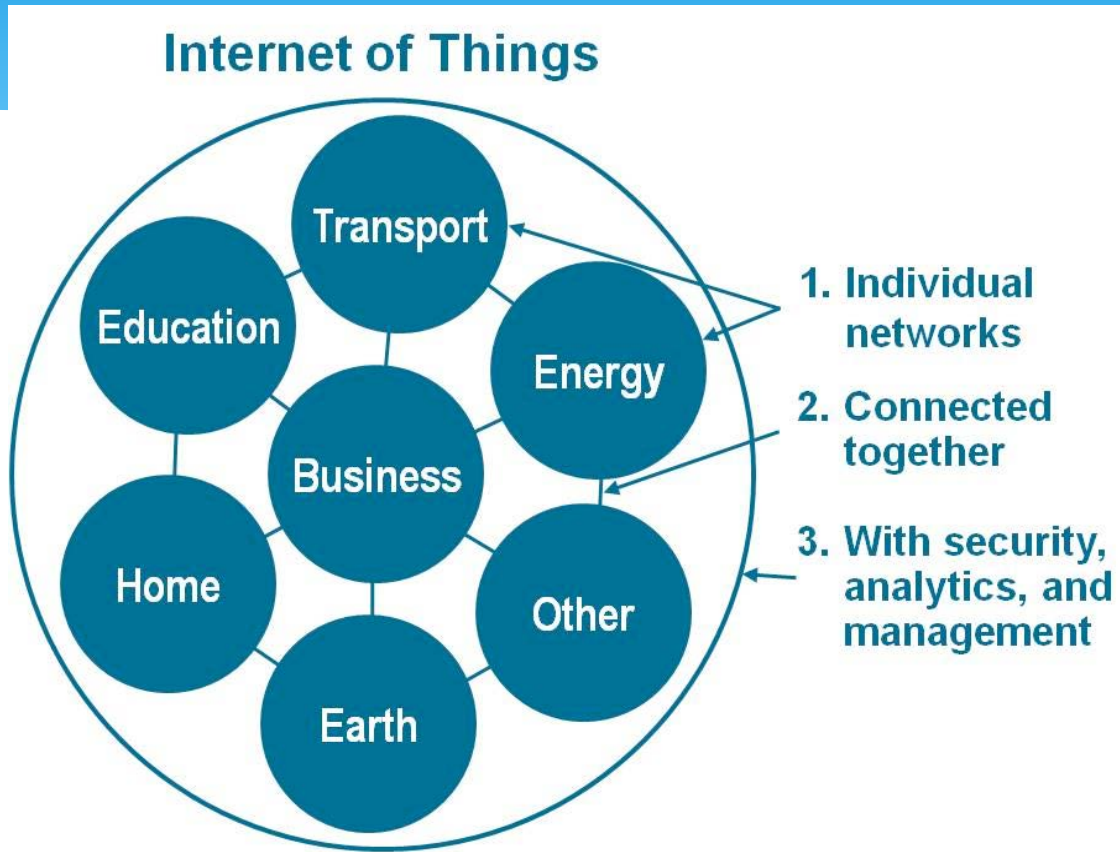

These networks connected with added security, analytics, and management capabilities. This will allow IoT to become even more powerful in what it can help people achieve.

The Internet of Things

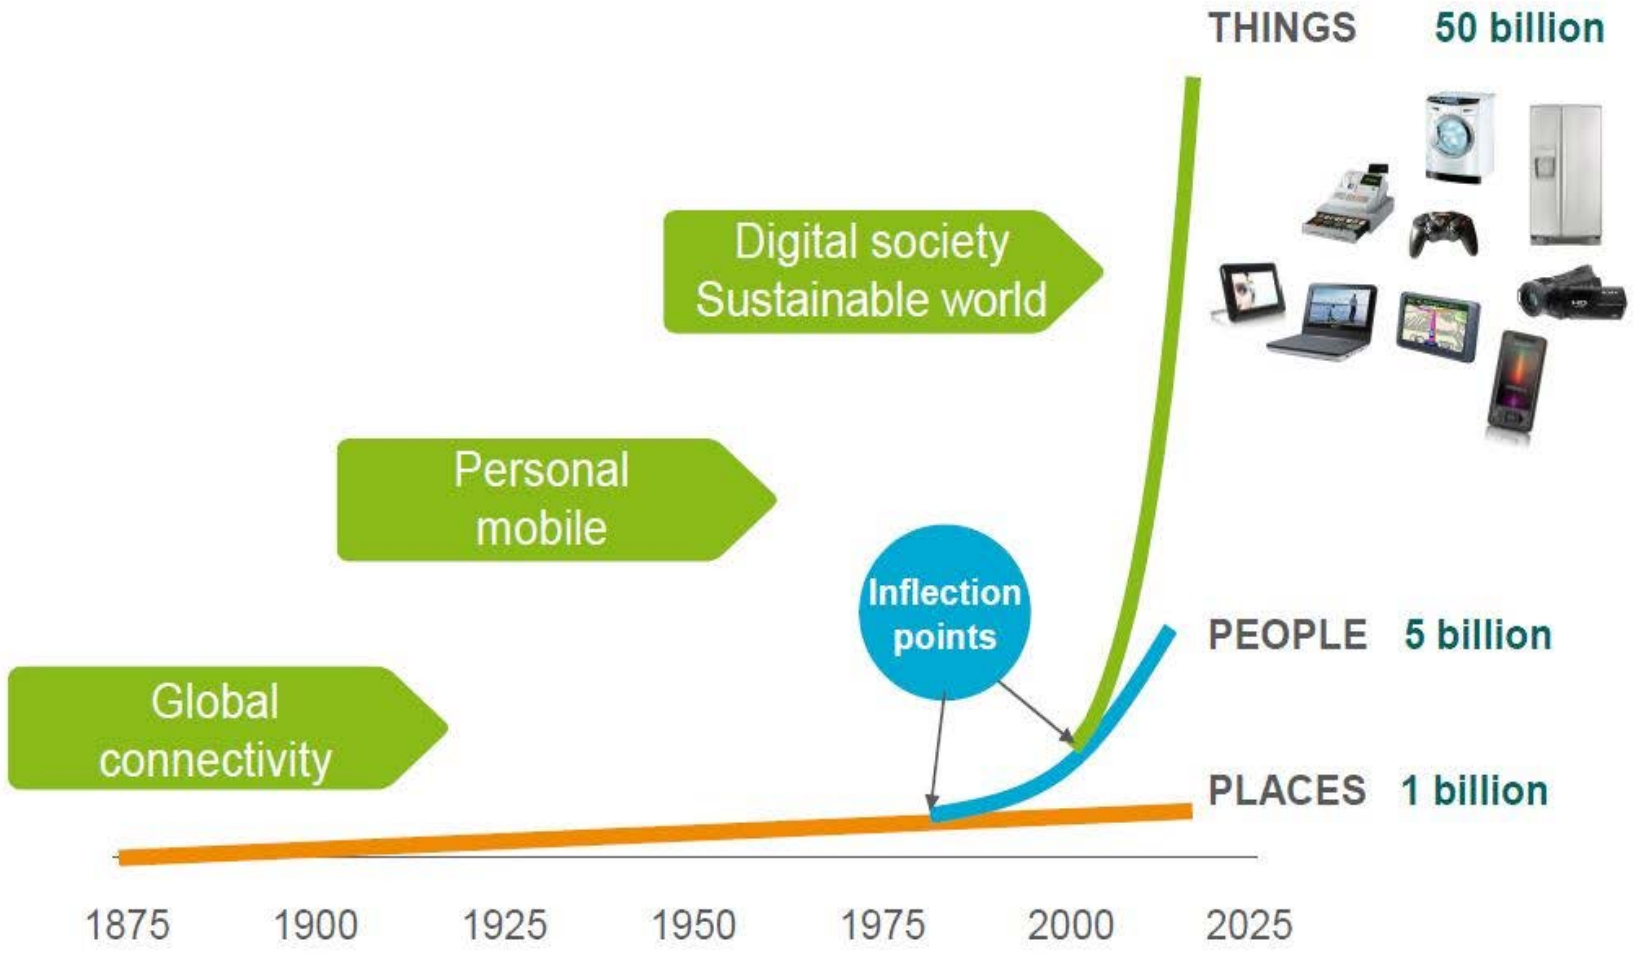

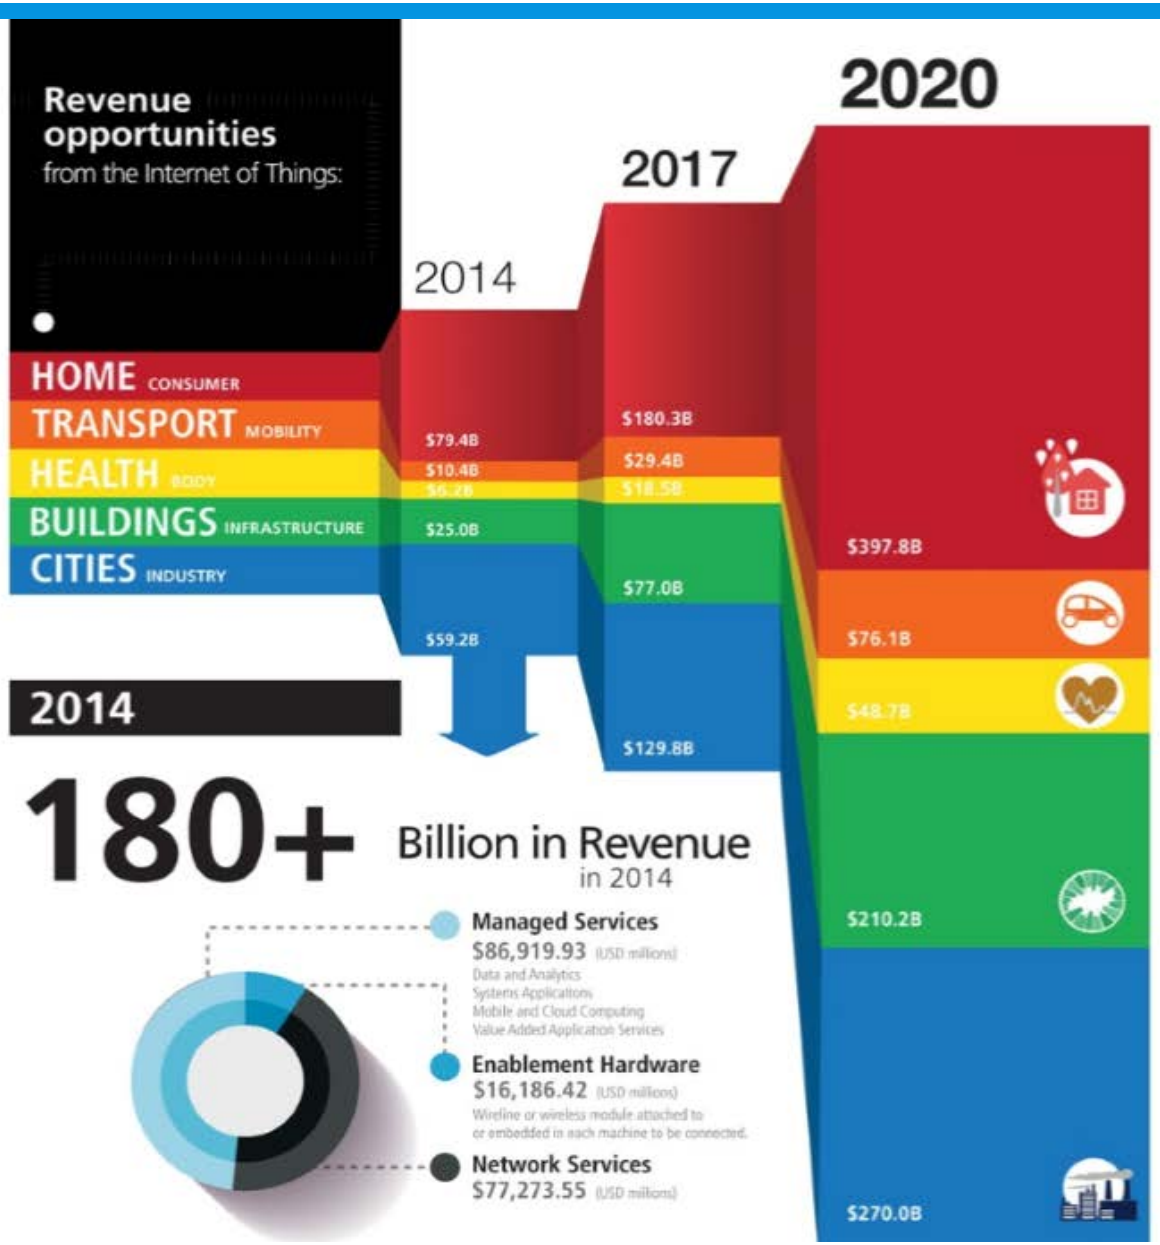
Revenue opportunities
from the Internet of Things:
2014
2017
2020
HOME CONSUMER
TRANSPORT MOBILITY
HEALTH BODY
BUILDINGS INFRASTRUCTURE
CITIES INDUSTRY
$79.4B
$10.4B
$25.0B
$59.2B
$180.3B
$29.4B
$77.0B
$129.8B
$397.8B
$76.1B
$48.7B
$210.2B
$270.0B
2014
180+ Billion in Revenue
in 2014
Managed Services
$86,919.93 (USD millions)
Data and Analytics
Systems Applications
Mobile and Cloud Computing
Value Added Application Services
Enablement Hardware
$16,186.42 (USD millions)
Wireline or wireless module attached to or embedded in each machine to be connected.
Network Services
$77,273.55 (USD millions)

2017 Top Tech Trends

- **1) Internet of Things**
- The IoT promises to be the most disruptive technological revolution since the advent of the World Wide Web. Projections indicate that up to 100 billion uniquely identifiable objects will be connected to the Internet by 2020, with enormous technical, socioeconomic, political, and even spiritual consequences.

IEEE

Gartner

Top Ten Strategic Technology Trends for 2014 :

- Mobile Device Diversity and Management
- Mobile Apps and Applications
- The Internet of Everything
- Hybrid Cloud and IT as Service Broker
- Cloud/Client Architecture
- The Era of Personal Cloud
- Software Defined Anything
- Web-Scale IT
- Smart Machines
- 3-D Printing

What is the Internet of Things?

- Cisco "Internet of Everything"
  - …the latest wave of the Internet -- connecting physical objects…to provide better safety, comfort, and efficiency
- IBM "Internet of Things"
  - …a completely new world-wide web, one comprised of the messages that digitally empowered devices would send to one another. It is the same Internet, but not the same Web.
- GE "Industrial Internet"
  - …convergence of machine and intelligent data…to create brilliant machines
- RTI "Your Systems. Working as One."
  - …an entirely new utility. As profound as the cell network, GPS, or the Internet itself. The Internet of Things and the Intelligent Systems it enables will fundamentally change our world.

Core Nervous System for the IoT

Applications of IoT

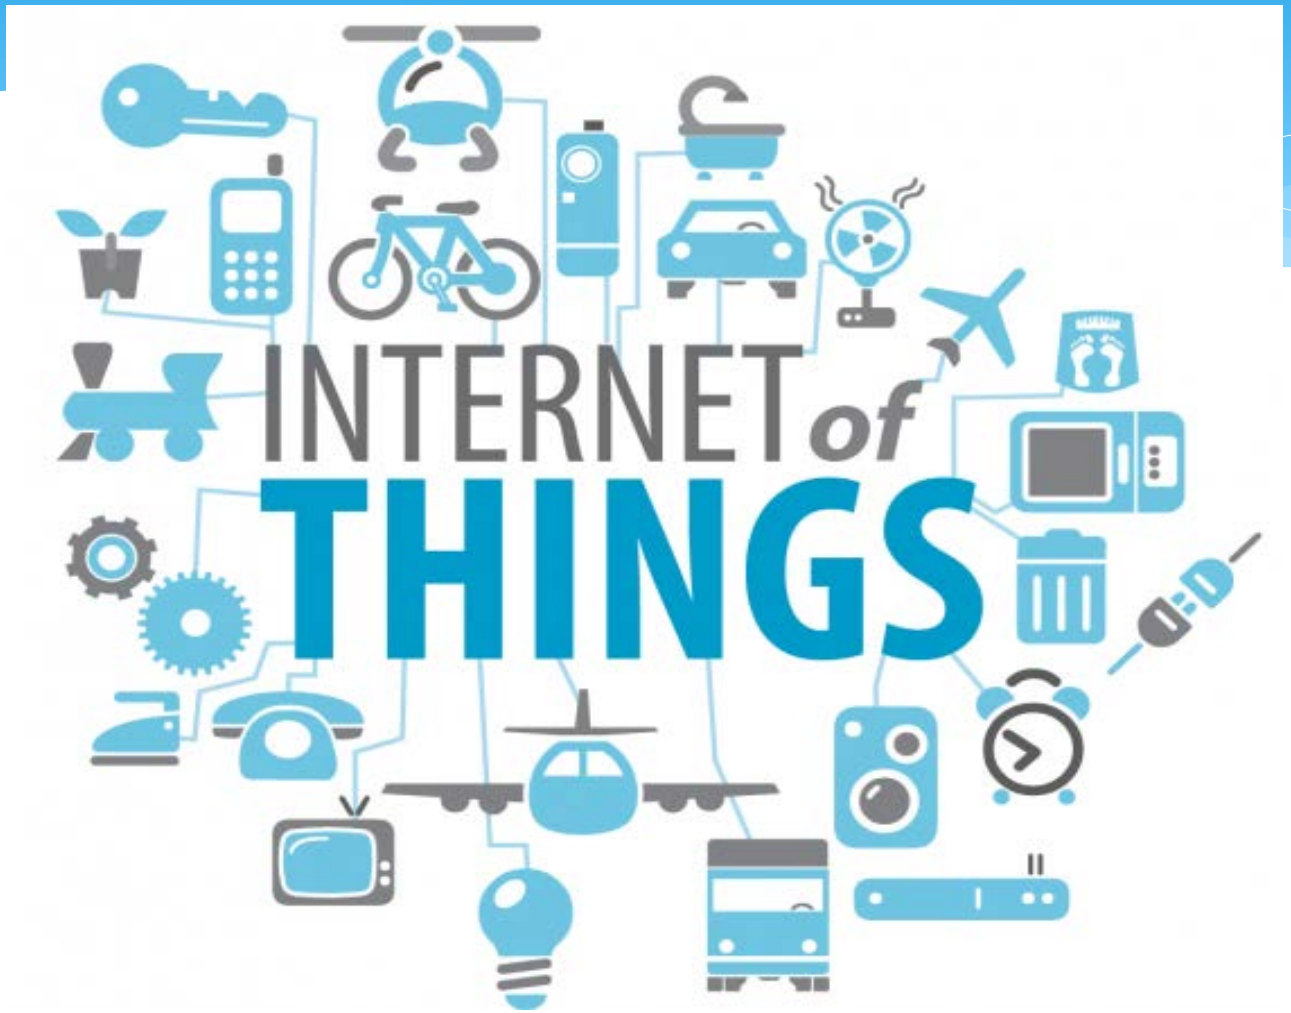

"The Ultimate Goal of IOT is to Automate Human Life."

Few Applications of IoT

- ✓Building and Home automation
- ✓Manufacturing
- ✓Medical and Healthcare systems
- ✓Media
- ✓Environmental monitoring
- ✓Infrastructure management
- ✓Energy management
- ✓Transportation
- ✓Better quality of life for elderly
- ✓… … …

You name it, and you will have it in IoT!

TO DIVERSE APPLICATIONS

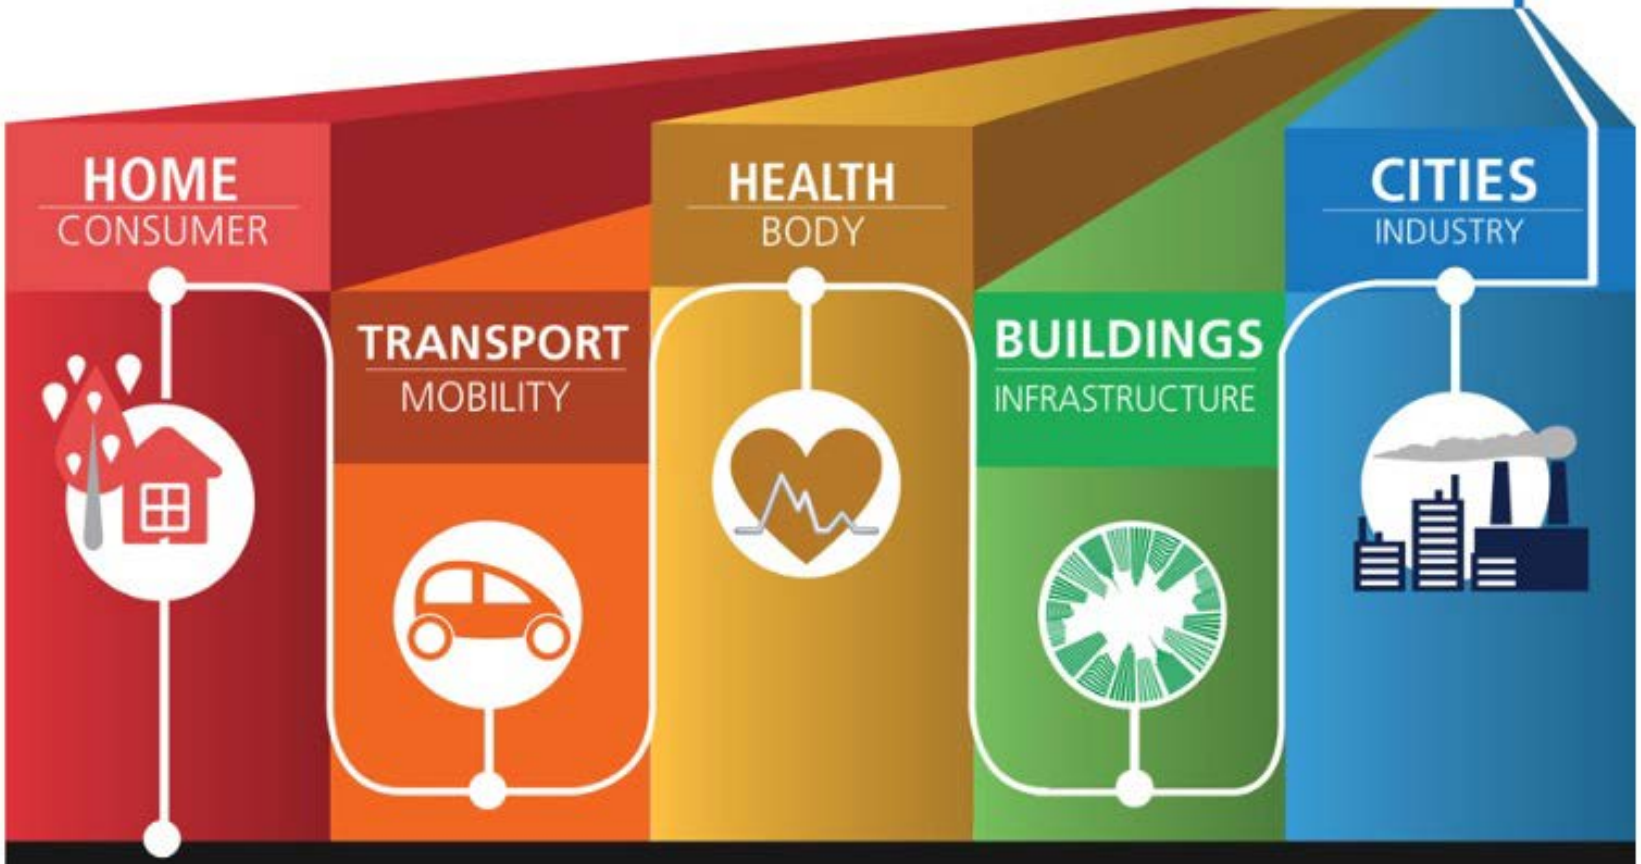

Light bulbs
Security
Pet Feeding
Irrigation Controller
Smoke Alarm
Refrigerator
Infotainment
Washer / Dryer
Stove
Energy Monitoring

Traffic routing
Telematics
Package Monitoring
Smart Parking
Insurance Adjustments
Supply Chain
Shipping
Public Transport
Airlines
Trains

Patient Care
Elderly Monitoring
Remote Diagnostic
Equipment Monitoring
Hospital Hygiene
Bio Wearables
Food sensors

HVAC
Security
Lighting
Electrical
Transit
Emergency Alerts
Structural Integrity
Occupancy
Energy Credits

Electrical Distribution
Maintenance
Surveillance
Signage
Utilities / Smart Grid
Emergency Services
Waste Management

IoT Emergency Example

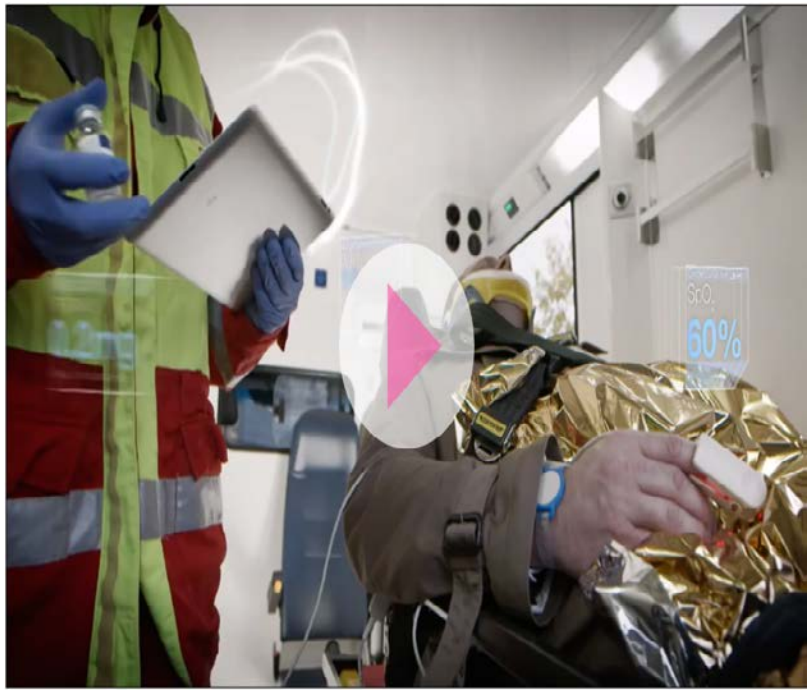

* IoT Emergency Example
* Byron's Story

Efficient Waste Management in Smart Cities Supported by the Sensing-as-a-Service

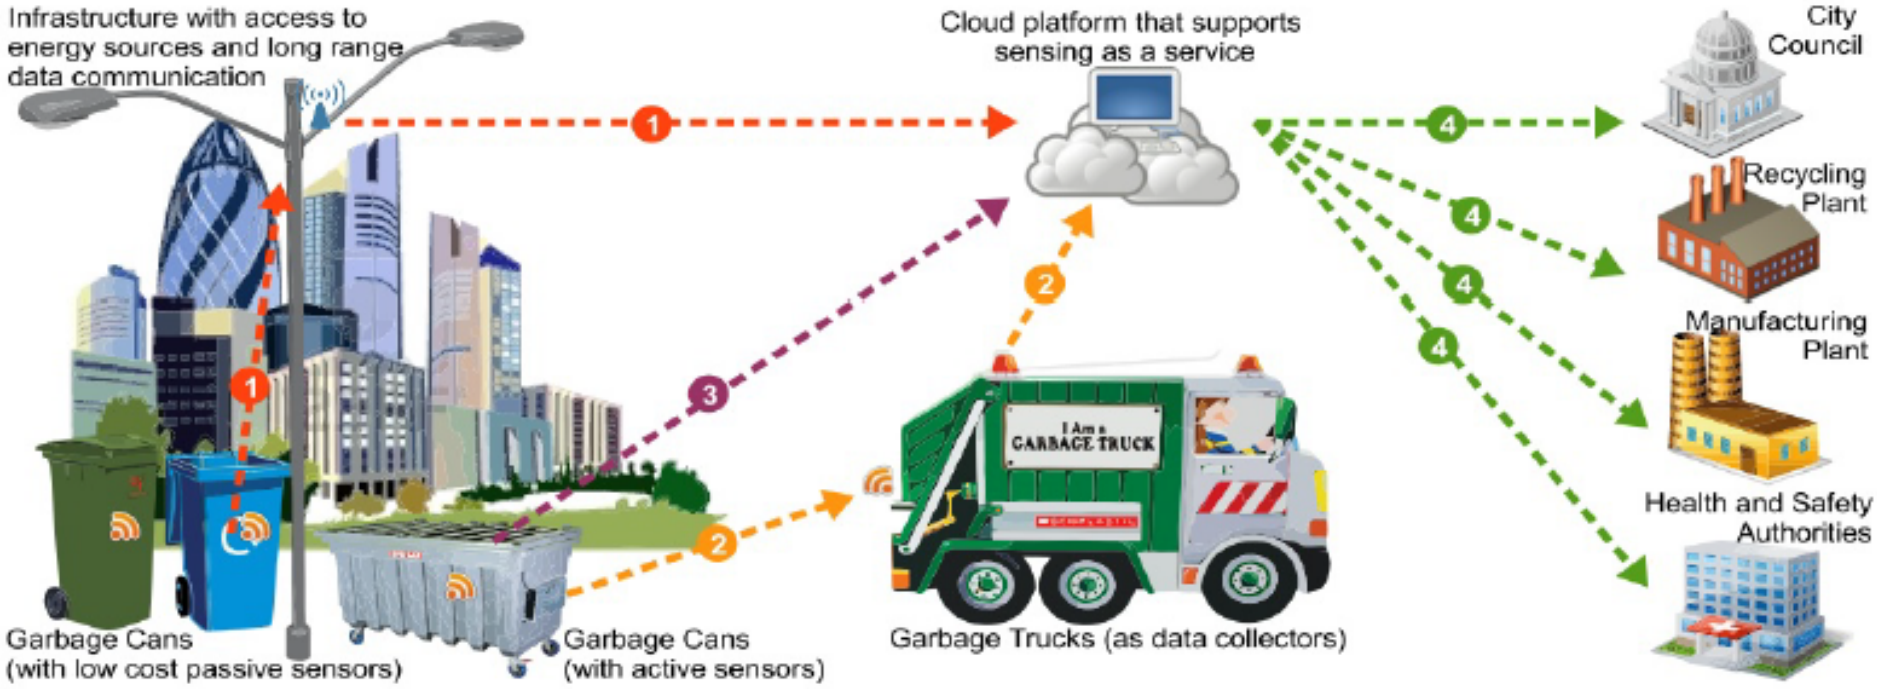

[Source: "Sensing as a Service Model for Smart Cities Supported by Internet of Things", Charith Perera et. al., Transactions on Emerging Telecommunications Technology, 2014]

IOT Application Scenario - Shopping

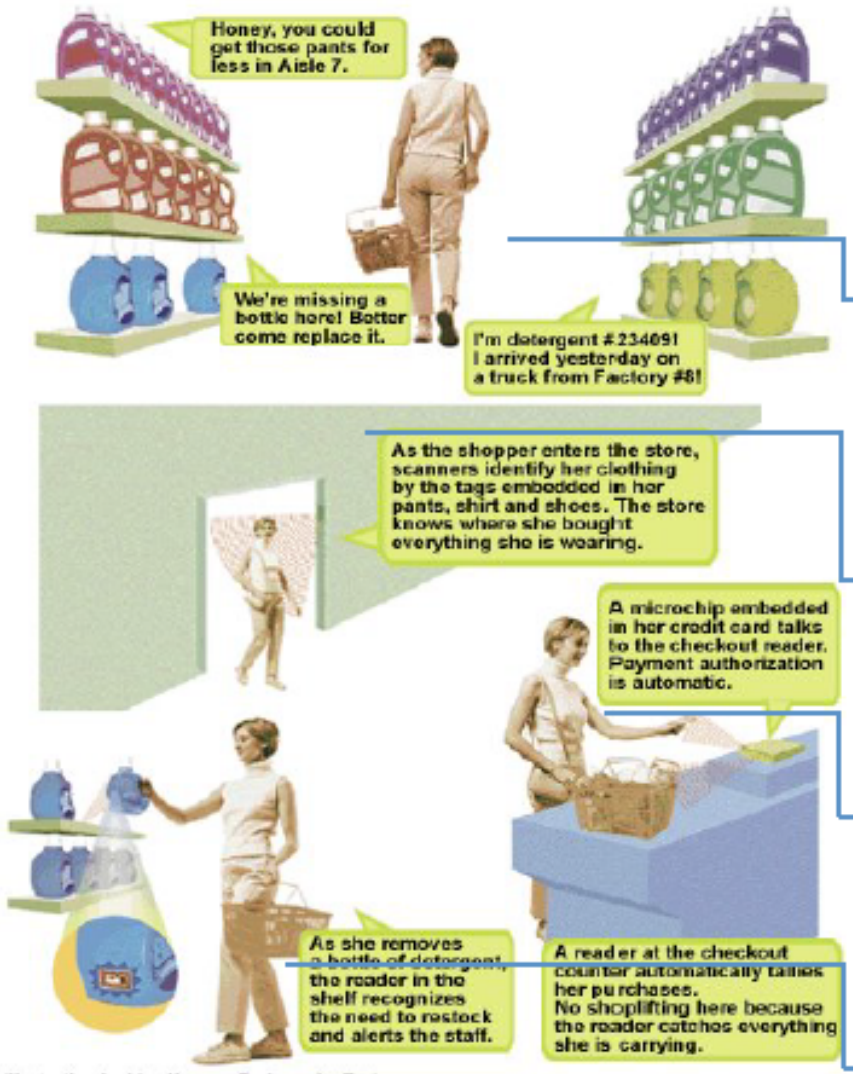

Illustration by Lisa Kreuse Brailman for Forbes

(2) When shopping in the market, the goods will introduce themselves.

(1) When entering the doors, scanners will identify the tags on her clothing.

(4) When paying for the goods, the microchip of the credit card will communicate with checkout reader.

(3) When moving the goods, the reader will tell the staff to put a new one.

How Well Do I Sleep?

Sleep

zzzZ

YOUR SLEEP EFFICIENCY
92%

Your sleep pattern asleep awake

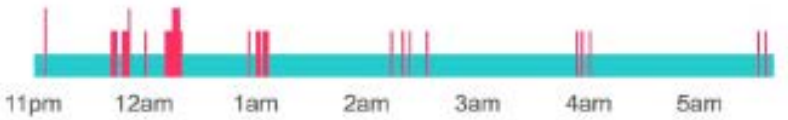

You went to bed at	Time to fall asleep	Times awakened	You were in bed for	Actual sleep time
11:00PM	0min	20	6hrs 40min	6hrs 6min

Sleep Stats

Time asleep over the past 30 days in hours

Times awoken over the past 30 days

8 h 50 mins asleep

1d 1w 1m 3m 1y

- Awake for 212 mins (81x)
- Restless for 278 mins (91x)

Thursday, February 27

fitbit flex
Wireless Activity + Sleep Wristband

The challenge of IoT

Total challenge of IOT

1.Technological Standardization in most areas are still remain fragmented.

2.Managing and fostering rapid innovation is a challenge for governments .

3. Privacy.

4.Absence of governance.

The challenge of IoT

Total challenge of IOT

1.Technological Standardization in most areas are still remain fragmented.

2.Managing and fostering rapid innovation is a challenge for governments .

3. Privacy.

4.Absence of governance.


IN 60 SECONDS...
1 NEW DEFINITION IS ADDED ON urban
1,600+ READS ON Scribd.
13,000+ HOURS MUSIC STREAMING ON PANDORA
12,000+ NEW ADS POSTED ON craigslist
370,000+ MINUTES VOICE CALLS ON skype
98,000+ TWEETS
20,000+ NEW POSTS ON tumblr.
THE LARGEST SOCIAL READING PUBLISHING COMPANY
New Craigslist Ads
320+ NEW twitter ACCOUNTS
100+ NEW Linked in ACCOUNTS
13,000+ iPhone APPLICATIONS DOWNLOADED
1 associatedcontent NEW ARTICLE IS PUBLISHED
THE WORLD'S LARGEST COMMUNITY CREATED CONTENT!!
QUESTIONS ASKED ON THE INTERNET...
100+ Answers.com
40+ YAHOO! ANSWERS
6,600+ NEW PICTURES ARE UPLOADED ON flickr
600+ NEW VIDEOS
50+ WORDPRESS DOWNLOADS
25+ HOURS TOTAL DURATION
70+ DOMAINS REGISTERED
695,000+ facebook STATUS UPDATES
125+ PLUGIN DOWNLOADS
60+ NEW BLOGS
168 MILLION EMAILS ARE SENT
694,445 SEARCH QUERIES
1,700+ Firefox DOWNLOADS
1,500+ BLOG POSTS
79,364 WALL POSTS
510,040 COMMENTS
Google
Google Search
GO-Globe.com

Future of IOT

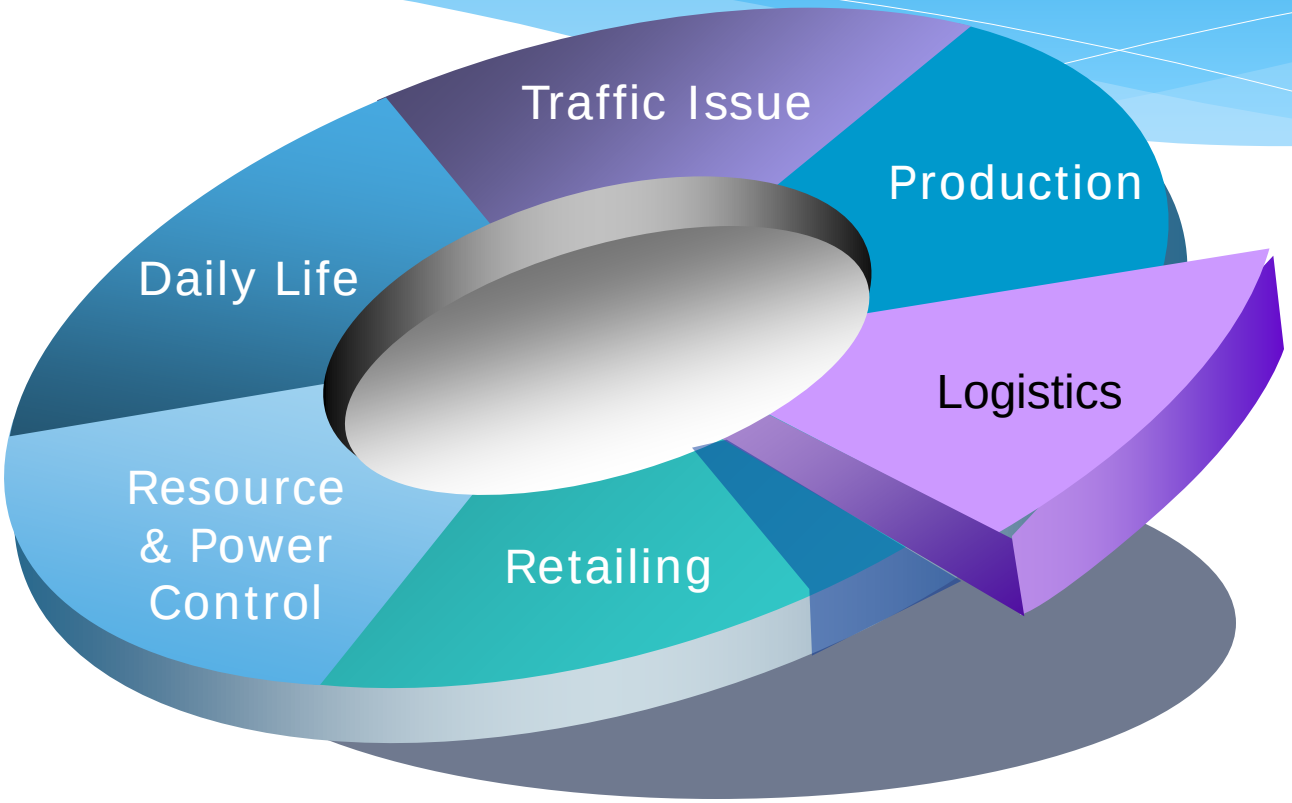

Future of IoT:

help organizations change and innovate faster
by making it easy to
connect the world's applications, data and devices

IoT Security

Focus: Securing for the future

2016 saw the rapid increase in the use of ransomware and a resurgence of distributed denial of service (DDoS) attacks using hijacked devices making up the internet of things (IoT) against a background of new and newly-reported breaches of personal information involving an ever-increasing number of online user accounts for a range of online services. In 2017, these trends are expected to continue as ransomware and IoT botnet malware evolve to become even more sophisticated, but there is also an increase in professional, advanced attacks – including attacks on cloud infrastructure – and the rise of data manipulation attacks, further underlining the need for a fresh approach to data security. Cyber security education and training is more important than ever to enable organizations to truly understand the risks and vulnerabilities, and then look for the most effect way to manage and mitigate them, which may include new security technologies such as behavioral analytics and artificial intelligence to tip the balance in favor of the cyber defenders.

IoT Security

Far from easing our lives, they warn, if we're not careful, the internet of things (IoT) could end them.

Cesare Garlati, chief security strategist for the prpl Foundation, an open source consortium working on next-generation datacentre software and architectures, says: "Most of these IoT devices are connected to, or directly control, physical objects – an elevator or heating system, for example. Therefore a breach doesn't just represent a traditional loss of data with resulting fines, but a physical attack that might involve human casualties or fatalities."

Potential to wreak havoc

* IoT Devices
  * smart thermostats to connected cameras,
  * medical implants to industrial controllers,
  * a succession of devices has been shown to be hackable, many with the potential to wreak economic,
    * domestic and physical havoc.
* And there are plenty of miscreants eager to gain such power over our lives, businesses and economy
  * criminals hoping to hold us to ransom for financial gain,
  * cyber terrorists bent on causing mayhem and state actors engaged in clandestine cyber warfare.

IoT Security Essentials

* IoT hardware manufacturers and integrators must:
  * Specify hardware to minimum requirements so a device is not capable of doing more than it needs.
  * Ensure all hardware is tamper-proof, with no internal or external USB ports, for instance.
  * Build equipment should be built around secure hardware such as Trusted Platform Module (TPM).
  * Ensure there is a secure path for firmware upgrades.

IoT Security Essentials

* IoT solution developers must:
  * Follow secure software development methodology.
  * Ensure any open-source software you choose has an active community addressing any security issues that arise.
  * Integrate with care: check all interfaces of components for security flaws, paying particular attentions to superfluous functionality that may be available via an API layer.

IoT Security Essentials

* IoT solution deployers must:
  * Ensure all deployed hardware is tamper-proof - particularly where left unsupervised or in public spaces.
  * Keep authentication keys safe after the deployment. Any compromised key can be used by a malicious device to masquerade as an existing device.

IoT Security Essentials

* IoT solution operators must:
  * Keep the system up to date with the latest OSs and drivers
  * Protect against malicious activity by securing device operating systems with the latest anti-malware capabilities.
  * Audit the IoT infrastructure often for security-related issues
  * Physically protect the infrastructure from malicious access
  * Protect cloud authentication credentials by changing passwords frequently, and not logging on from public machines.

IoT hacks that hit the headlines

* 2010: Stuxnet (believed to have been created by Israeli intelligence) vibrates centrifuges in Iran nuclear plant.
* 2011: Hacker takes wireless control of insulin pumps.
* 2014: Hackers commandeer hundreds of webcams and baby monitors.
* 2015: Researchers remotely take over and crash Cherokee jeep.
* 2015: Plane flight controls hacked via in-flight entertainment system.
* 2016: Smart thermostats hacked to host ransomware

IoT hacks that hit the headlines Continued

* 2016: Outages caused by result of connected devices such as DVRs and IP cameras that hackers turned into botnets
  * Hackers repeatedly send web requests to overwhelm web servers and take them down
  * Causes
    * Username/password hardcoded into the firmware
    * IoT's setup to use default username/password
* 2016: In February 2016, the Hollywood Presbyterian Medical Center paid a $17,000 bitcoin ransom for the decryption key for patient data.
  * The Hospital was infected by the delivery of an email attachment disguised as a Microsoft Word invoice.
  * There appears to be a trend in ransomware being used to attack hospitals and it appears to be growing

Why IoT sensors are game changers in supply chain analytics

* How IoT sensors differ
  * Big Data
  * Analytics
* IoT devices add context, automatic response
  * includes tablets and smartphones to access, analyze and display the data among a much wider universe of people, anywhere, anytime.
* Where are they?

"Big Data is not magic. It doesn't matter how much data you have if you can't make sense of it."

Characteristics of Big Data

* Volume
  * Large amounts of data
  * eCommerce, social media, sensors
* Velocity
  * Rate at which data flows
  * Streaming data, Batch/Real-time
* Variety
  * Different formats of data
  * Structured, unstructured, semi-structured
* Veracity
  * Accuracy of data
  * Data quality, trust associated with data

Hadoop

* Batch Mode
* Real-time Streaming
  * Apache Spark 2.0
  * Apache Kudu

Cloudera

* “As far back as 2012, Cloudera recognized the analytic gap in the Hadoop ecosystem that was leading architects to create complex hybrid architectures for real-time analytics. With the Apache Kudu 1.0 launch, the original vision is coming to fruition as users can now rely on a single, simplified project for fast analytics on fast data. We’ve seen the community quickly adopt Kudu and apply it to numerous high-scale, real-time analytic use cases.”

Spark 2.0

* With Spark 2.0, organizations are better able to take advantage of streaming data, develop richer machine learning models and deploy them in real time, enabling more workloads to go into production.

Spark 2.0 features include

* Better performance and enhanced usability with the new Dataset API
* Structured Steaming for better performance and easier ingest of traditional structured data, for time series, tabular and Internet of Things (IoT) data
* Compile-time type safety for user defined functions, for improved reliability in mission-critical applications
* Machine learning model, pipeline persistence and newly supported machine learning libraries to take on new data sets and analytic applications

Apache Kudu

* Kudu offers fast scans across data for analytics, and instant read/write capabilities for frequent updates and searches.
* Kudu also enables enterprises to adopt real-time use cases at a greater rate.
* Along with its integration with Spark, Kudu 1.0 is also tightly integrated with MapReduce and Impala to enable best-in-class processing.

Apache Kudu features include

* A simplified architecture that enables very fast batch and stream processing
* Fault tolerance and scalability into the hundreds of nodes
* A columnar structure that enables analytic analysis on the latest data, for real-time use cases such as time series data, machine data analytics and online reporting

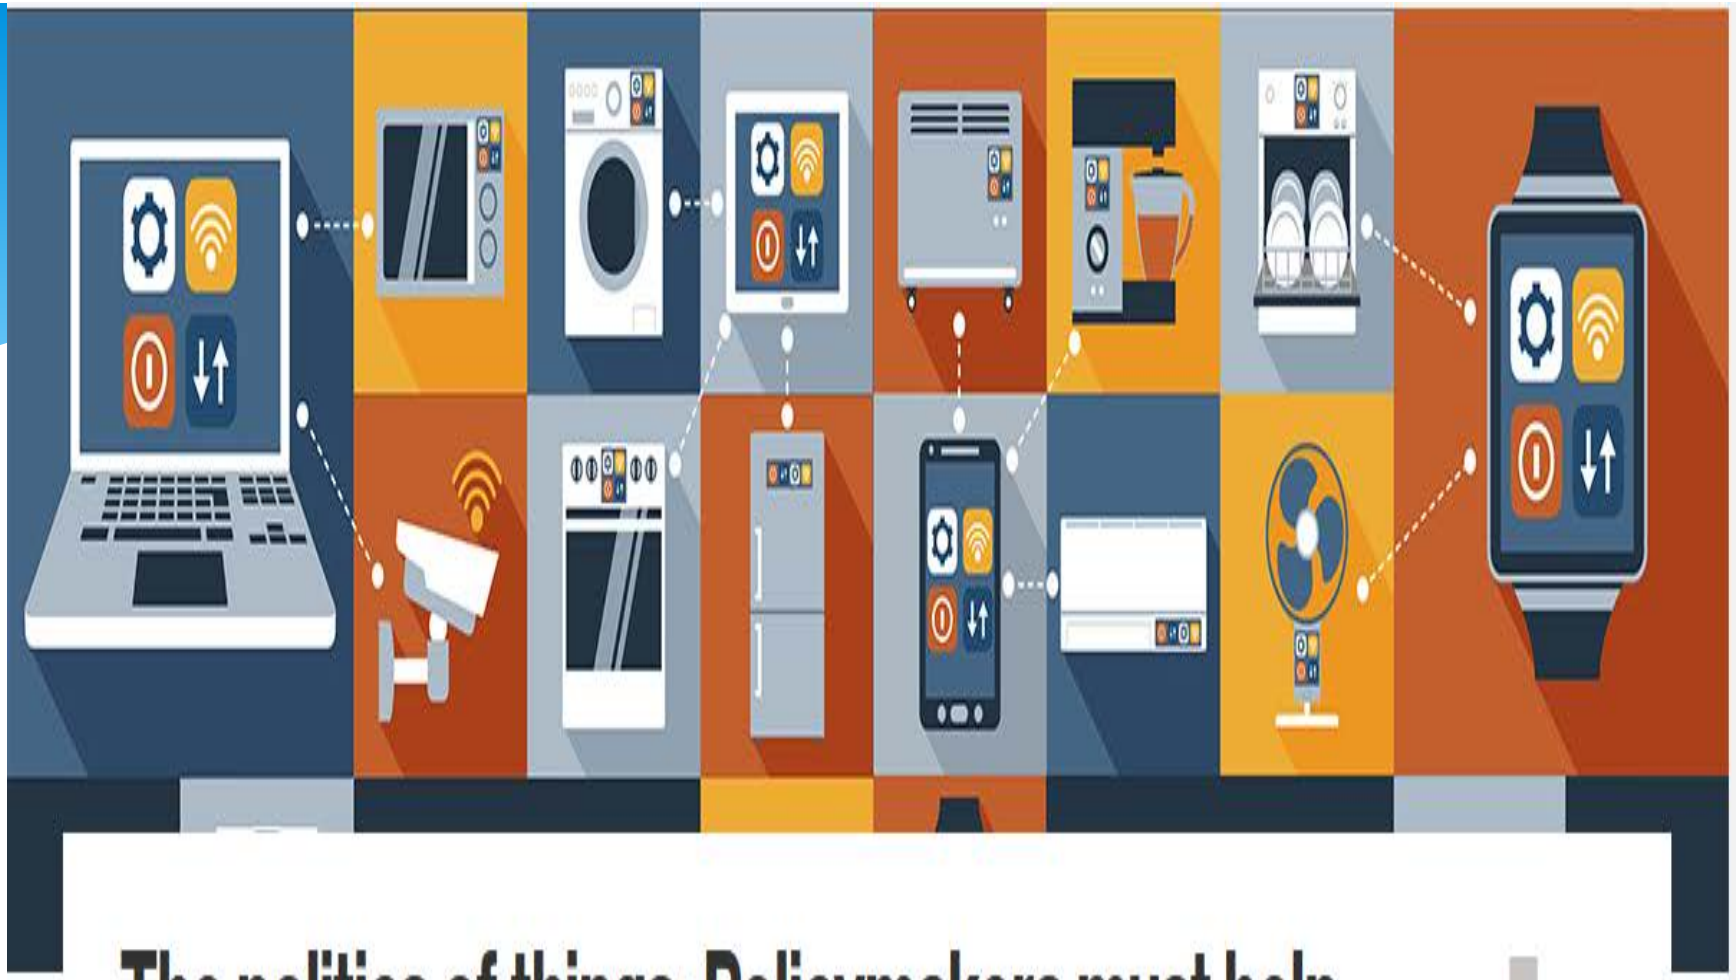

The politics of things: Policymakers must help shape the IoT

IoT HealthCare

* The internet of things has numerous applications in healthcare,
  * from remote monitoring
  * smart sensors and
  * medical device integration.
* It has the potential to not only keep patients safe and healthy,
  * but to improve how physicians deliver care as well.
* Healthcare IoT can also boost patient engagement and satisfaction by allowing patients to spend more time interacting with their doctors.

Boston Medical Center tap IoT Patient Care

* Hospital uses sensors for security purposes
  * Newborn babies are given wristbands
    * allowing a wireless network to locate them at any time
    * If a newborn is taken too close to an exit door without being signed out, elevators will stop and exit doors will lock
  * In Neonatal intensive care unit
    * nurses receive critical alerts on hospital cell phones about their patients' medical conditions,
      * including heart rate and oxygen changes that sensors have detected,
        * allowing them to get to patients' bedsides more quickly

Boston Medical Center

At IBM's Watson lab, customers marry the power of AI with the IoT

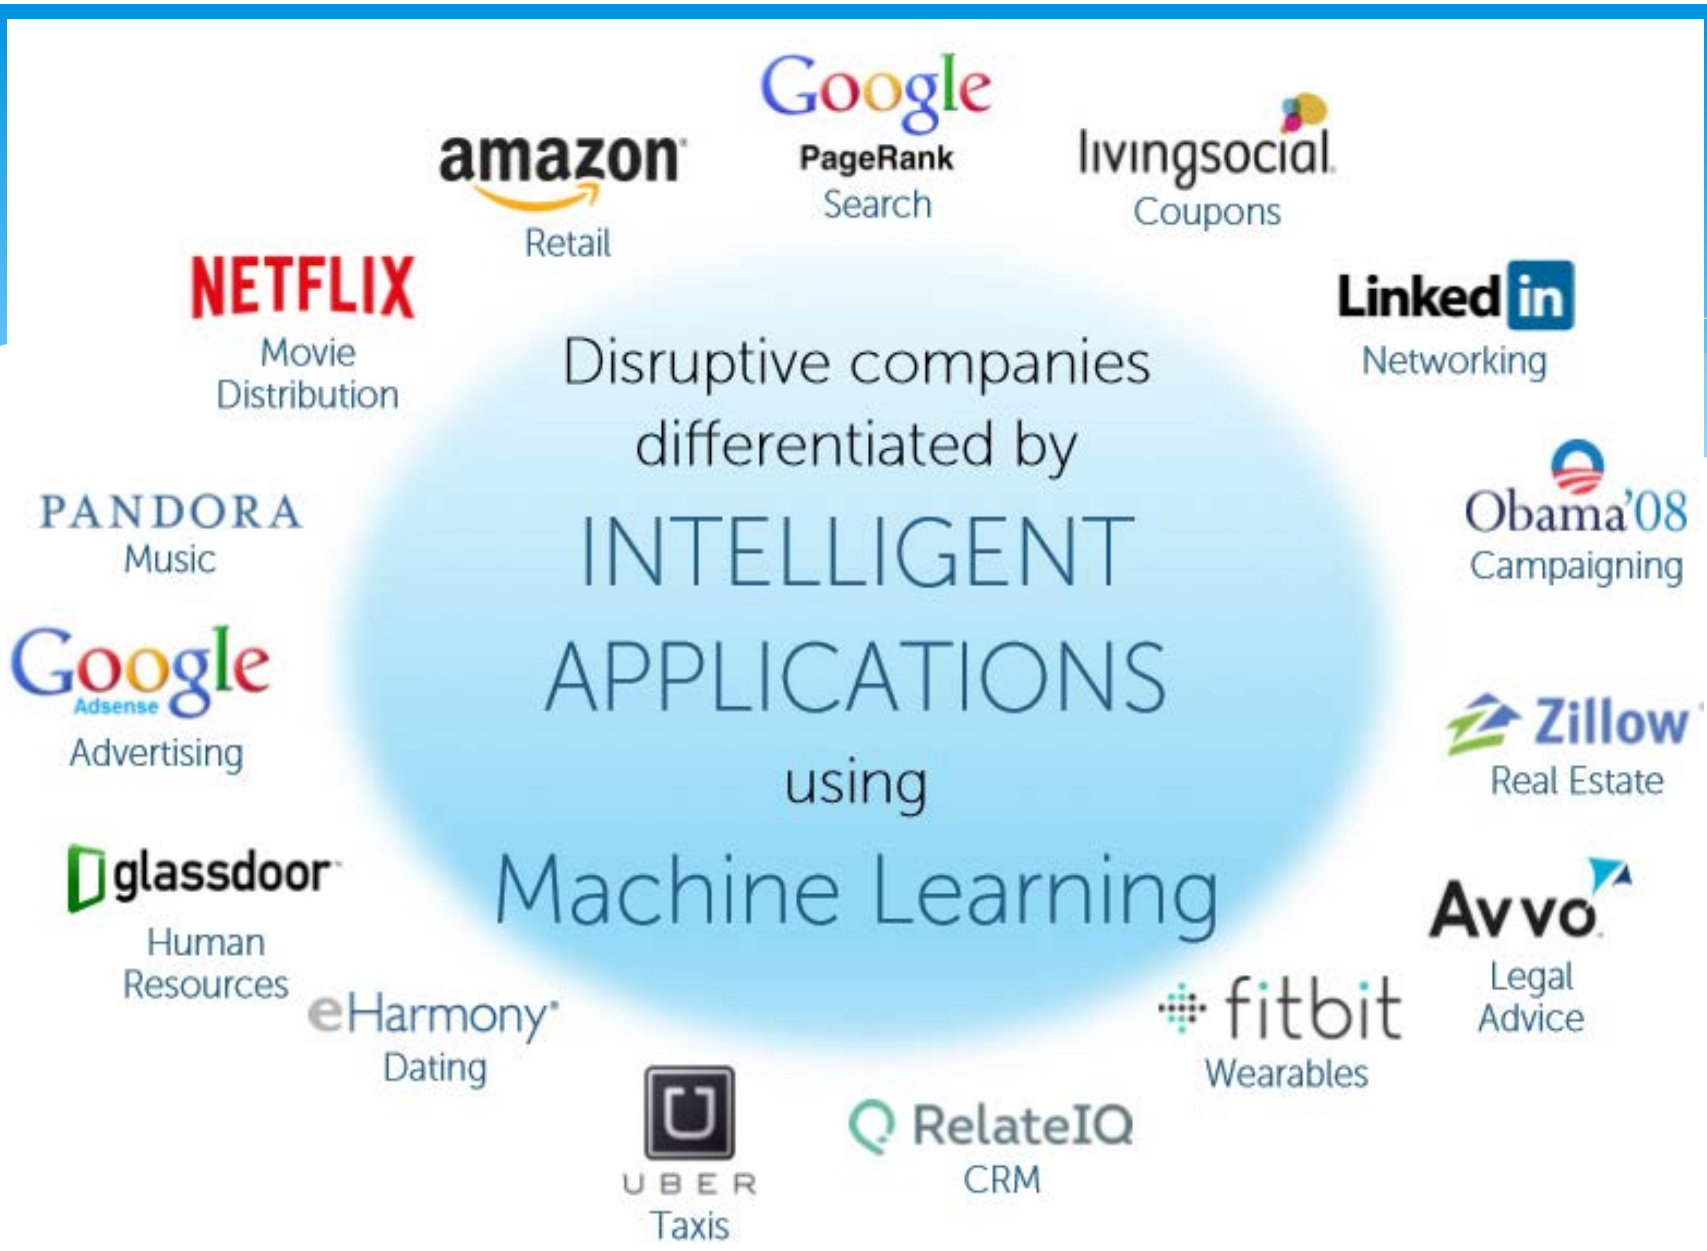
Google
PageRank
Search
amazon
Retail
livingsocial
Coupons
NETFLIX
Movie
Distribution
LinkedIn
Networking
Disruptive companies
differentiated by
INTELLIGENT
APPLICATIONS
using
Machine Learning
PANDORA
Music
Obama'08
Campaigning
Google
Adsense
Advertising
Zillow
Real Estate
glassdoor
Human
Resources
Avvo
Legal
Advice
eHarmony
Dating
fitbit
Wearables
UBER
Taxis
RelateIQ
CRM

Tutorials

Python Turi GraphLab

* IPython Notebooks
  * Document retrieval
  * Song Finder
    * Build songs that you might want to listen too!
  * Deep Learning

Document retrieval

Product recommendation

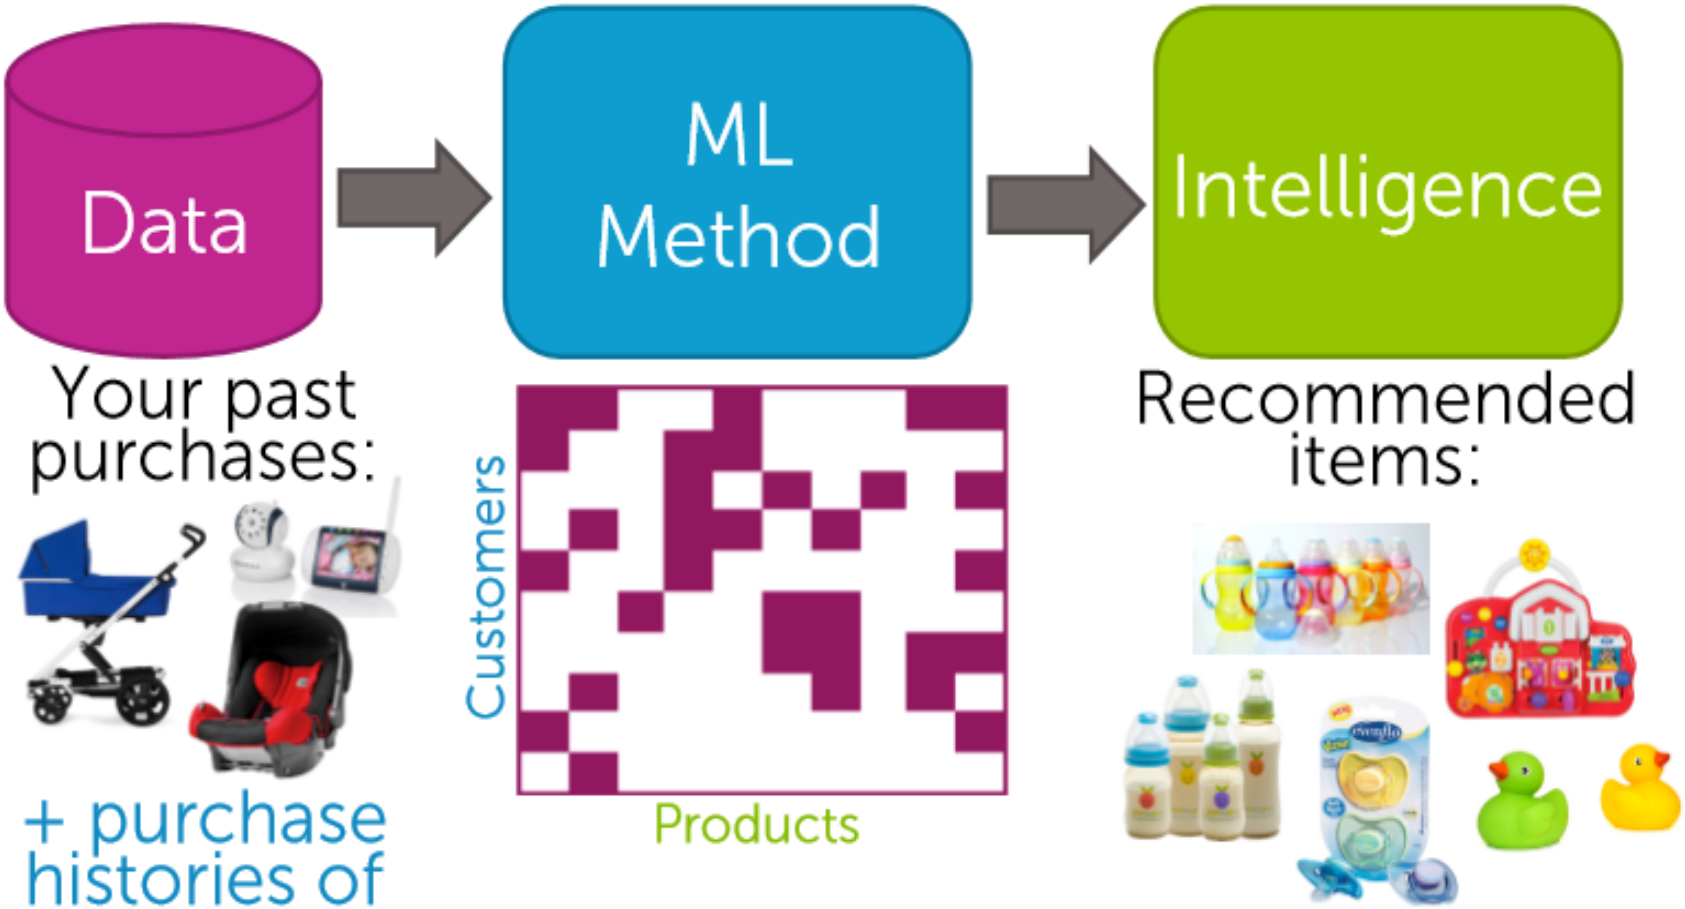

Visual product recommender

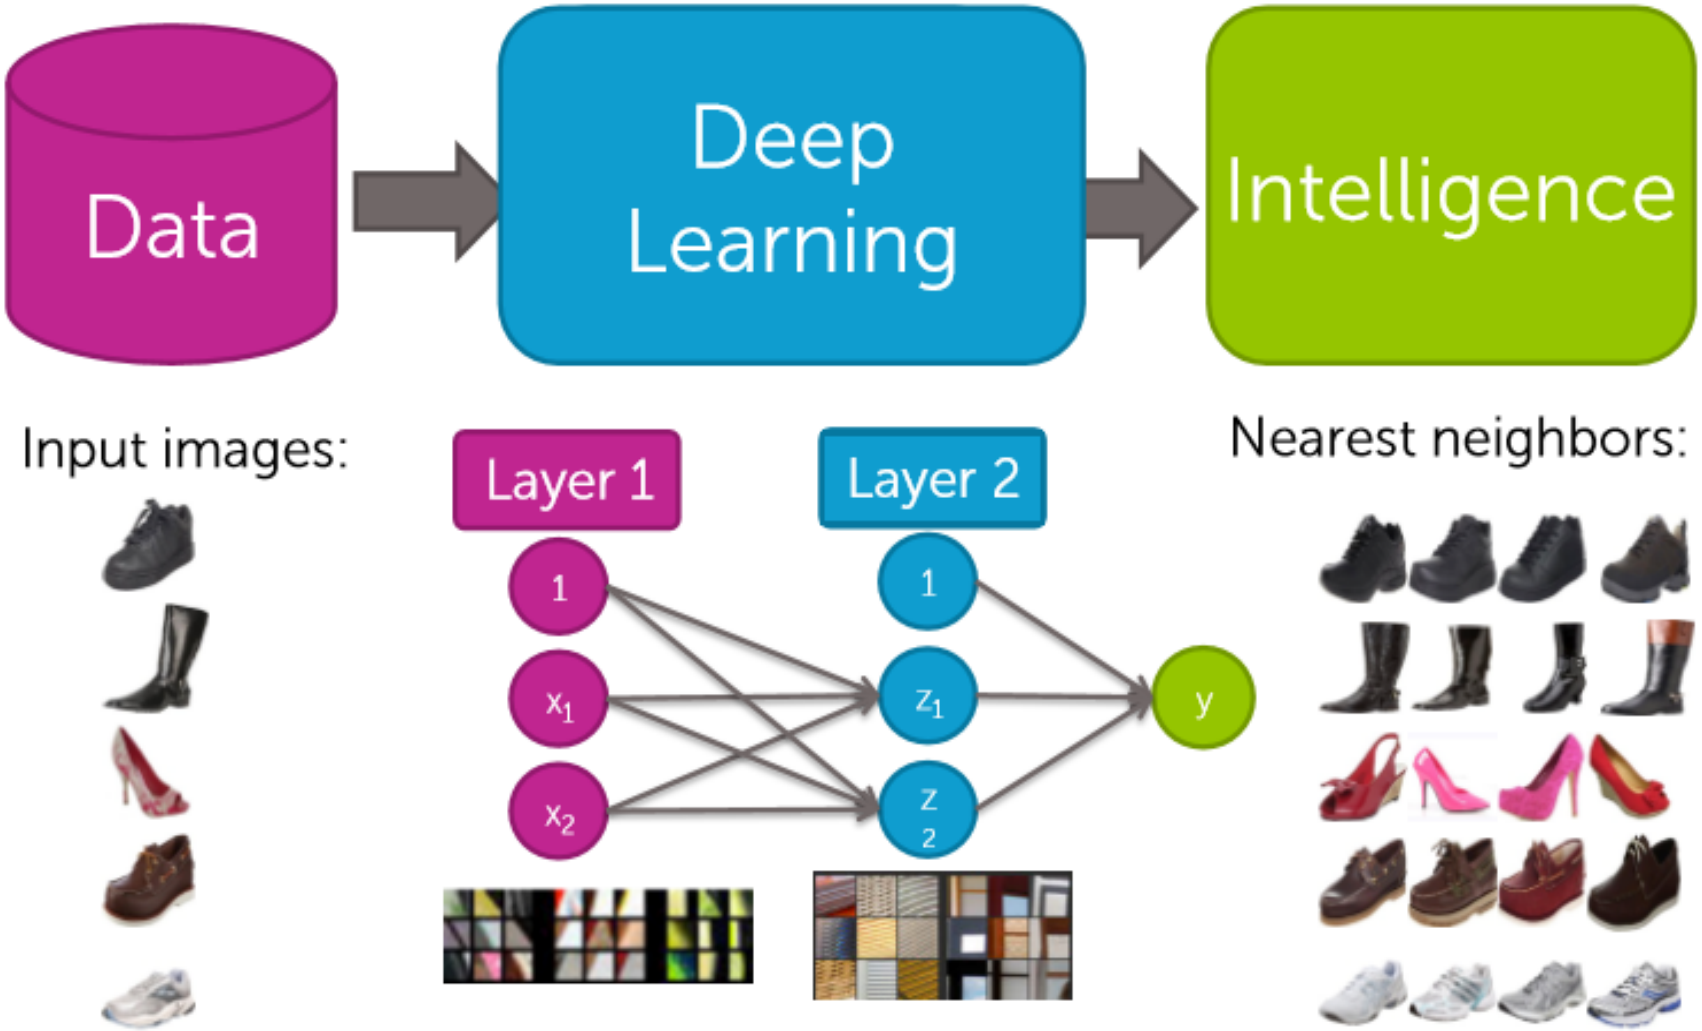

Document Retrieval

* Load data
* Compute TF-IDF for the corpus
  * graphlab.text_analytics.count_words
* Manually compute distances between people
  * Cosine similarity
* Build a knn_model nearest neighbor model for document retrieval
  * Knn_model = graphlab.nearest_neighbors.create()
* Apply nearest-neighbors model for retrieval
  * Query model to see who is closest
    * Knn_model.querry()

Nearest neighbor search

- Query article:

- Corpus:

- **Specify:** Distance metric
- **Output:** Set of most similar articles

Nearest neighbor search

- Query article:

- Corpus:

- **Specify:** Distance metric
- **Output:** Set of most similar articles

Song Recommender

* Create song recommender
  * Test data
  * Training data
* Build popularity_model
  * Based on graphlab.popularity_recommender
  * Same for all users
* Build song recommender personalized_model for personalization
  * Based on graphlab. item_similarity_recommender
  * Use get_similar_items()

Deep Learning

* Load the CIFAR-10 dataset
  * Popular benchmark dataset in computer vision
  * Takes about 5 to 10 minutes to download
    * Dataset is already split into training and test
* Train Knn_model via nearest-neighbor model for images using deep features
  * Knn_model = graphlab.nearest_neighbors.create(..)
* Find similar images
  * Knn_model.query(..)
* Show neighbors
  * Graphlab show_neighbors()

Image features

- Features = local detectors
  - Combined to make prediction
  - (in reality, features are more low-level)

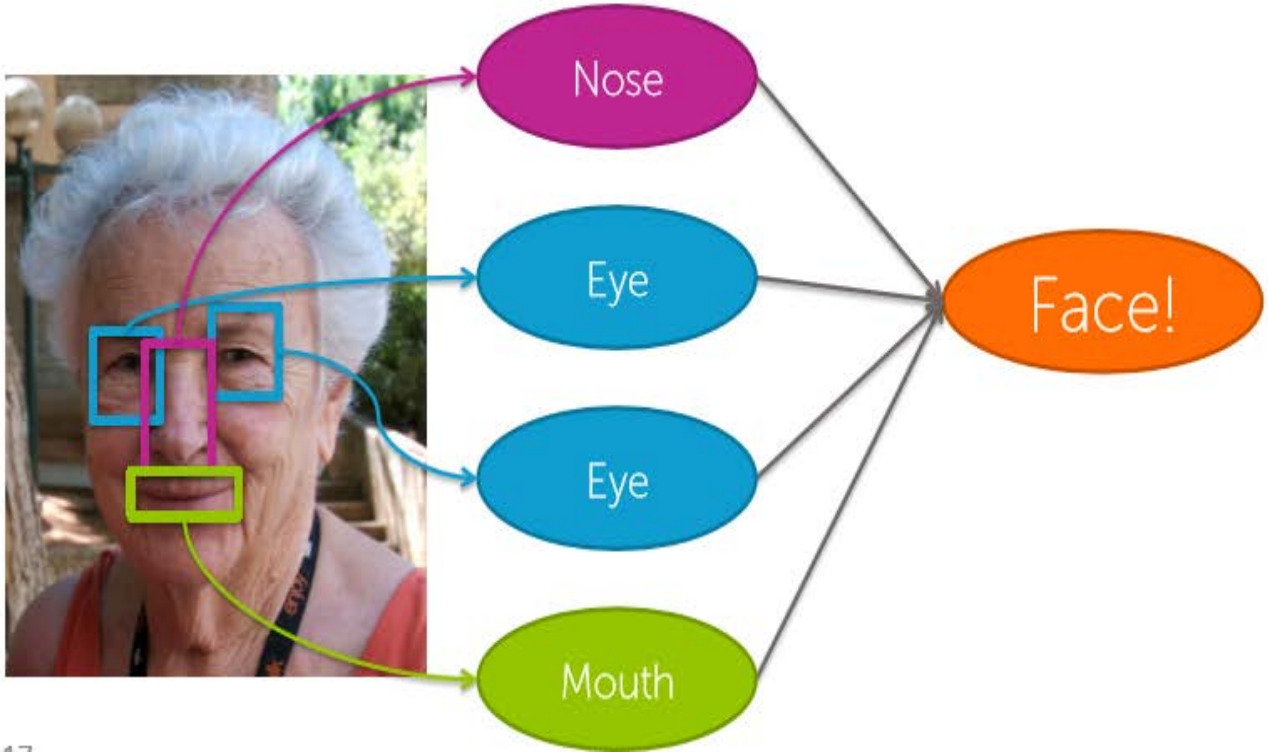

Typical local detectors look for locally "interesting points" in image

- *Image features*: collections of locally interesting points
  - Combined to build classifiers

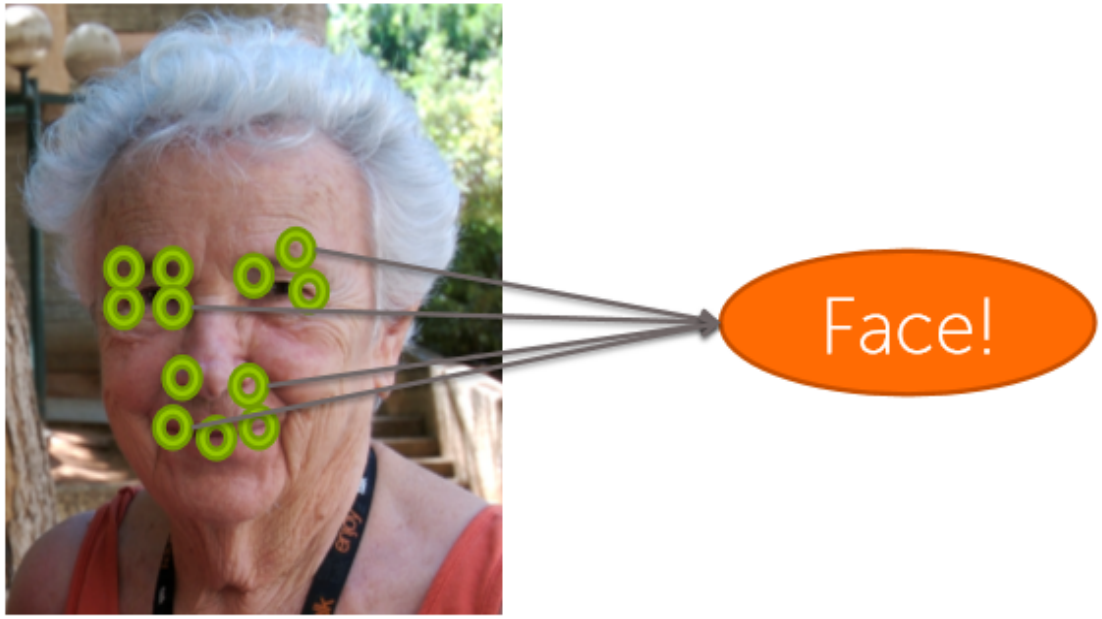

Many hand created features exist for finding interest points…

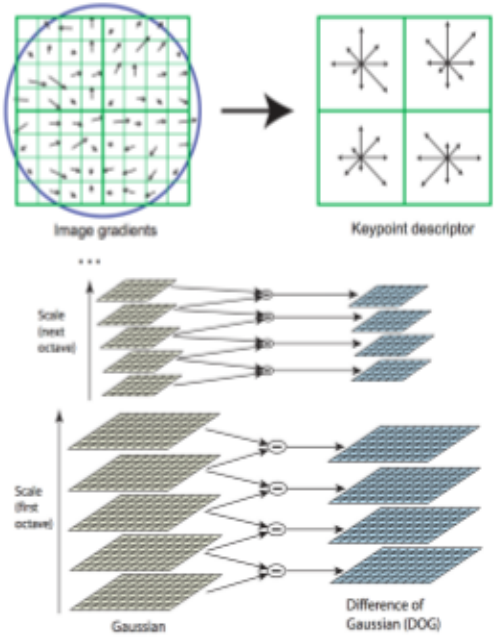

SIFT [Lowe '99]

- *Spin Images* [Johnson & Herbert '99]
- *Textons* [Malik et al. '99]
- *RIFT* [Lazebnik '04]
- *GLOH* [Mikolajczyk & Schmid '05]
- *HoG* [Dalal & Triggs '05]
- …

Deep learning score card

Pros

- Enables learning of features rather than hand tuning
- Impressive performance gains
  - Computer vision
  - Speech recognition
  - Some text analysis
- Potential for more impact

Deep learning workflow

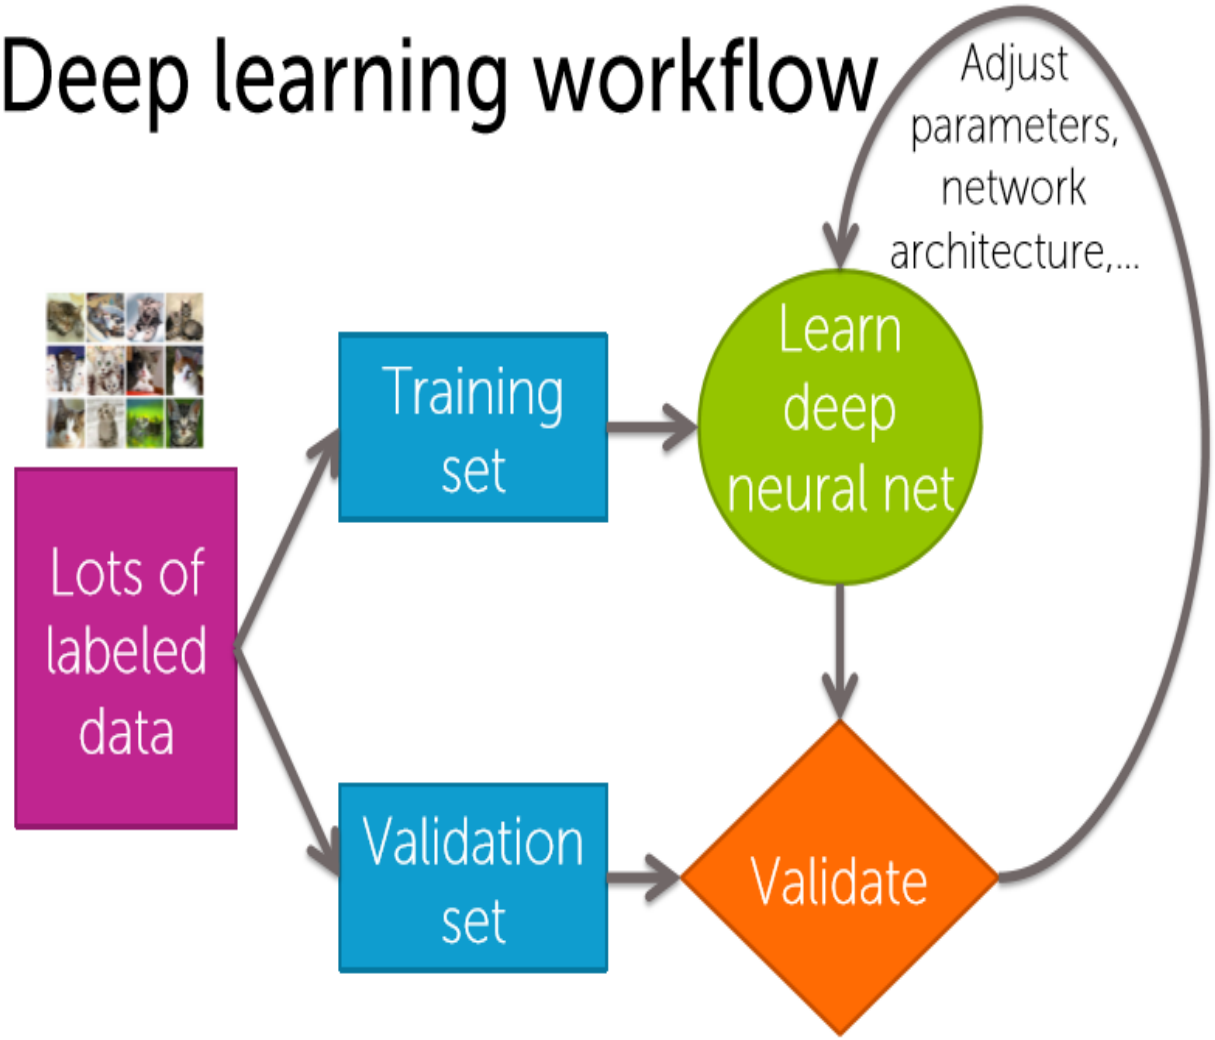

Deep learning score card

Pros

- Enables learning of features rather than hand tuning
- Impressive performance gains
  - Computer vision
  - Speech recognition
  - Some text analysis
- Potential for more impact

Cons

- Requires a lot of data for high accuracy
- Computationally really expensive
- Extremely hard to tune
  - Choice of architecture
  - Parameter types
  - Hyperparameters
  - Learning algorithm
  - …

Computational cost+ so many choices
=
incredibly hard to tune

April 25, 2017 | 7:30 AM – 3:45 PM
The Newseum, Washington, D.C.

Agenda

7:30 - 8:30 AM	**Registration & Networking Breakfast**
8:30 - 8:45 AM	**Welcome and Introductory Remarks** **William Sullivan**, Vice President and General Manager, North American Public Sector, Cloudera
8:45 - 9:00 AM	**Machine Learning, Artificial Intelligence and the Future of Big Data Analytics** **Mick Hollison**, Chief Marketing Officer, Cloudera
9:00 - 9:30 AM	**Health IT & the Role of Data as Fuel for Discovery** **Shawn Dolley**, Health & Life Science Big Data Expert, Cloudera **Matthew Quinn**, Senior Advisor for Health Information Technology, Health Resources and Services Administration
9:30 - 10:00 AM	**Networking Break**
10:00 - 10:45 AM	**Leading with Data for Government Results** **Vimesh Patel**, Director of the Office of Data Strategy and Innovation, National Counterterrorism Center **Adam Wilson**, Chief Executive Officer, Trifacta **Christine Kerns**, Regional Director, National Security Programs, Cloudera (Moderator)
10:45 - 11:15 AM	**Cloud Strategies that Work for Government** **Charles Zedlewski**, Senior Vice President, Products, Cloudera
11:15 AM - 12:00 PM	**Making the Most of Your Data—The Practical and the Possible** **Jessica Kahn**, Director, Data and Systems Group, Centers for Medicare and Medicaid Services, Department of Health and Human Services **Matt Keating**, Vice President of Customer Solutions, Zoomdata **Peter Kovac**, Manager of High Performance Computing and Data Engineering, Office of Financial Research, Department of the Treasury **Jeffrey Thomas**, Information Security Program Manager, Johns Hopkins Applied Physics Laboratory **Jim Pruden**, Senior Director, Cloudera Government Solutions (Moderator)
12:00 - 1:00 PM	**Networking Lunch**
1:05 - 2:00 PM	**Instrument Your Agency—The Potential for Connected Data** **Murthy Mathiprakasam**, Director, Product Marketing, Big Data, Informatica **Tracy Rausch**, Founder and Chief Executive Officer, DocBox **Matt Tarascio**, Director, Intelligence Technologies, Analytics, and Sustainment, Sikorsky, a Lockheed Martin Company **Dave Shuman**, Industry Leader, Retail, Manufacturing & IOT, Cloudera (Moderator)
2:05 - 3:00 PM	**Lower Government Risk and Address the Invisible Threats** **Joseph Mitchell**, Section Chief, Enterprise Data Analytics Section, Federal Bureau of Investigation **Dr. Ron Ross**, Fellow, National Institute of Standards and Technology **Dr. William Vanderlinde**, Chief Scientist, Intelligence Advanced Research Projects Activity **Dr. Melvin Greer**, Principal Engineer, Director, Data Science and Analytics, Intel (Moderator)

Questions

* Thanks!!!