IDC
Analyze the Future

UNIVERSITY OF OXFORD

Update: HPC and Industry

Steve Conway

HPC User Forum

University of Oxford

September 29-30, 2016

Market Growth

HPC Moves into Industry (Examples)

Government & Academic Research

Ford

Morgan Stanley

NOVARTIS

P&G

GEICO

Baidu

TOYOTA

DreamWorks Animation SKG

Walmart

facebook

BOEING

bp

SAMSUNG

PayPal

MERCK

PING

Genentech

AIRBUS

Panasonic

MAYO CLINIC

UNITED STATES POSTAL SERVICE

amazon

1970 1980 1990 2000 2010 Today

IDC Analyze the Future

IDC Studies of HPC in Industry Began 15 Years Ago…

Research in Progress

- **DOE:** 3-year global study on ROI related to HPC investments
- **NSF-NCSA:** global study on best practices in public-private HPC partnerships
- **Multi-client study:** annual global HPC/HPDA end-user study
- **Government:** in-depth study #2 on cyber security practices in the U.S. private sector
- **Government (multiple studies):** exascale initiatives around the world
- **Government:** study on algorithms used in high performance data analysis (HPDA)
- **Private:** HPC center user satisfaction and requirements
- **Private:** global study on the HPC market for cloud computing
- **Private:** global study on the HPC market for machine learning/ deep learning
- **Private:** global study on the HPC market for SMEs

2015 Book…

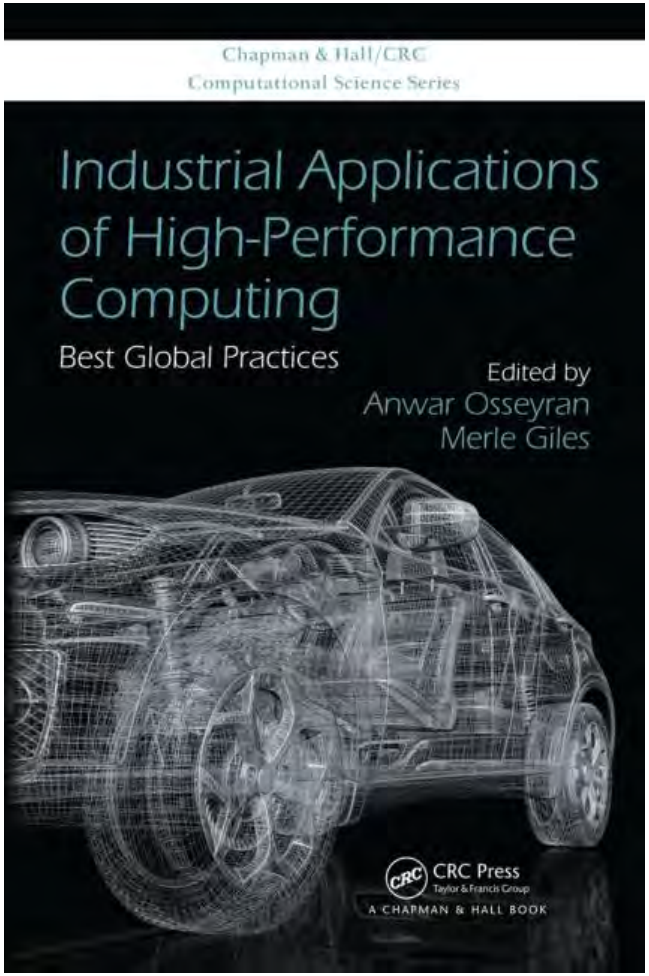

- **IDC contributed two chapters.**
  - Review of the global HPC community by country
  - Industrial use of high performance data analysis (HPDA)
- **We discussed the need for a study with NSF.**

Best Practices in HPC Public-Private Partnerships

New IDC Study for NSF/NCSA (2016)

IDC
Analyze the Future

Rationale for Global Study

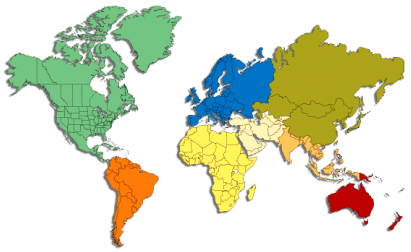

- **NSCI tasks NSF to help "maximize the benefits of HPC for economic competitiveness and scientific discovery.“**
  - NSCI Objective 5: …ensure that the benefits of the R&D advances are, to the greatest extent, shared between the United States Government and industrial and academic sectors."
  - “Industry” means primarily commercial users, not HPC vendors.
- **Across the world, governments are asking national HPC centers to start/operate public-private partnership programs.**
  - Success is not easy – many centers have little experience with industrial/commercial users.
- **IDC asked to collect and analyze existing PPP practices—what has worked well, what has not. NCSA will put this in a larger perspective.**
- **Report will be published October 2016 – openly available**

Interviewees and Questions

- National HPC centers worldwide (40)
- Industrial/commercial users of these centers (34)

- History of partnership
- Goals
- Funding
- Marketing (attracting industry)
- Satisfaction
- Recommendations

National HPC Centers: Types of Relationships with Industrial HPC Users

1: Science, No Industry

2: Science, PIs Bring Industry Collaborators

3. Industry Allocations for Open Science

4. Industry Allocations, Proprietary Work

IDC
Analyze the Future

HPC Applications Used in Manufacturing: Factors Limiting Performance

IDC Study for NSF/NCSA (2012)

IDC Analyze the Future

Executive Summary

- IDC conducted the study in 2012
- In-depth interviews + focus group
- 30 U.S. companies
- From 22 to 164,000 employees
- R&D expenditures: 3% to 45% of revenues.
- Average R&D investment 6.5 times more, as a % of revenues, than for all U.S. companies as a group

IDC
Analyze the Future

Goals of the NCSA / NSF Study

- Document HPC applications the manufacturers are using
- Determine the scalability of these applications (number of cores).
- Identify the barriers preventing greater scalability
- Determine what methods the manufacturers are using to overcome scalability limits
- Recommend areas where additional investment could improve the scalability of applications, especially where investment could be leveraged to benefit both scientific and industrial HPC users.

IDC
Analyze the Future

Research Methodology

- IDC conducted intensive interviews, in person or by phone, with individuals from 30 organizations who were selected for their known or assumed deep understanding of the attributes of key software applications in their areas of activity.
- These individuals came from a broad enough spectrum of industrial segments to be representative of the HPC-reliant industry sector as a whole.
- IDC convened a focus group in conjunction with our HPC User Forum held in Richmond, Virginia. Invited participants included noted experts in multiple fields of computational science and engineering.
- These individuals were also known to have a strong understanding of the relationships between their HPC applications and the science underpinning the applications.

Executive Summary

- 15.5% - current applications would meet their requirements for the next 5 years.
- 38.9% - underlying mathematical model/algorithm needs to be improved.
- 31.2% - underlying science needs to be improved.
- 30% of the organizations named specific ways in which they now have to "dumb down" their problems in order to complete the runs in reasonable amounts of time
- 63% of the responses favored a high degree of involvement with NCSA/NSF in pursuing needed improvements.

IDC
Analyze the Future

Executive Summary

30% of the companies had to “dumb down” their problems to finish in reasonable times:

- Use coarser meshes
- Employ fewer elements
- Not fully exploit the known science
- Employ fewer time steps or reduce the length of the investigated timeframe
- Run more costly physical experiments

IDC
Analyze the Future

High Performance Data Analysis (HPDA) Market

IDC Global Forecast (2015-2020)

HPDA = Data-Intensive Computing Using HPC

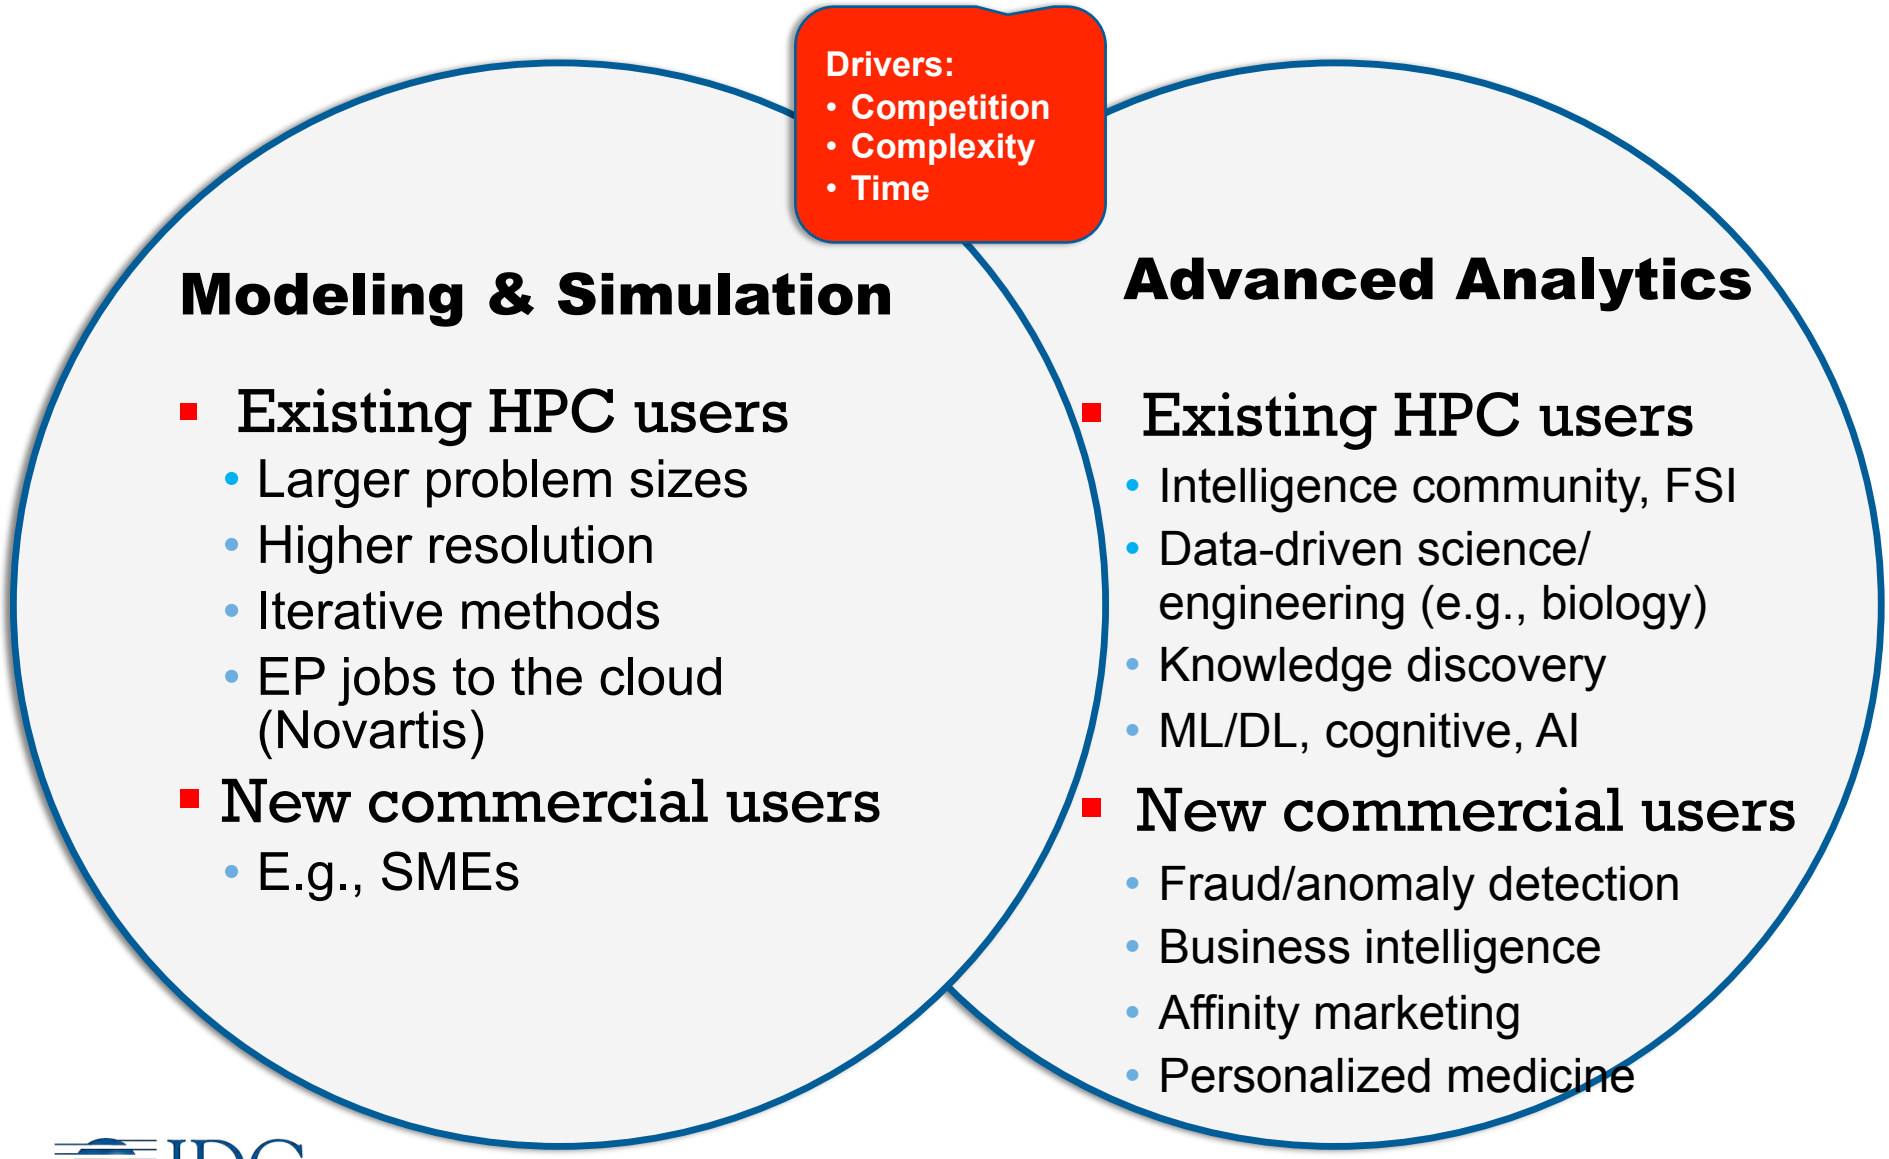

HPDA Includes Cumulative Results of Iterative Methods

- Parametric modeling (product design)
- Stochastic modeling (financial)
- Ensemble modeling (weather/climate)

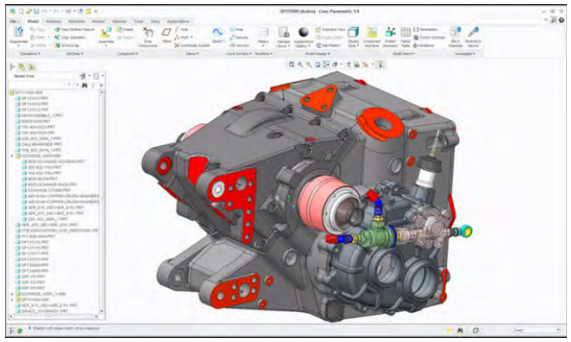

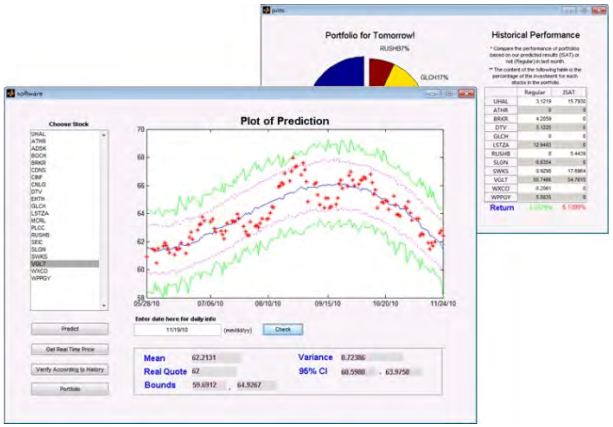

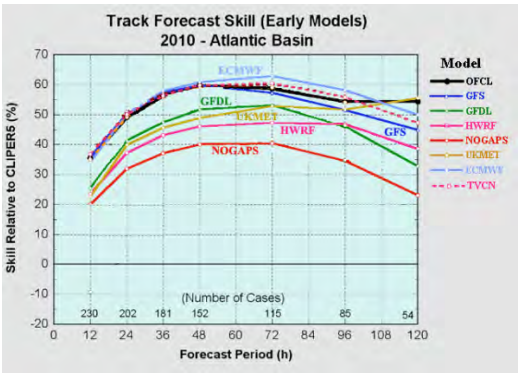

IDC
Analyze the Future

HPDA Server Market Forecast to 2020

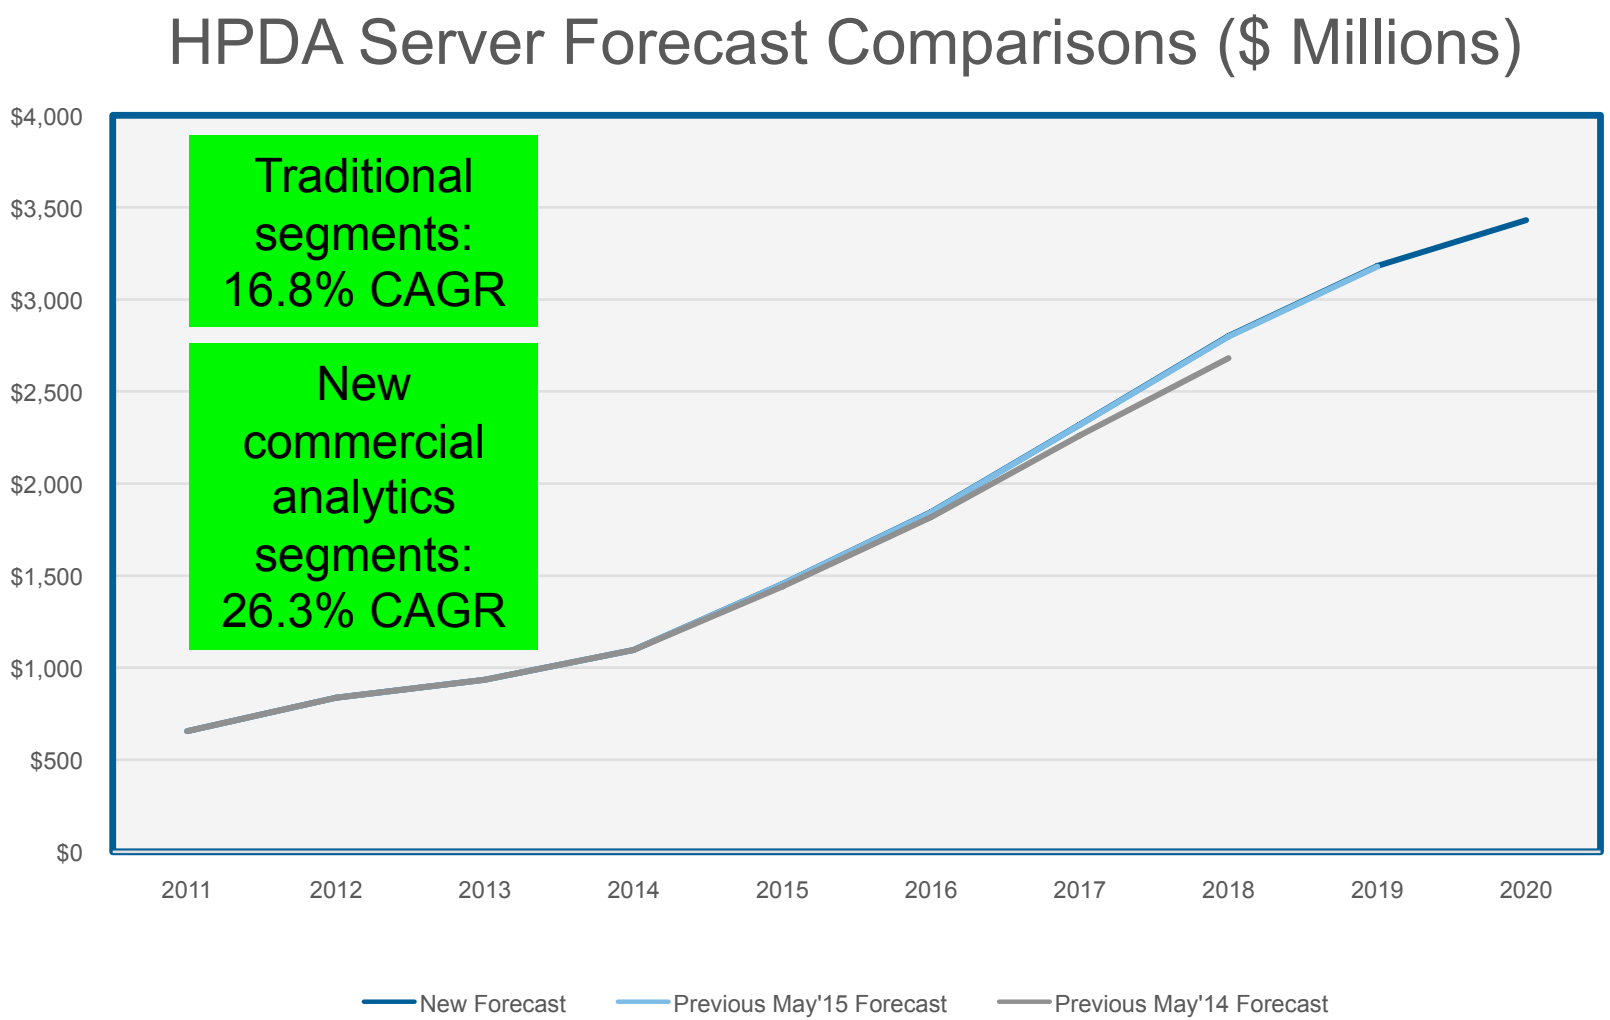

HPDA As a % of Total HPC Utilization (All Systems)

Of all the workloads on your systems, what percentage of your total system utilization is devoted to Big Data analytics (MapReduce/Hadoop, graph analytics, semantic analysis, etc.)?			
	Percentage of System in Government	Percentage of System in Academia	Percentage of System in Industry
Currently used for Big Data analytics - Percent	18.2%	8.4%	32.6%
Expected 6 to 18 months from now - Percent	19.8%	15.6%	33.4%
N = 128			

Source: IDC, 2015

Where HPDA Workloads Are Run

Type System Used for HPDA Workloads

	Number of Responses	Percentage of Respondents
We currently run these types of workloads on the same HPC system used for other types of workloads and we plan to continue as is for at least the next 6-18 months	56	52.8%
We currently run these types of workloads on the same HPC system used for other types of workloads but we plan to migrate them over to a separate HPC system or Big Data appliance within the next 6-18 months	25	23.6%
We currently run these types of workloads on a separate HPC system or Big Data appliance and we plan to continue as is for at least the next 6-18 months	26	24.5%
We currently run these types of workloads on a separate HPC system or Big Data appliance but plan to migrate them over to the same HPC system used for other types of workloads within the next 6-18 months	3	2.8%
Other	3	2.8%
Total	113	
N = 106		

Source: IDC, 2015

IDC Analyze the Future

Leverages HPC and enterprise architectures to deliver HPDA for making real time business decisions.

Fraud detection in real time…

Hybrid HPC and enterprise systems provide fraud analytics for safeguarding millions of transactions per day in real time. High volume, highly scalable and highly available.

- 12 Million+ logins / day
- 15 Million+ transactions /day

Operations analytics in real time…

Real time anomaly detection in correlated event streams minimize operational downtime and maximize customer satisfaction by using predictive analytics.

- 3 Million events / second
- 25 Tb of data ingested / hour

IDC
Analyze the Future

Processing in Memory

- **USPS reports processing very large data sets in memory can boost performance up to 6 orders of magnitude:**
- **Processing outside of memory can cut performance:**
  - ~50% moving data from one blade to neighboring blade
  - ~67% moving data from top blade to middle blade in a rack
  - 25x moving data from top blade to bottom blade in a rack
  - Much more when moving data between racks

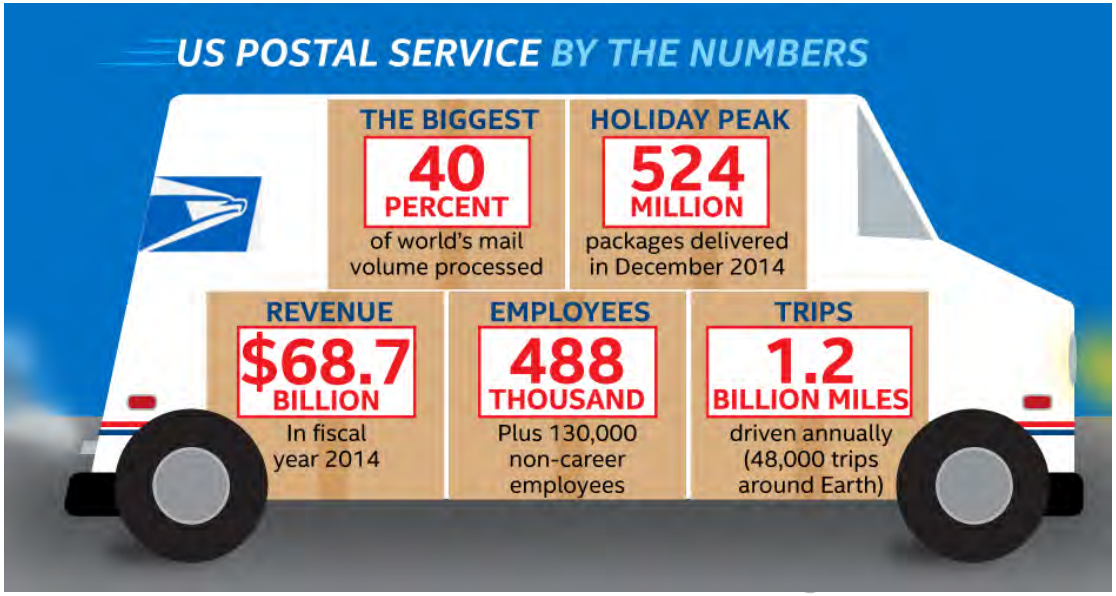

HPDA Application

Healthcare Innovations

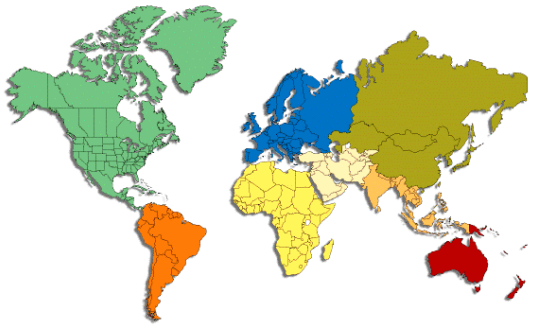

- Hangzhou CognitiveCare-IBM Watson (China): cancer (21 hospitals)
- Childrens Mercy Hospital (U.S.): rare childhood diseases
- Mannheim Frauenklinik (Germany): childbirth
- Frédéric Joliot Hospital (France): PET scan tracer optimization
- University of Oslo (Norway): bowel and prostate cancer
- Lobachevsky State University of Nizhni Novgorod (Russia): 3D virtual clones of patients
- University of Toronto SickKids Center (Canada): genomic medicine
- Neuroblastoma and Medulloblastoma Translational Research Consortium, Michigan (U.S.)
- University of Rochester Medical Center (U.S.): infant seizures
- Ariana Pharma (France): long-term cancer case management
- Victor Chang Cardiac Research Institute (Australia): heart arhythmias

HPDA Application:
Smart City

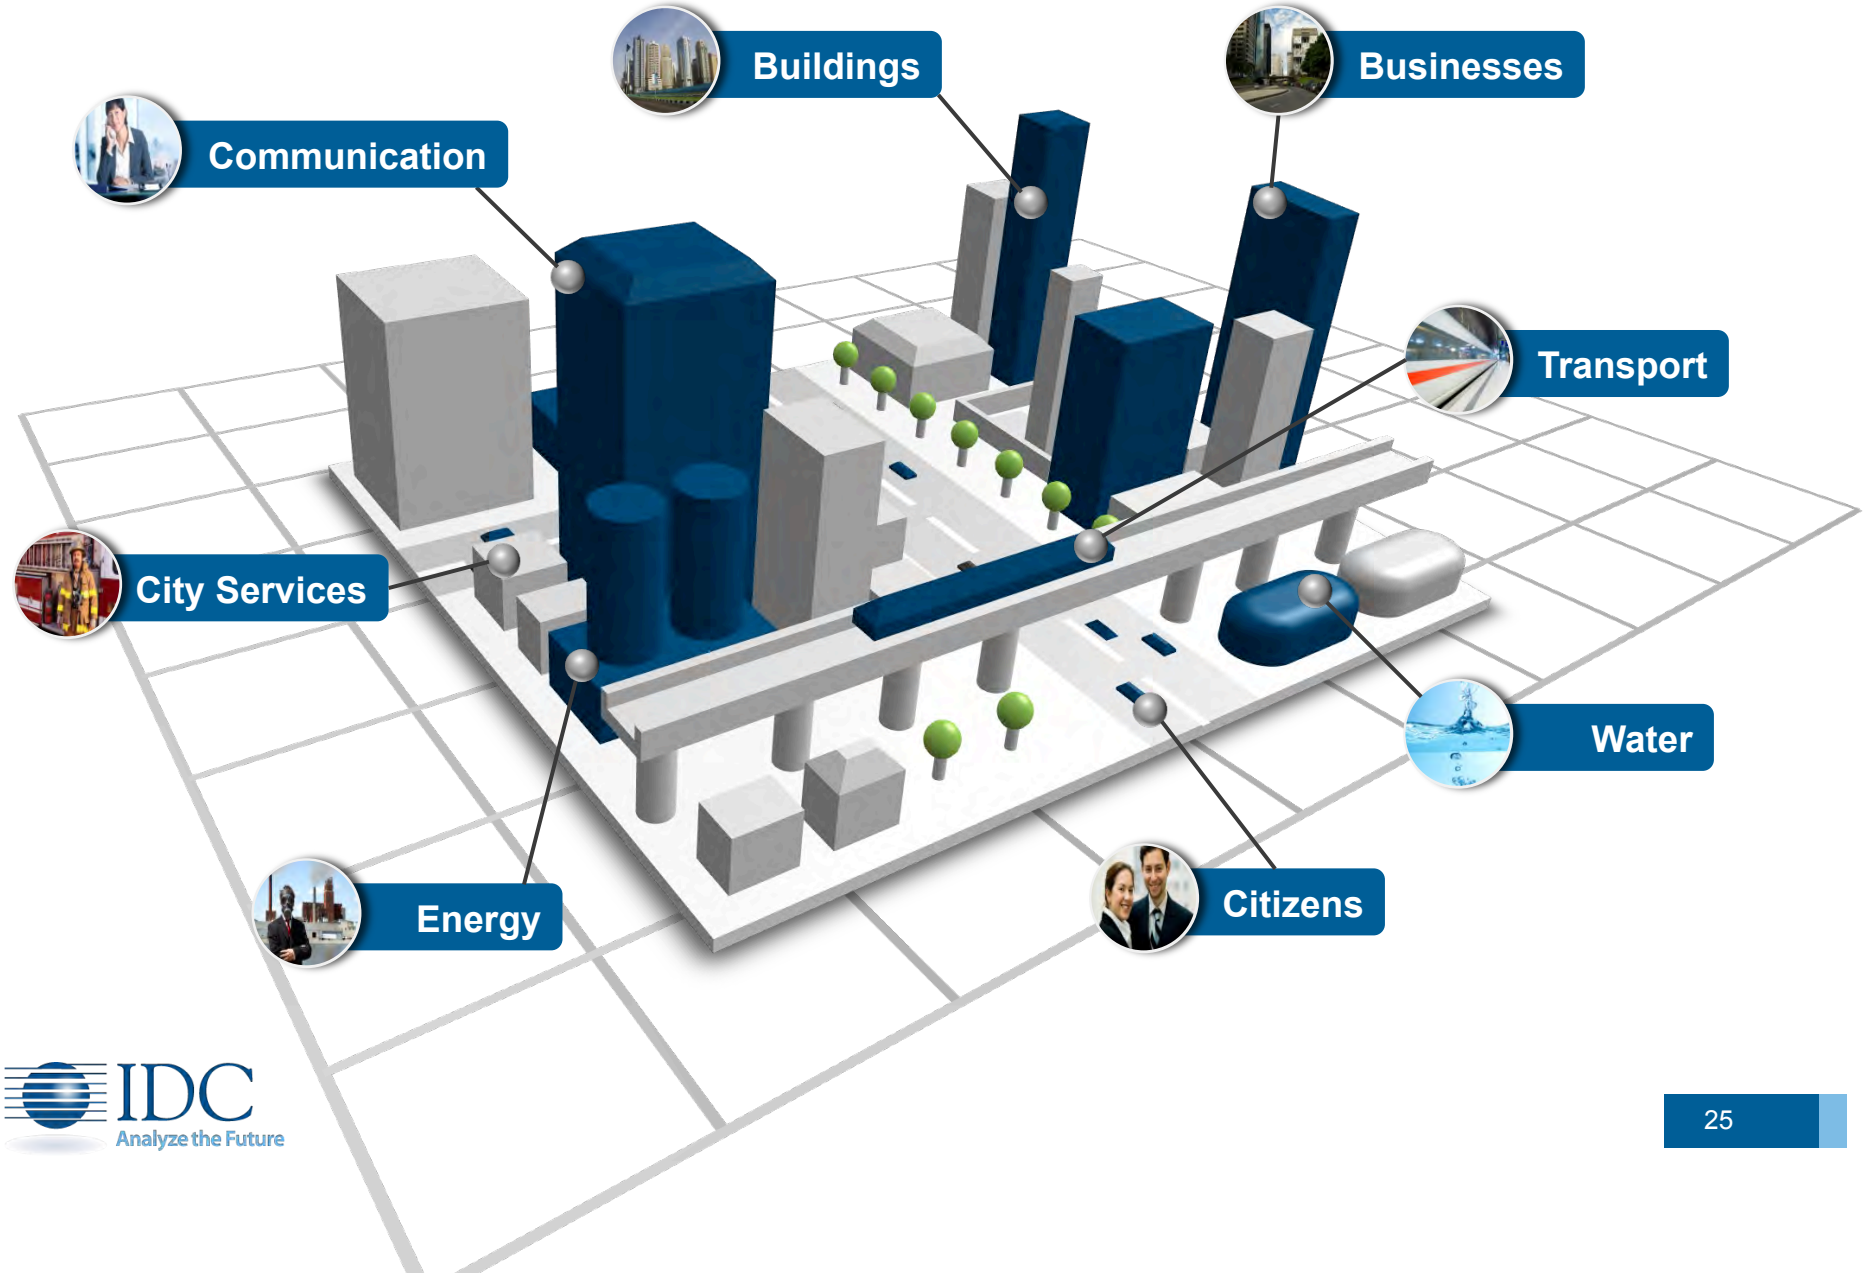

IDC
Analyze the Future

HPDA Application:

Internet of Things (IoT)

- IoT Network Management
  - Wellness
  - Security
  - China's IoT plan
- Dense local nodes
  - Driverless traffic
  - Smart power grid
  - IoT edge computing