

The New World of Data

IOT + Big Data + Cloud

Denny Lee
Principal Program Manager Lead
SQL Customer Advisory Team



2008

From Database to Complete Data Platform
The Evolving Field of Data Management

2009

SQL Server 2008
Steps Towards a Complete Data Platform

2010

Microsoft ♦ Cloud
Delivering IT as a Service

2012

The New World of Data
IOT + Big Data + Cloud

Agenda

Scenarios

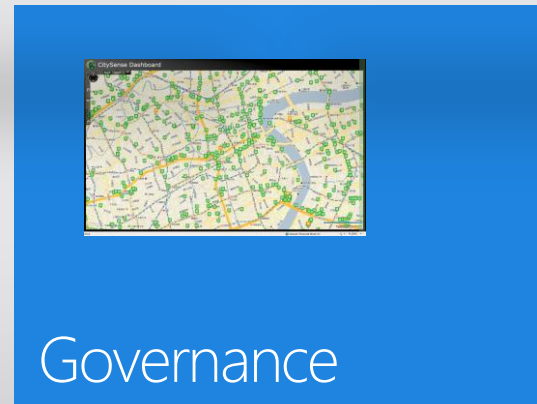

Internet of Things (IOT)



Big Data

IOT +
Big Data +
Cloud

Scenarios: How Data is Changing the World



Business

▶ Supply Chain

Tracking every package all the way from source to destination

Ensuring there is no tampering anywhere along the supplying chain.

▶ Manufacturing

Tracking machinery to predict / prevent failures


Quality control

▶ Utilities

Tracking demand at a fine-grain level to dynamically balance load


Enable differential pricing at different time of day

Enable prediction of power usage patterns




Finance

- ▶ Detect suspicious stock trading patterns (e.g. insider trading, market manipulation, ...)
- ▶ Trace cross-border currency movement in violation of money laundering / anti-terror laws
- ▶ Technical analysis to detect shifting correlations for algorithmic trading ("risk on" / "risk off"; USD / Euro / gold / stocks)



Science

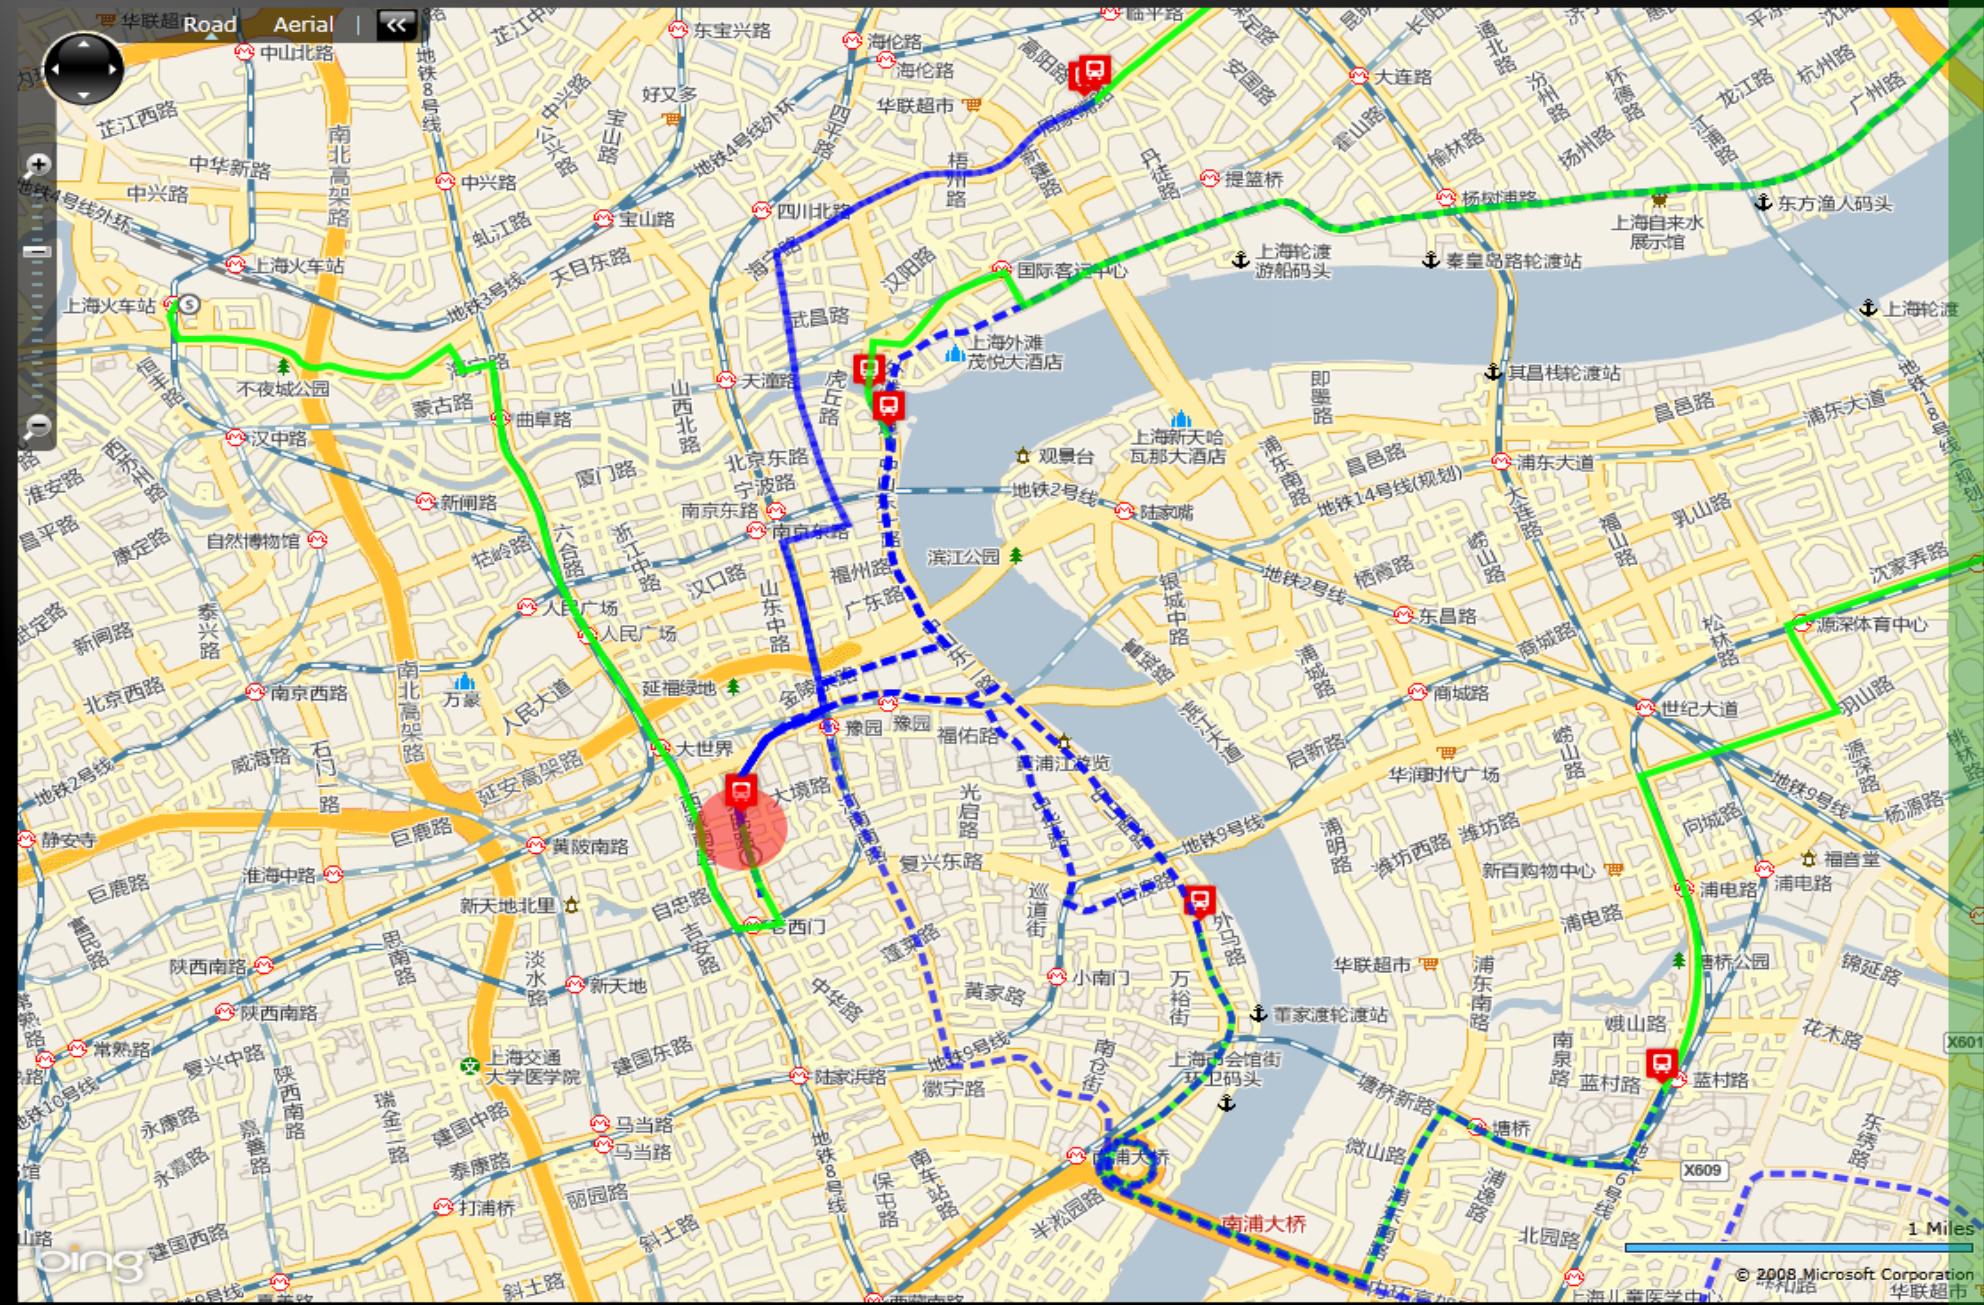
- ▶ Measuring, Visualizing & Analyzing Natural Phenomena
 - ▶ Precipitation, Earthquakes, ...



City Governance

- ▶ Incorporate real-time traffic information into navigation systems
- ▶ Optimize traffic planning
- ▶ Implement an optimized emergency-response system
- ▶ Provide secure surveillance / anti-terrorism system







Overview



Delay 57



Accident 20



Out of route

! Accident Info

- > Accident27 (54) impacts
- > Accident61 (6) impacts
- > Accident73 (11) impacts
- > Accident82
- > Accident96 (5) impacts
- > Accident100 (7) impacts
- > Accident102 (7) impacts
- > Accident105 (2) impacts
- > Accident106 (7) impacts
- > Accident107 (14) impacts
- > Accident109 (6) impacts
- > Accident110 (2) impacts
- > Accident111 (5) impacts
- > Accident112 (3) impacts
- > Accident113 (7) impacts
- > Accident114 (9) impacts
- > Accident115 (4) impacts
- > Accident116 (3) impacts
- > Accident117 (9) impacts
- > Accident118

Sustainability

- ▶ Monitoring Air, Water and Noise Pollution: <http://eyeonearth.cloudapp.net/>

Built on Windows® Azure™

English ▼ [about](#) | [providers](#) | [disclaimer](#)

EYE ON EARTH

Search  

STATIONS
• water stations

European Environment Agency 

WaterWatch  

500 miles 1000 km

© 2012 Microsoft Corporation Image courtesy of NASA

Internet of Things (IoT)

Wireless network technology advances

Standardized Comm. Protocol (e.g. IP)

Small, cheap capable, low power chips

Massive storage and compute through cloud


Internet of Things




cloud computing



communications



sensors



actuators

The basic idea of the IOT is that virtually every physical thing in this world can ... feature tiny computers that are connected to the Internet (ITU, 2005). When they do so, they are often called smart things, because they can act smarter than things that have not been tagged

Prof. Elgar Fleisch, ETH, Zurich



Reduce operation risk & diminish cost



Process automation











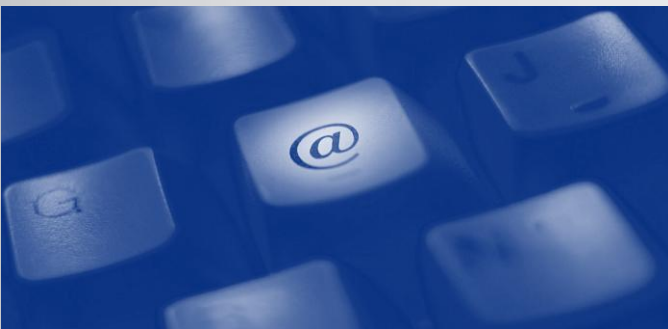



Process optimization



New way to create value & new business model



Internet vs. Internet of Things

<h2>Internet of things</h2>	 <p>Invisible devices</p>	<p>Trillions of computer-enabled devices which are part of the IoT</p> <p>Trillions of networked nodes</p>	 <p>100kBit/sec</p> <p>Low bandwidth last-mile connection</p>		
 <p>Mostly addressed by local schemes</p>	 <p>Machine-centric</p>	 <p>Sensing-focus</p>	 <p>Laptops / tablets / smartphones</p>	<p>6+ billion people 1.5 billion use net US: 4.3 devices per adult</p> <p>Billions of networked devices</p>	 <p>Cable: 1MBs+ Fiber: 50-100Mbs</p> <p>High-bandwidth access</p>
	 <p>Global addressing</p>	 <p>User-centric</p>	 <p>Communication-focus</p>	<h2>Internet</h2>	

Microsoft Offerings for Embedded Devices

Windows Embedde d Standard

Componentized Windows

- GB-sized footprint
- X86, x64
- Compatibility with existing Windows apps and drivers

Windows Embedde d Compact

Evolution of Windows CE

- Real-time OS
- MBs-sized footprint
- ARM, MIPS, x86

.NET Micro- framewor k

Adaptation of .NET Framework for resource-constrained environments

- ~100Ks memory, no external MMU
- Integrated VS developer environment
- 32-bit processors

Windows Embedded

IoT and Information Analysis



Tracking Behavior

Example: presence-based advertising and coupons based on consumer's location



Enhanced Situation Awareness

Example: intruder detection combining video, audio, vibration and other technologies into a modern surveillance system



Sensor-driven Decision Analytics

Example: continuous monitoring of chronic diseases to help doctors determine best treatments

Putting IoT to Work: Automation & Control



Process Automation

Example: Continuous, precise adjustments in manufacturing lines



Optimize Resource Consumption

Example: Datacenter management to optimize energy, storage and processor utilization,



Complex Autonomous System

Example: Collision avoidance system to sense objects and automatically apply breaks

Big Data




New Data Sources



Large Data Volumes



New Technologies



Non-traditional Data Types

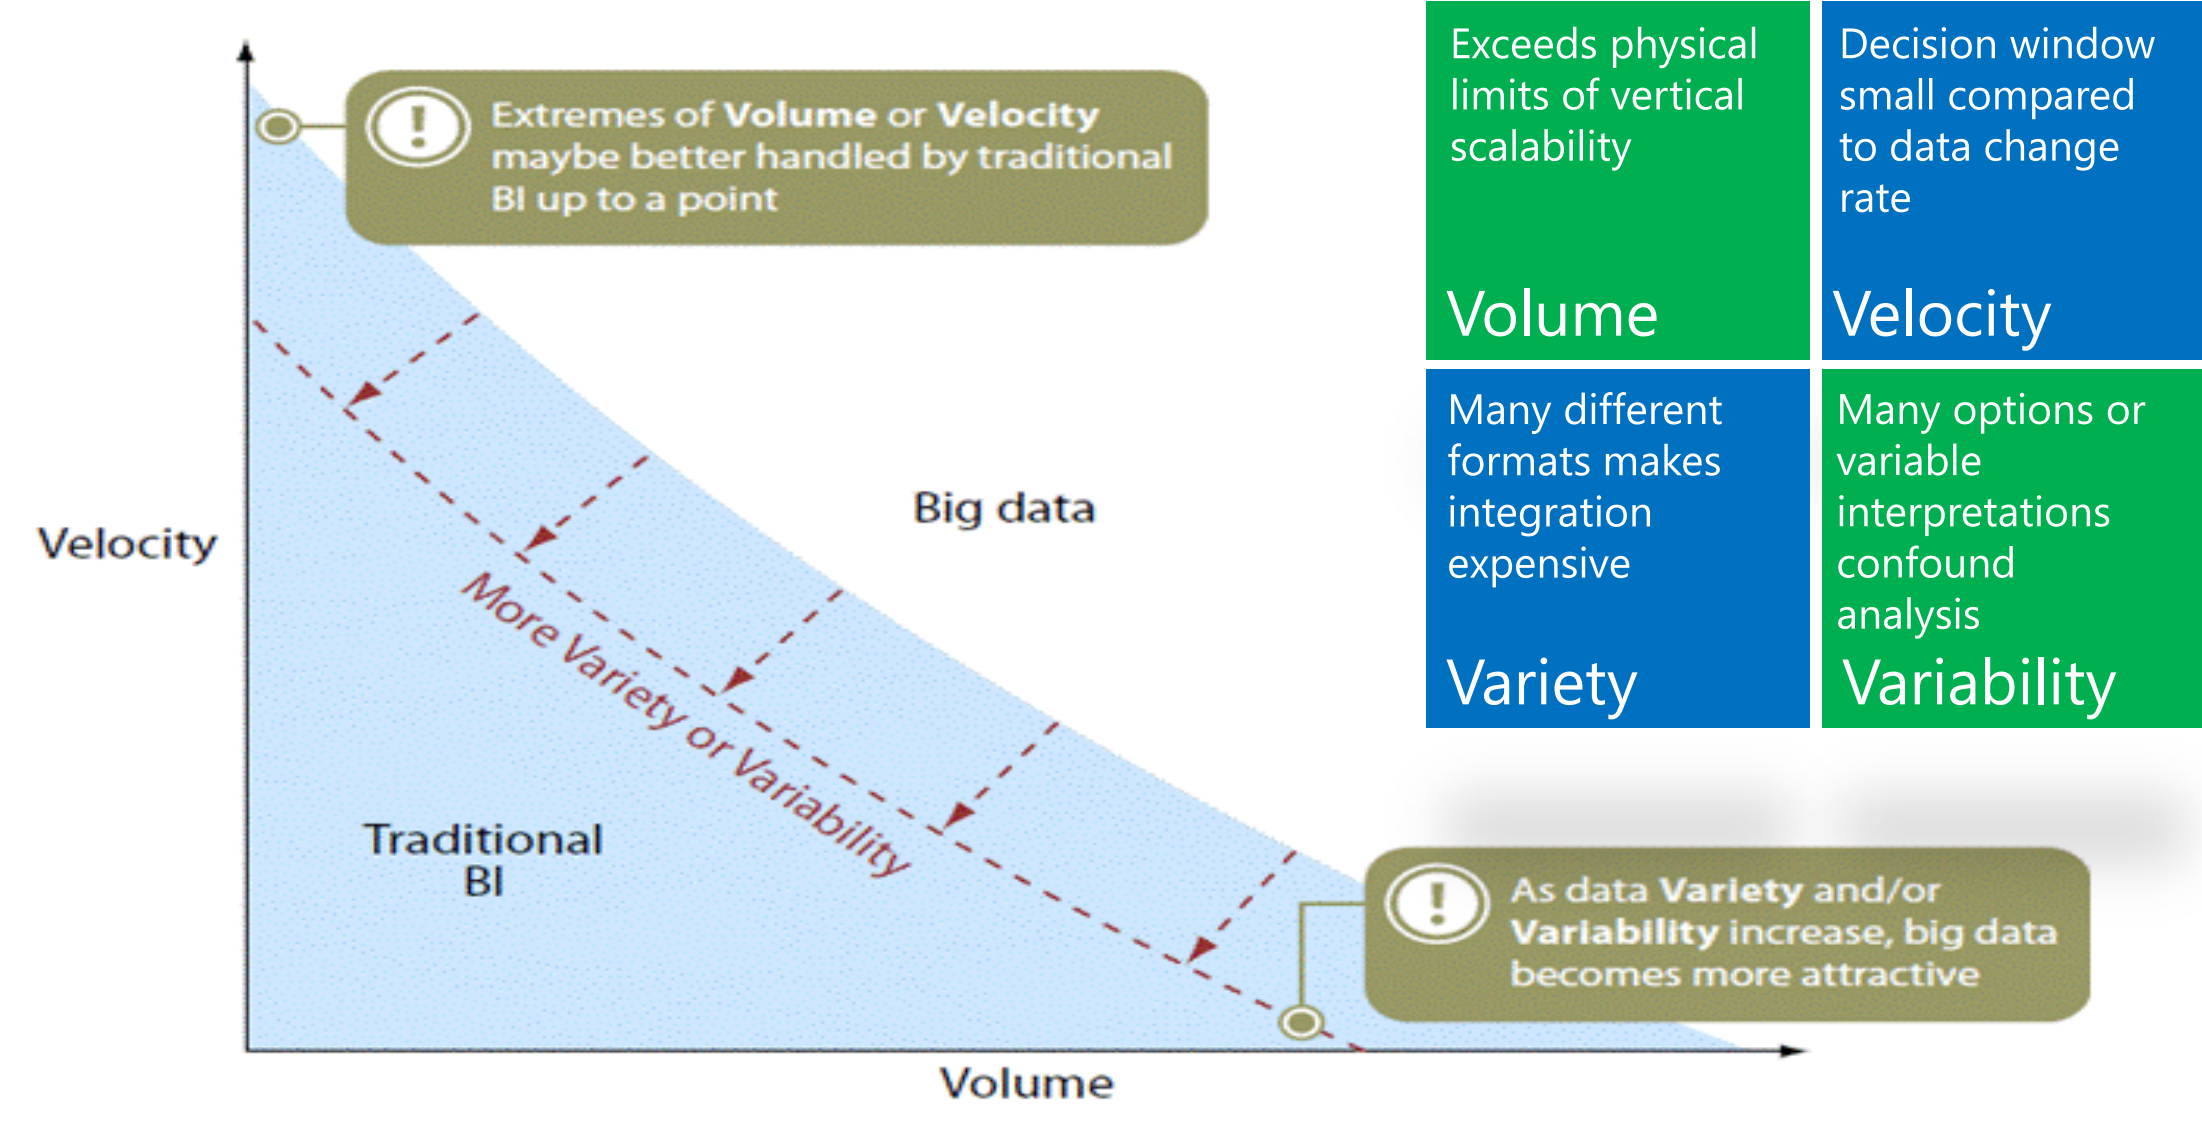


New Economics

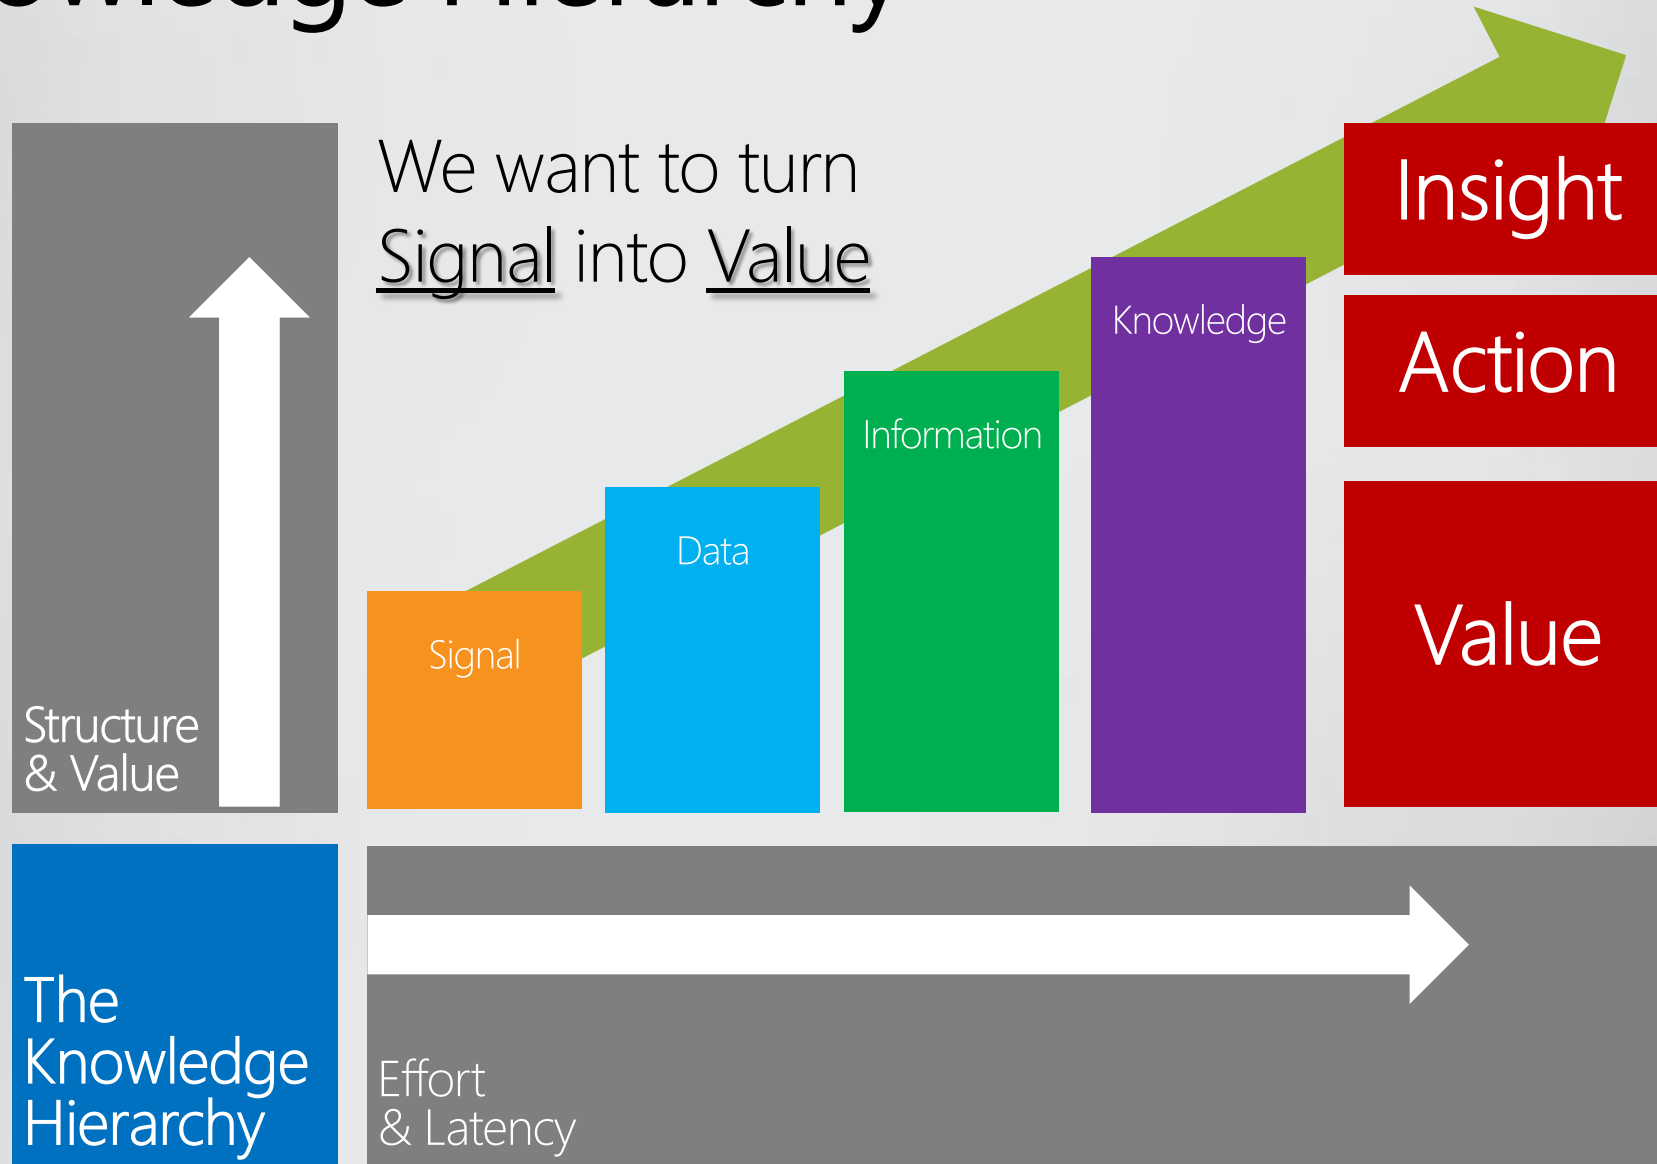
New Questions & New Insights




Big Data: 4 Vs



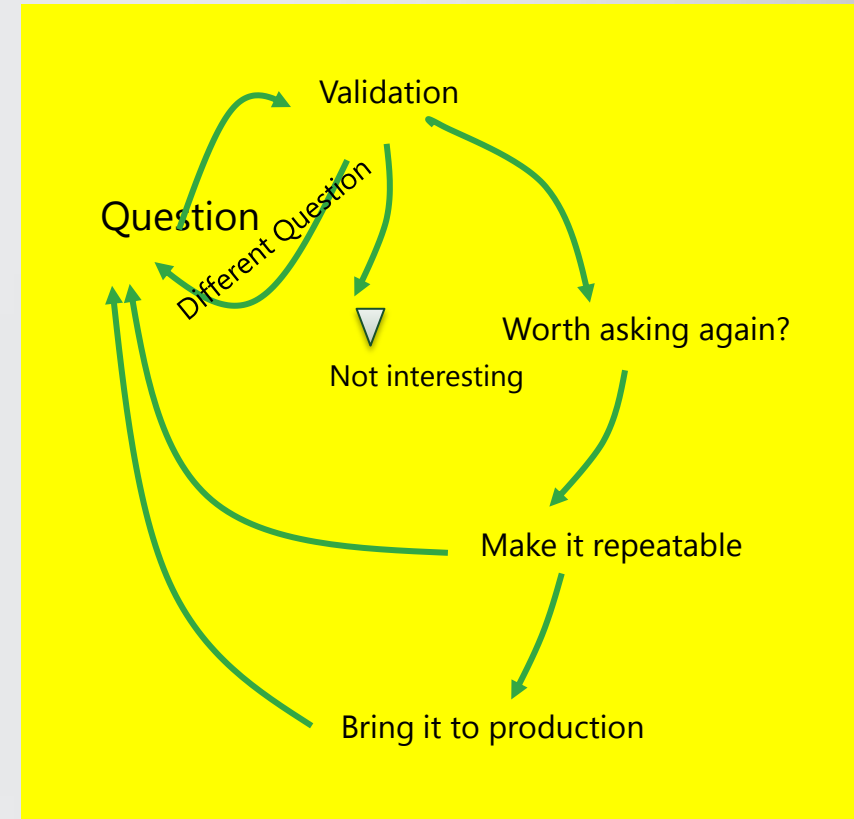
The Knowledge Hierarchy



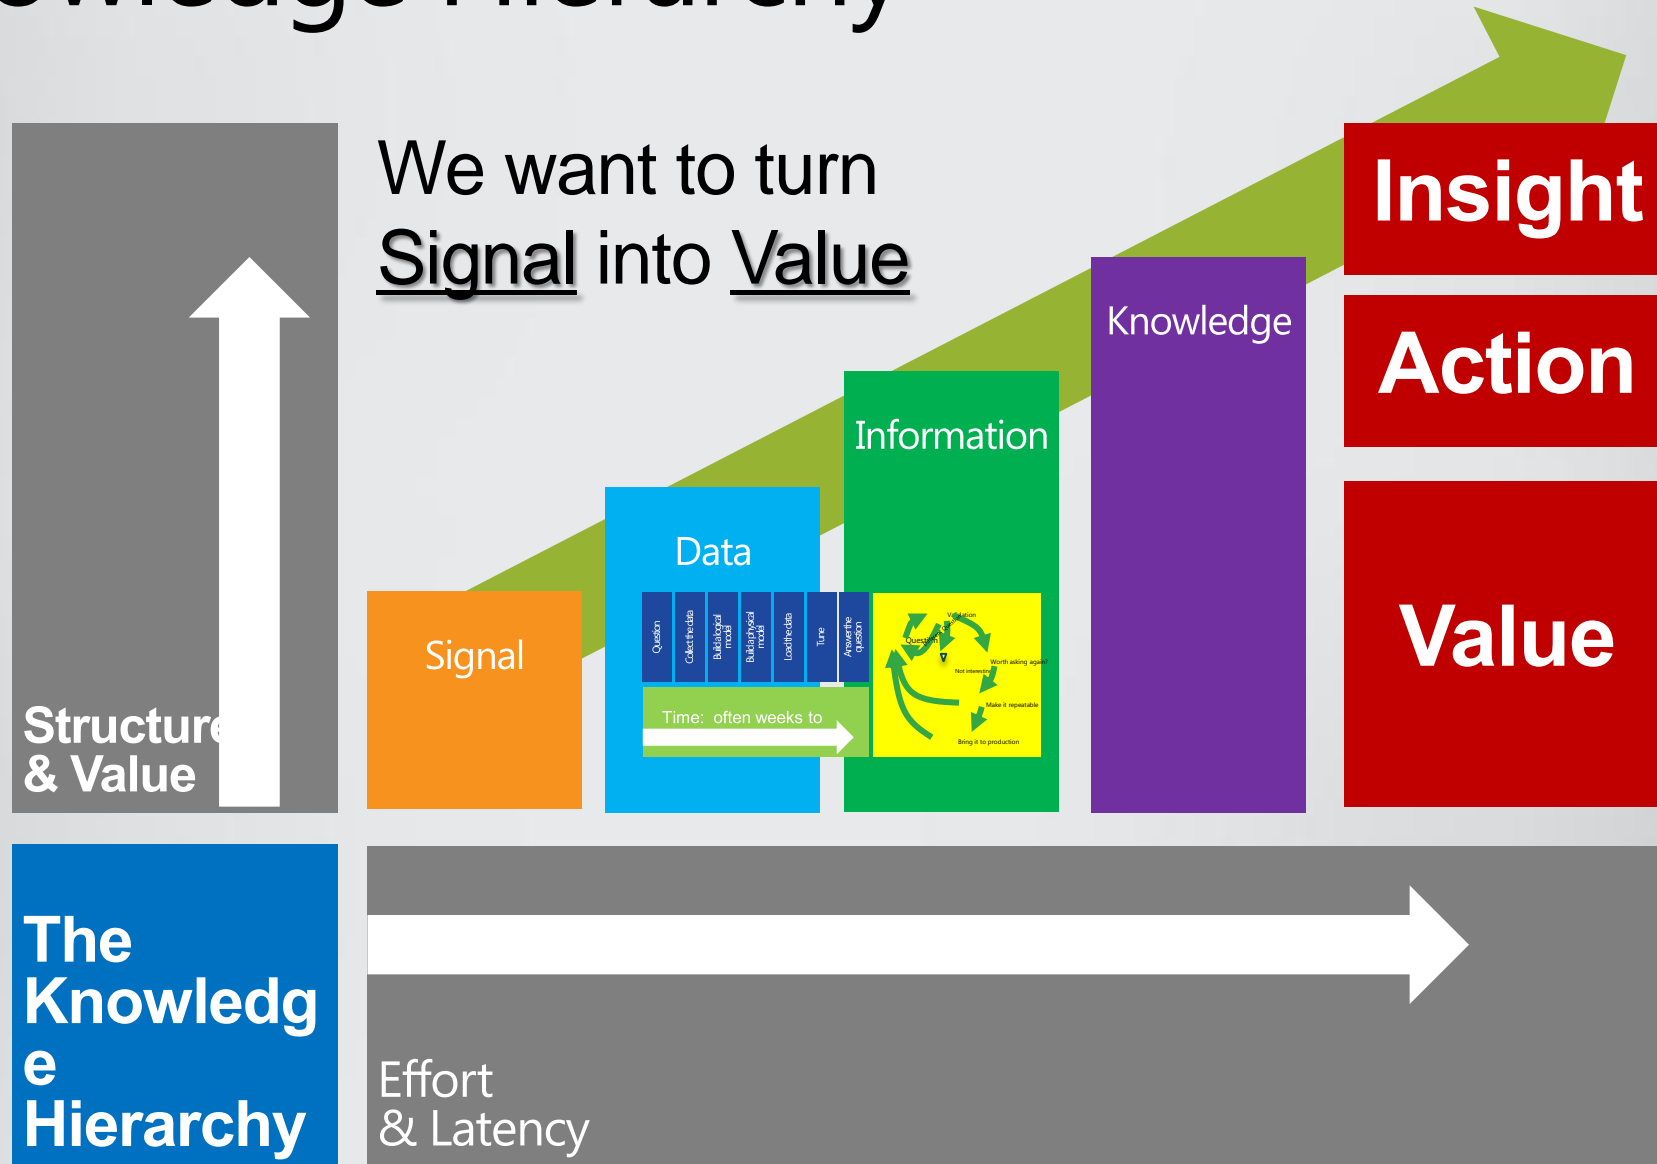
The Knowledge Hierarchy



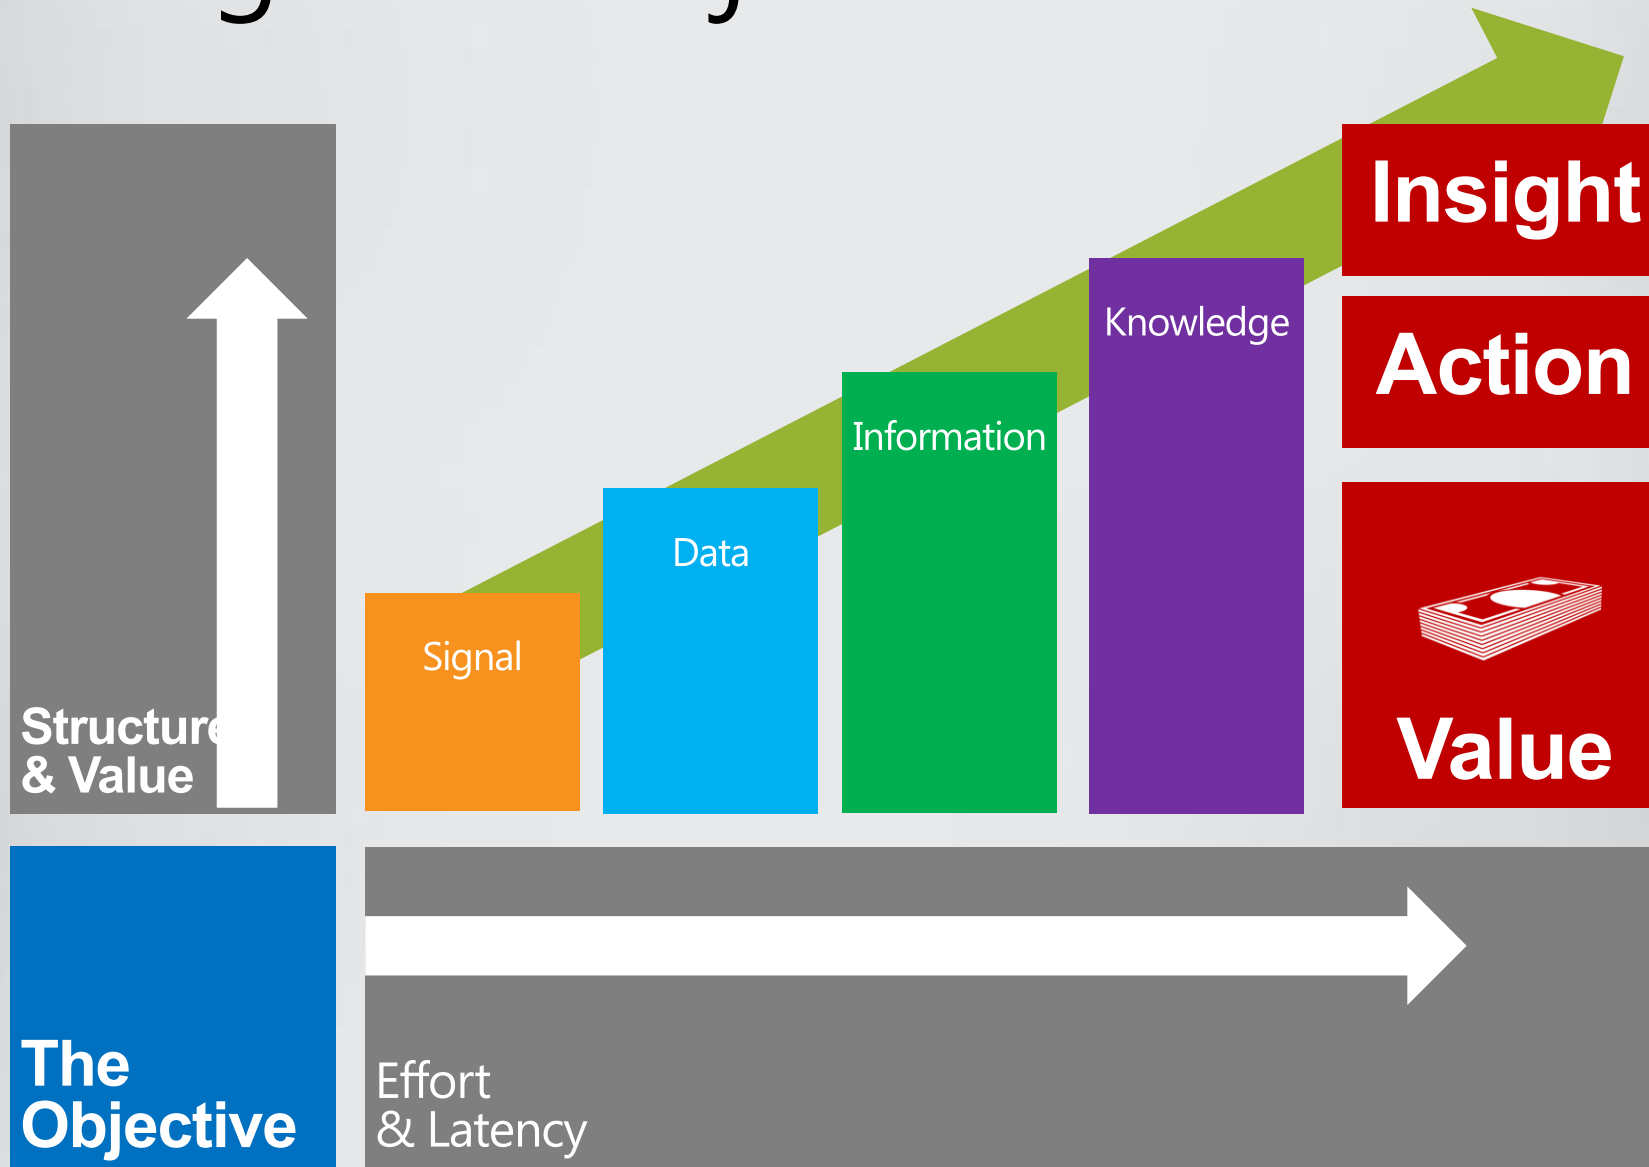
Time: often weeks to months




The Knowledge Hierarchy



The True Big Data Objective



The True Big Data Objective

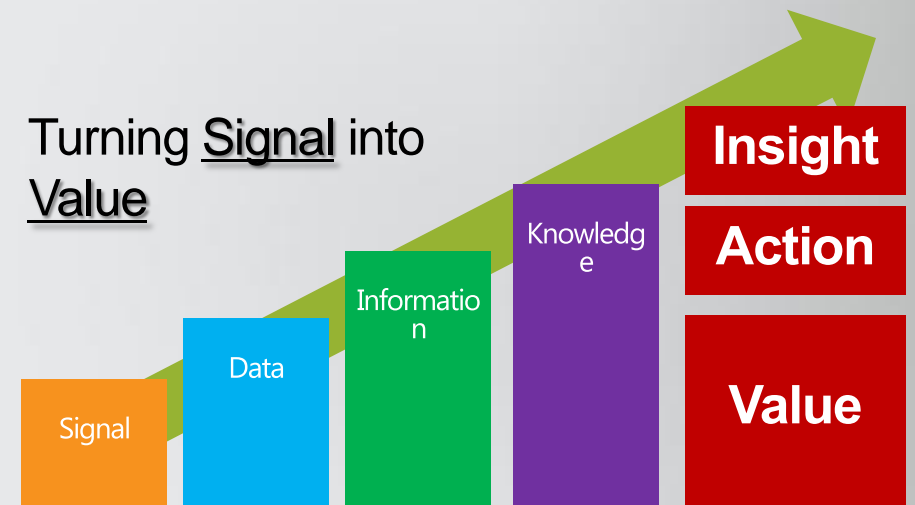


The "Big Data" Approach

Big Data processes both remake, and complement existing analytic workflows by

- ▶ Simplifying production of structured information from emerging "ambient" data sources
- ▶ Enabling rapid sense-making over un-enriched and un-modeled data
- ▶ Enabling analytics at scale over ambient data
- ▶ Enabling creation of ambient data driven models

- Existing systems enable sense-making over modeled data
- There is tremendous potential value in making sense of ambient data



Putting it Together: IOT, Big Data, Cloud



Internet
of
Things

Big
Data

Cloud Optimize Every Business

Public

 Windows Azure™ Platform

**Common
Technologies**


Identity ▪ Virtualization ▪ Management ▪ Development

Private

 Windows Server™
with Hyper-V™

 Microsoft®
System Center

Big Data Platform in Windows Azure


Microsoft
Distribution
of Hadoop

HPC

Data
Mining /
Predictive
Analytics

Data
Marketplace

Analysis Svcs
Reporting Svcs

StreamInsight

PowerPivot

PowerView

Business
Intelligence

Complex
Event
Processing

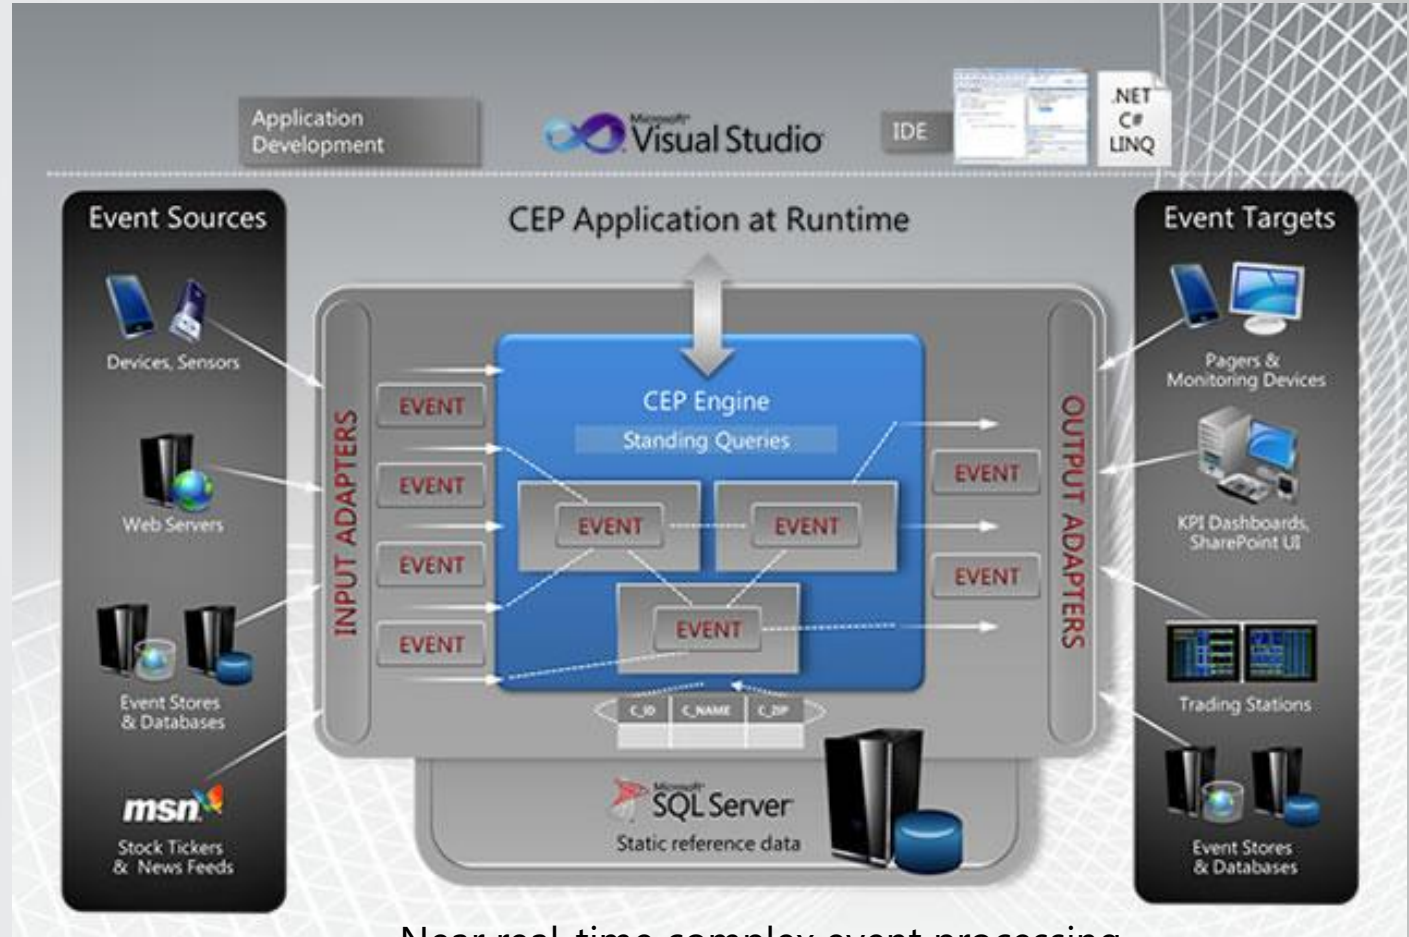
Self-Service
Analytics

Data
Visualization

Windows Azure Platform

Microsoft StreamInsight


- ▶ Process large event volumes across multiple data streams.
- ▶ Gain continuous insight through historical data mining
- ▶ Lower development costs by utilizing existing skill sets
- ▶ Tailor deployment to meet business needs.



Near real-time complex event processing

Hadoop at Microsoft

- ▶ Many wanted to get from Hadoop into Microsoft Information Platform
 - Yahoo – Populating 20+ TB Analysis Services cube from Hadoop
 - Klout – Creating social relevance scoring with Hadoop and SQL Analysis Services
- ▶ Data warehousing market demanded a Hadoop solution, wanted a great port of Hadoop for Windows
- ▶ At Microsoft
 - Work closely with Apache community and key vendors
 - Create great Hadoop experience on Windows
 - Contribute to Hadoop ecosystem – e.g. JavaScript Map/Reduce



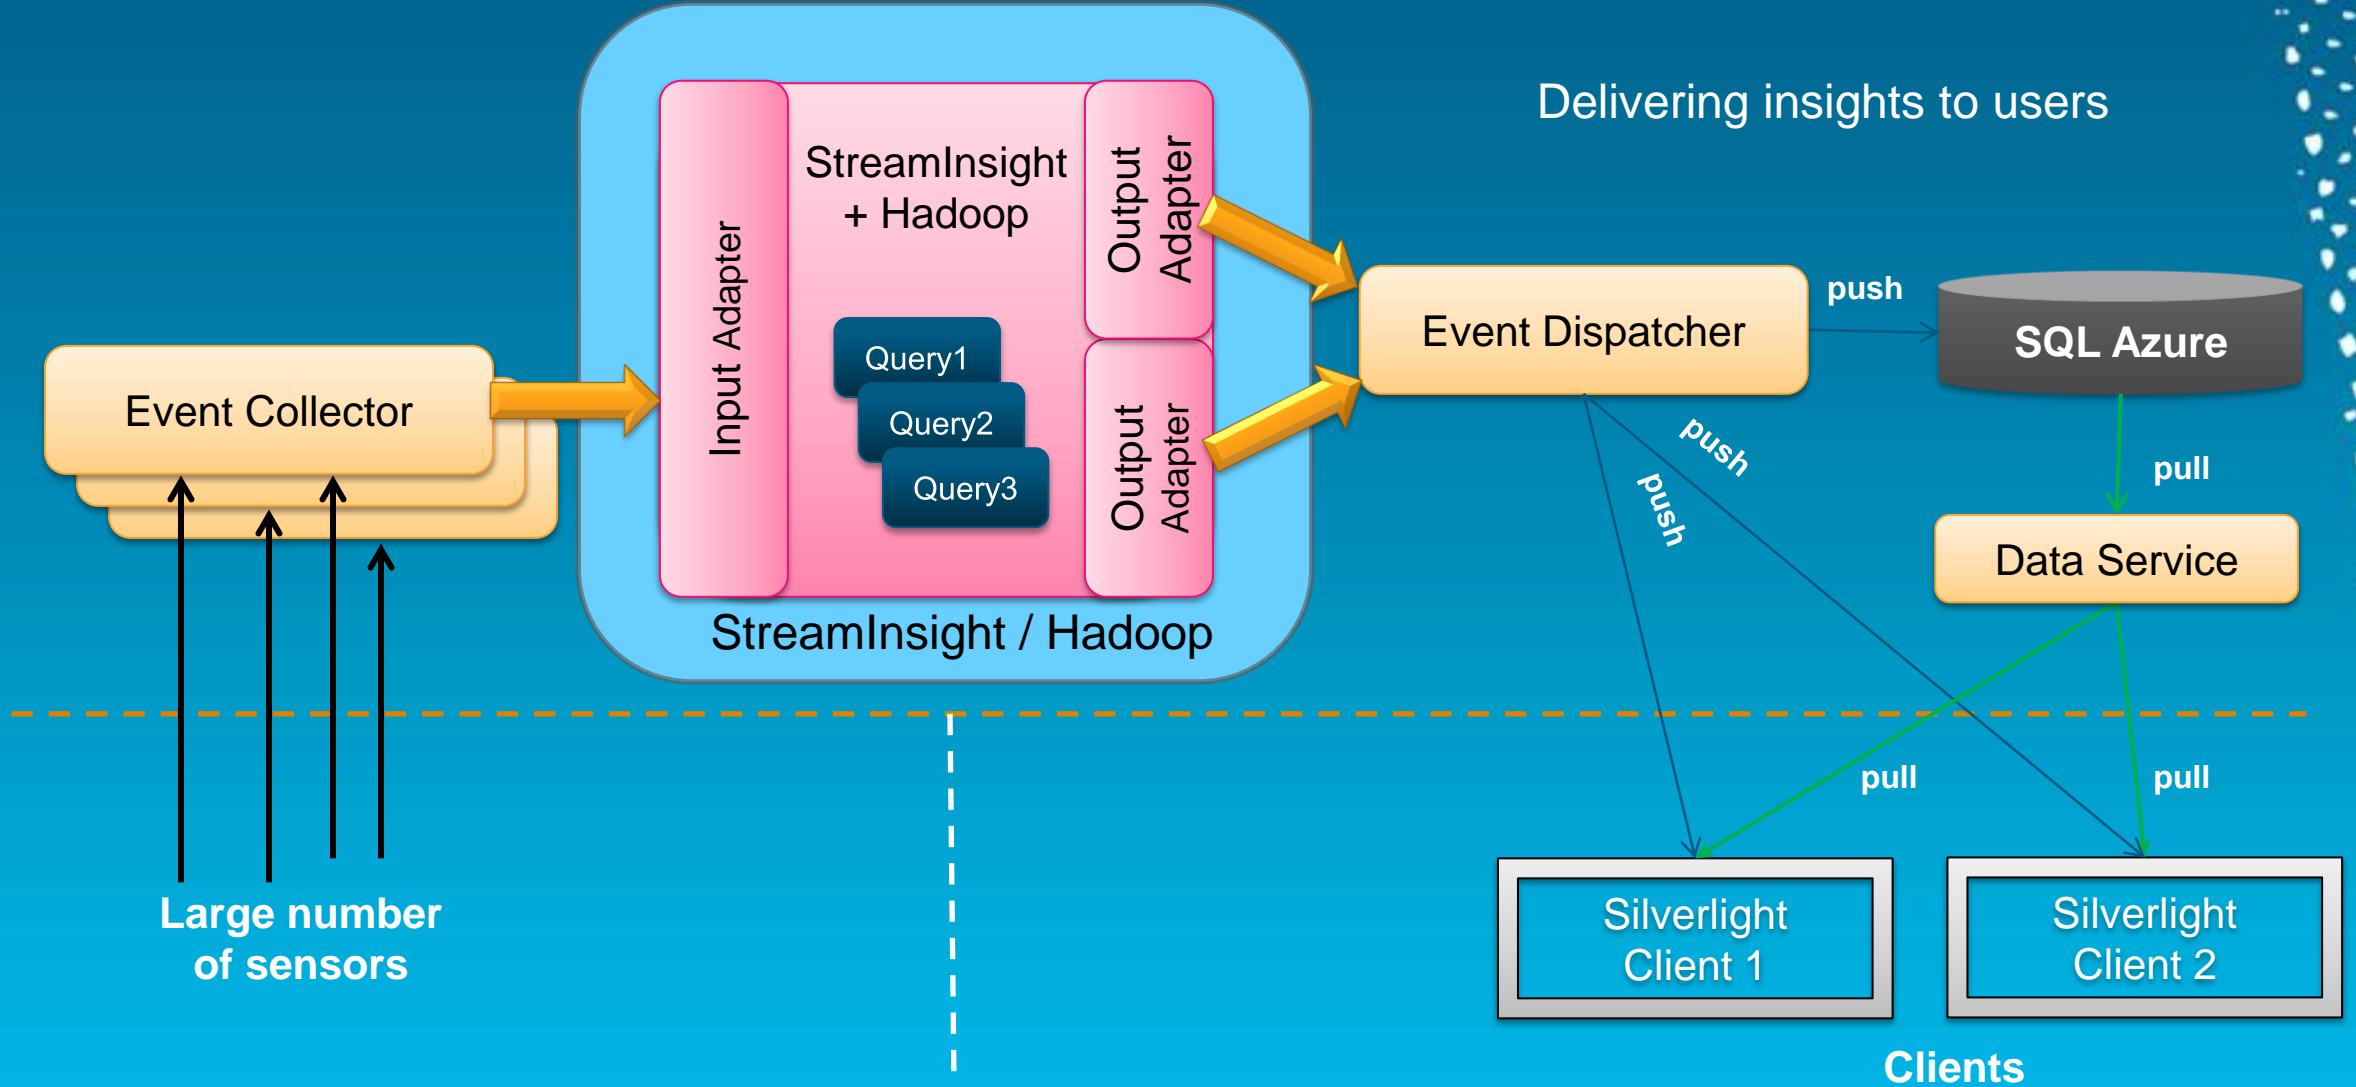
Making it Work: The Internals behind "CitySense"

Web Role
Worker Role

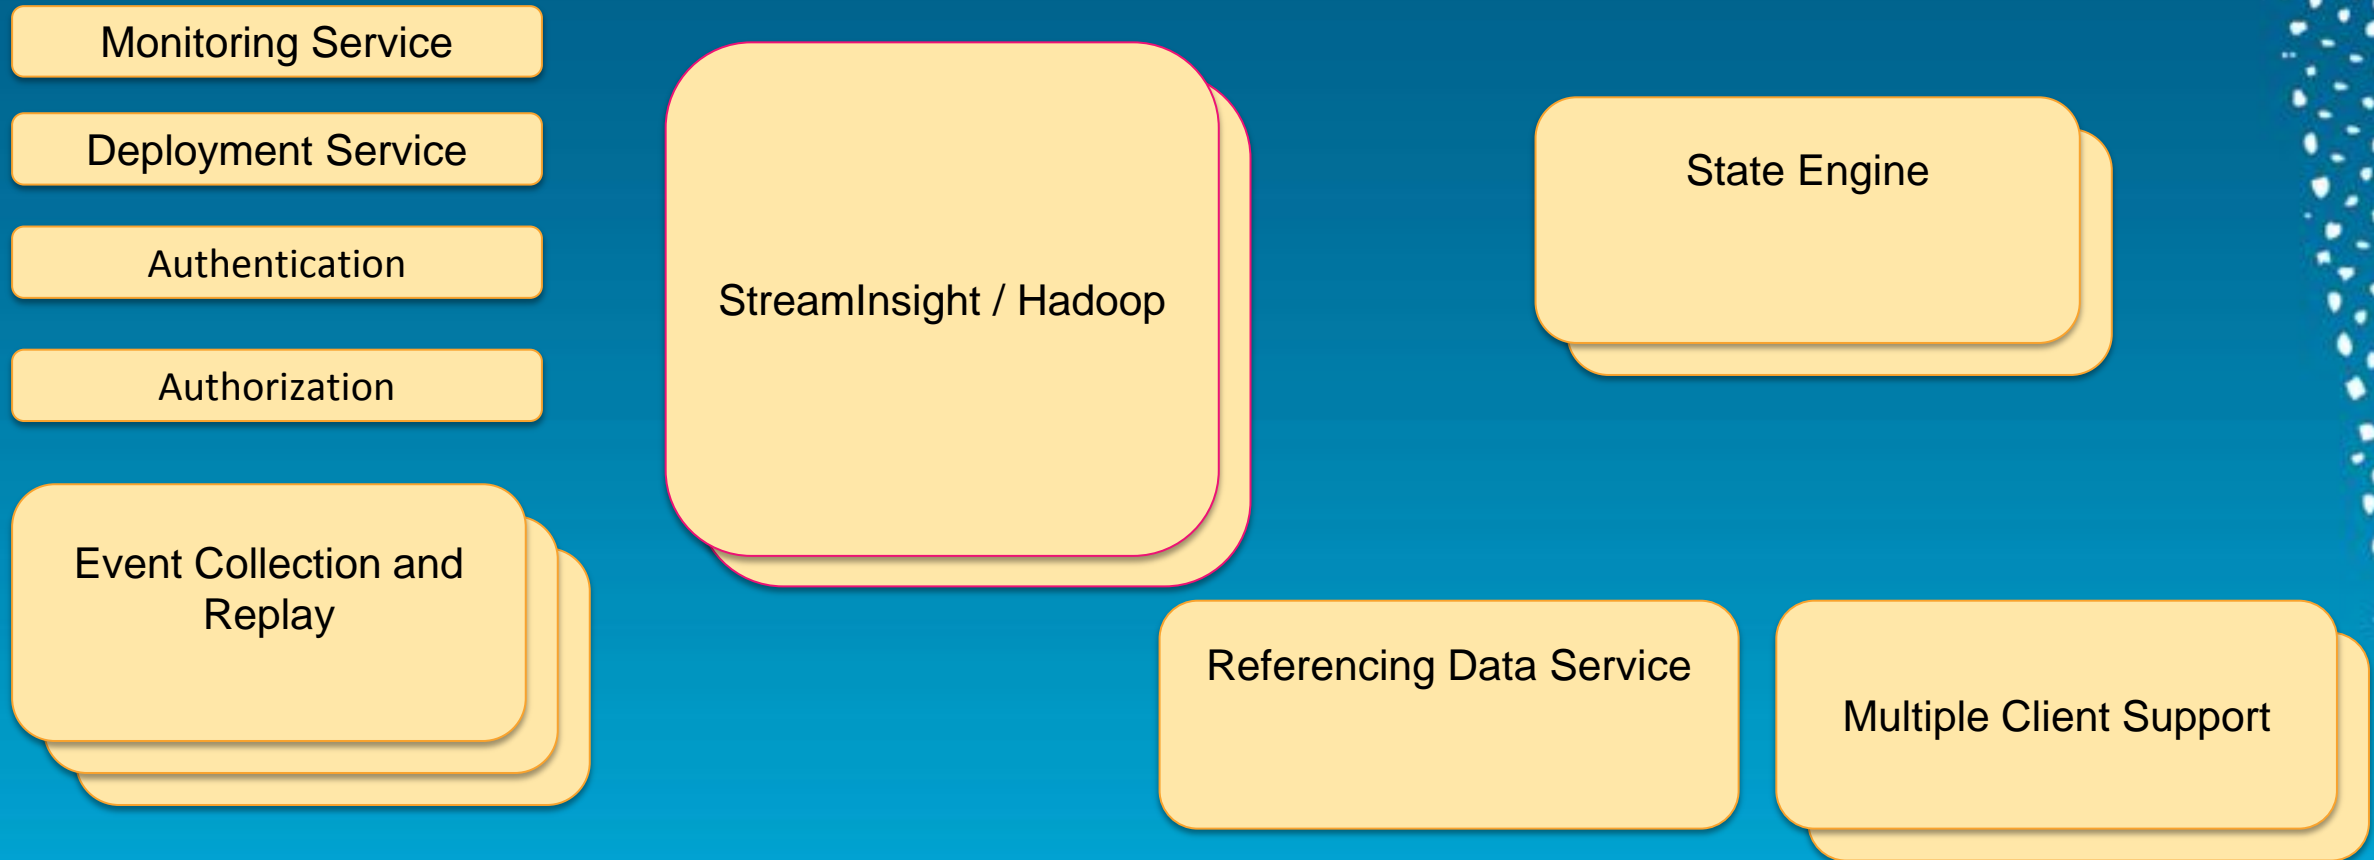
Real-time data processing and analytics

Delivering insights to users

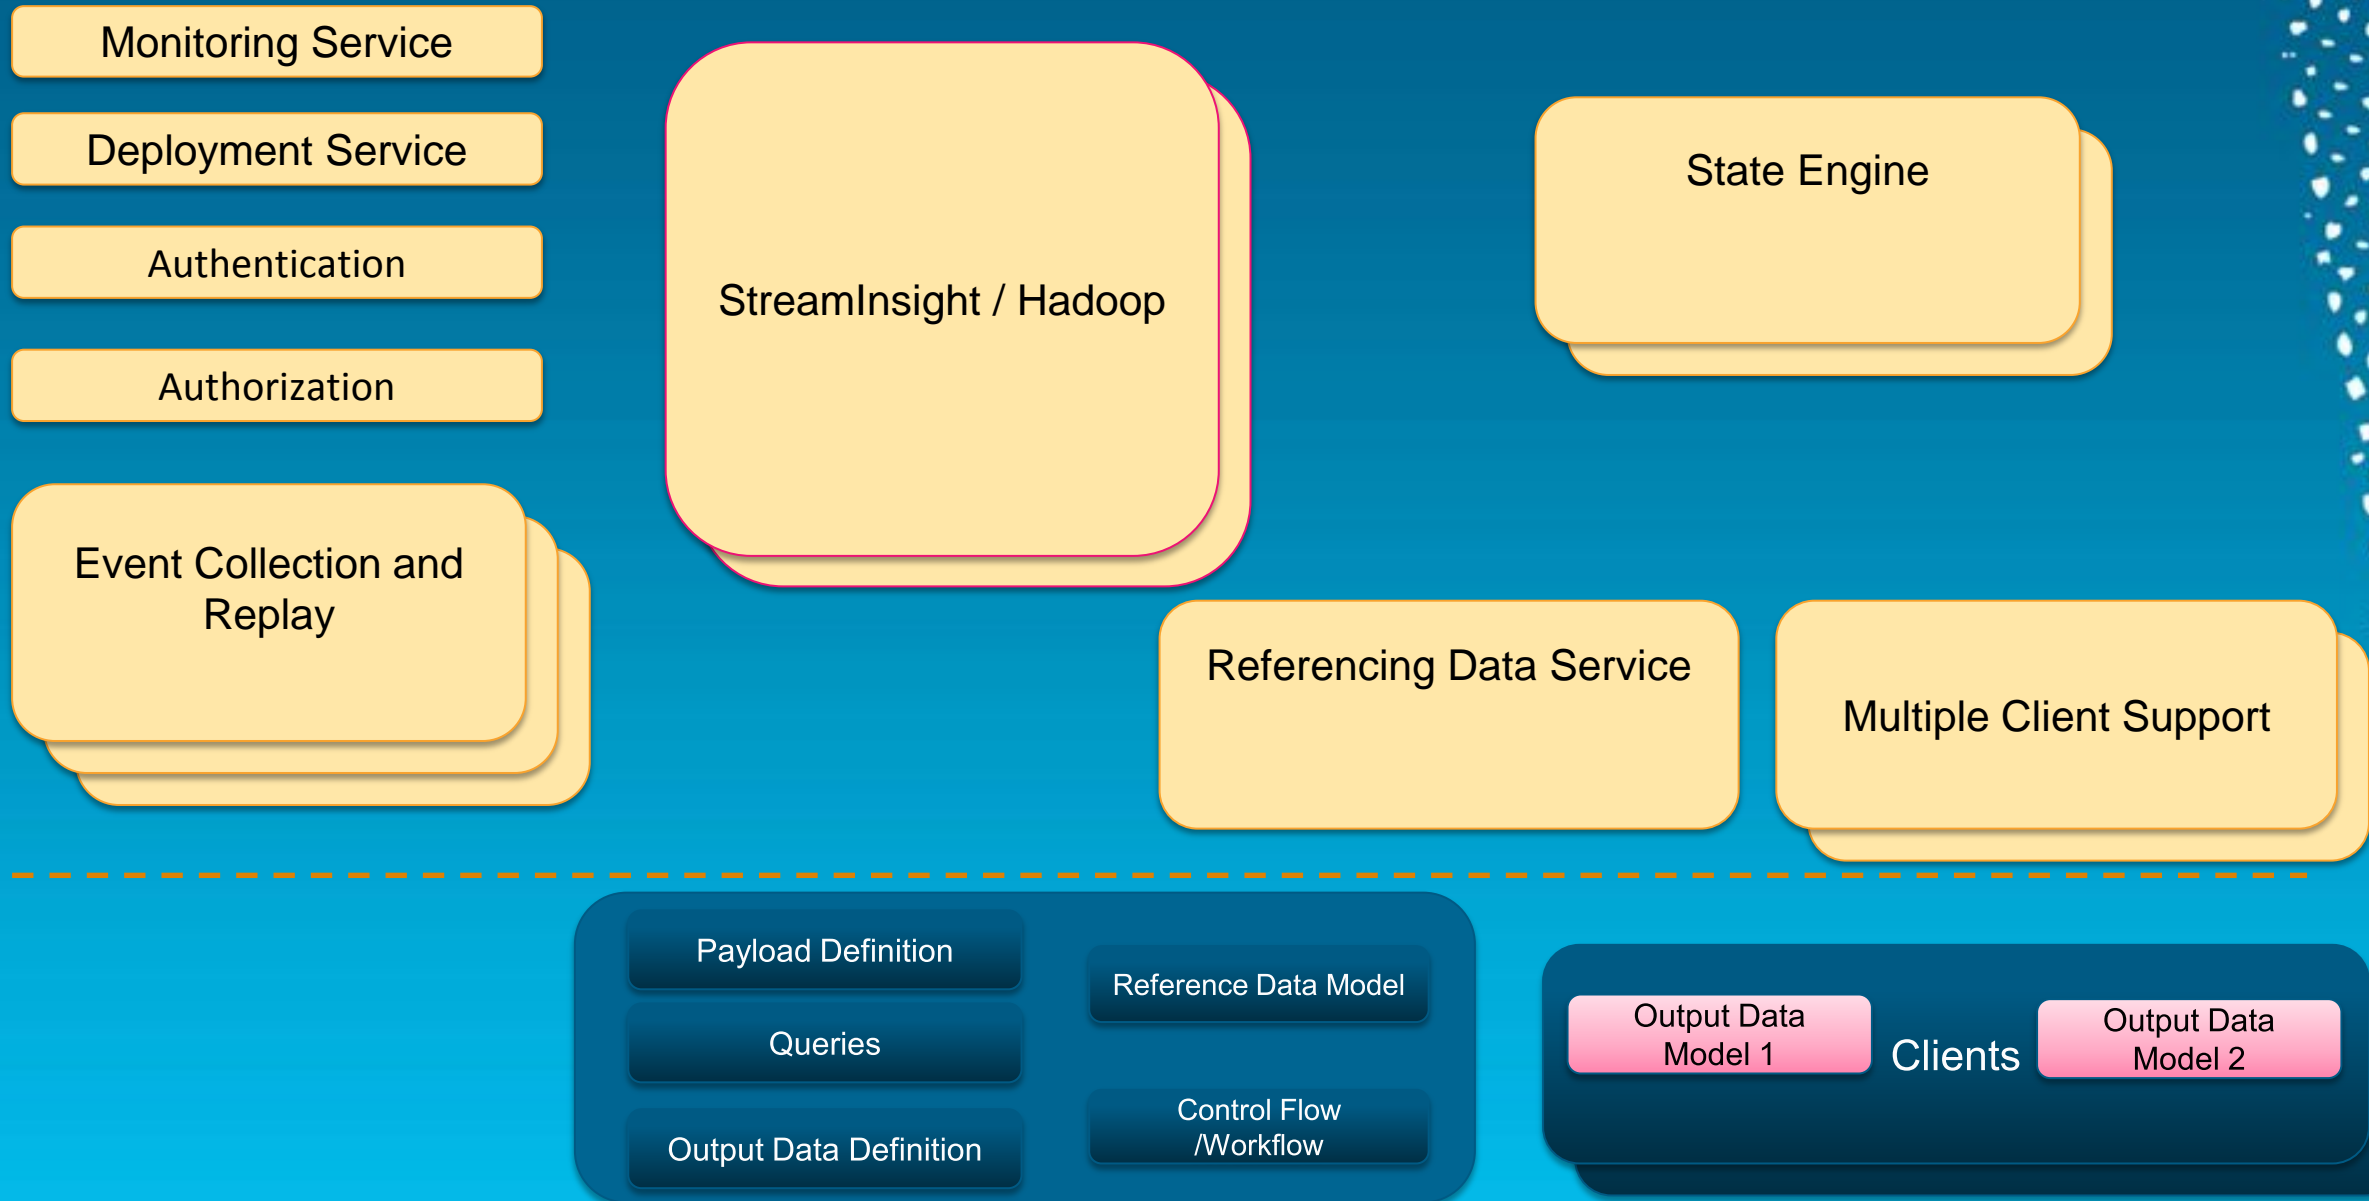
Data Acquisition




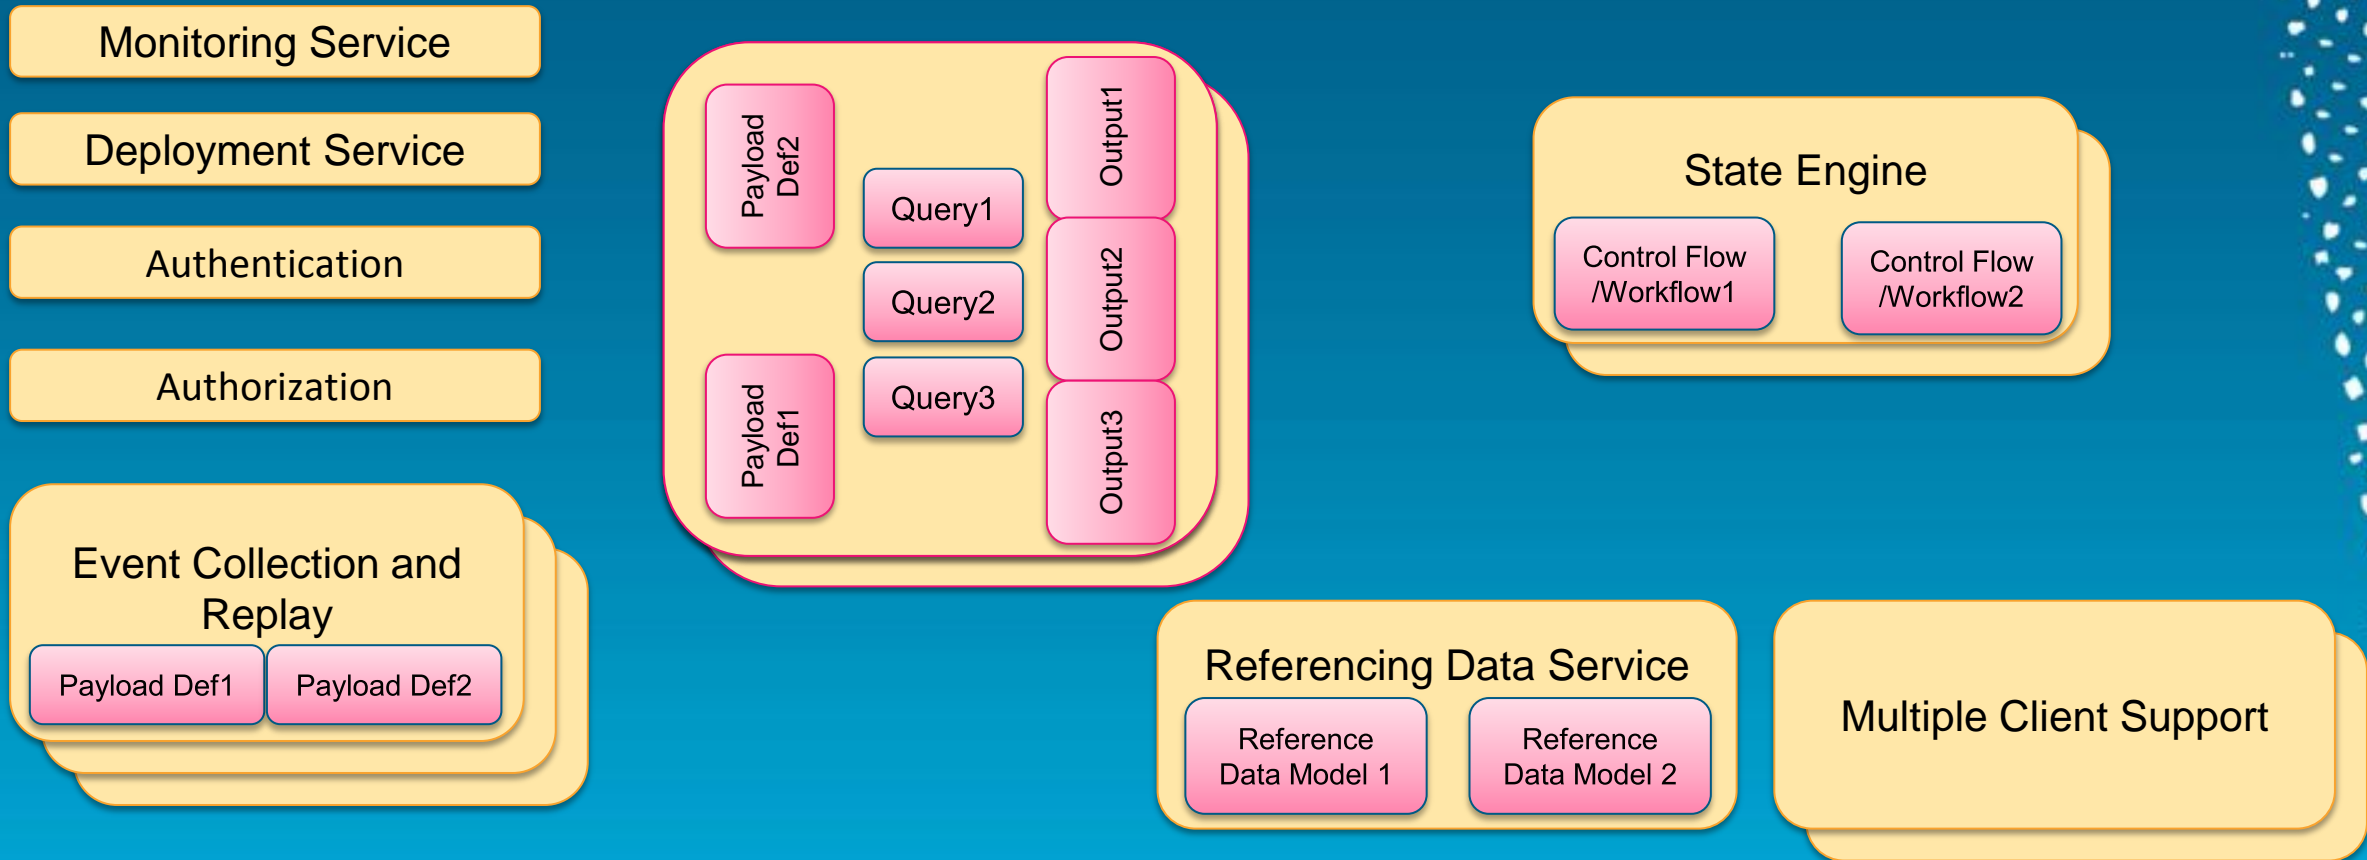
Turning “CitySense” into a Generalized Service



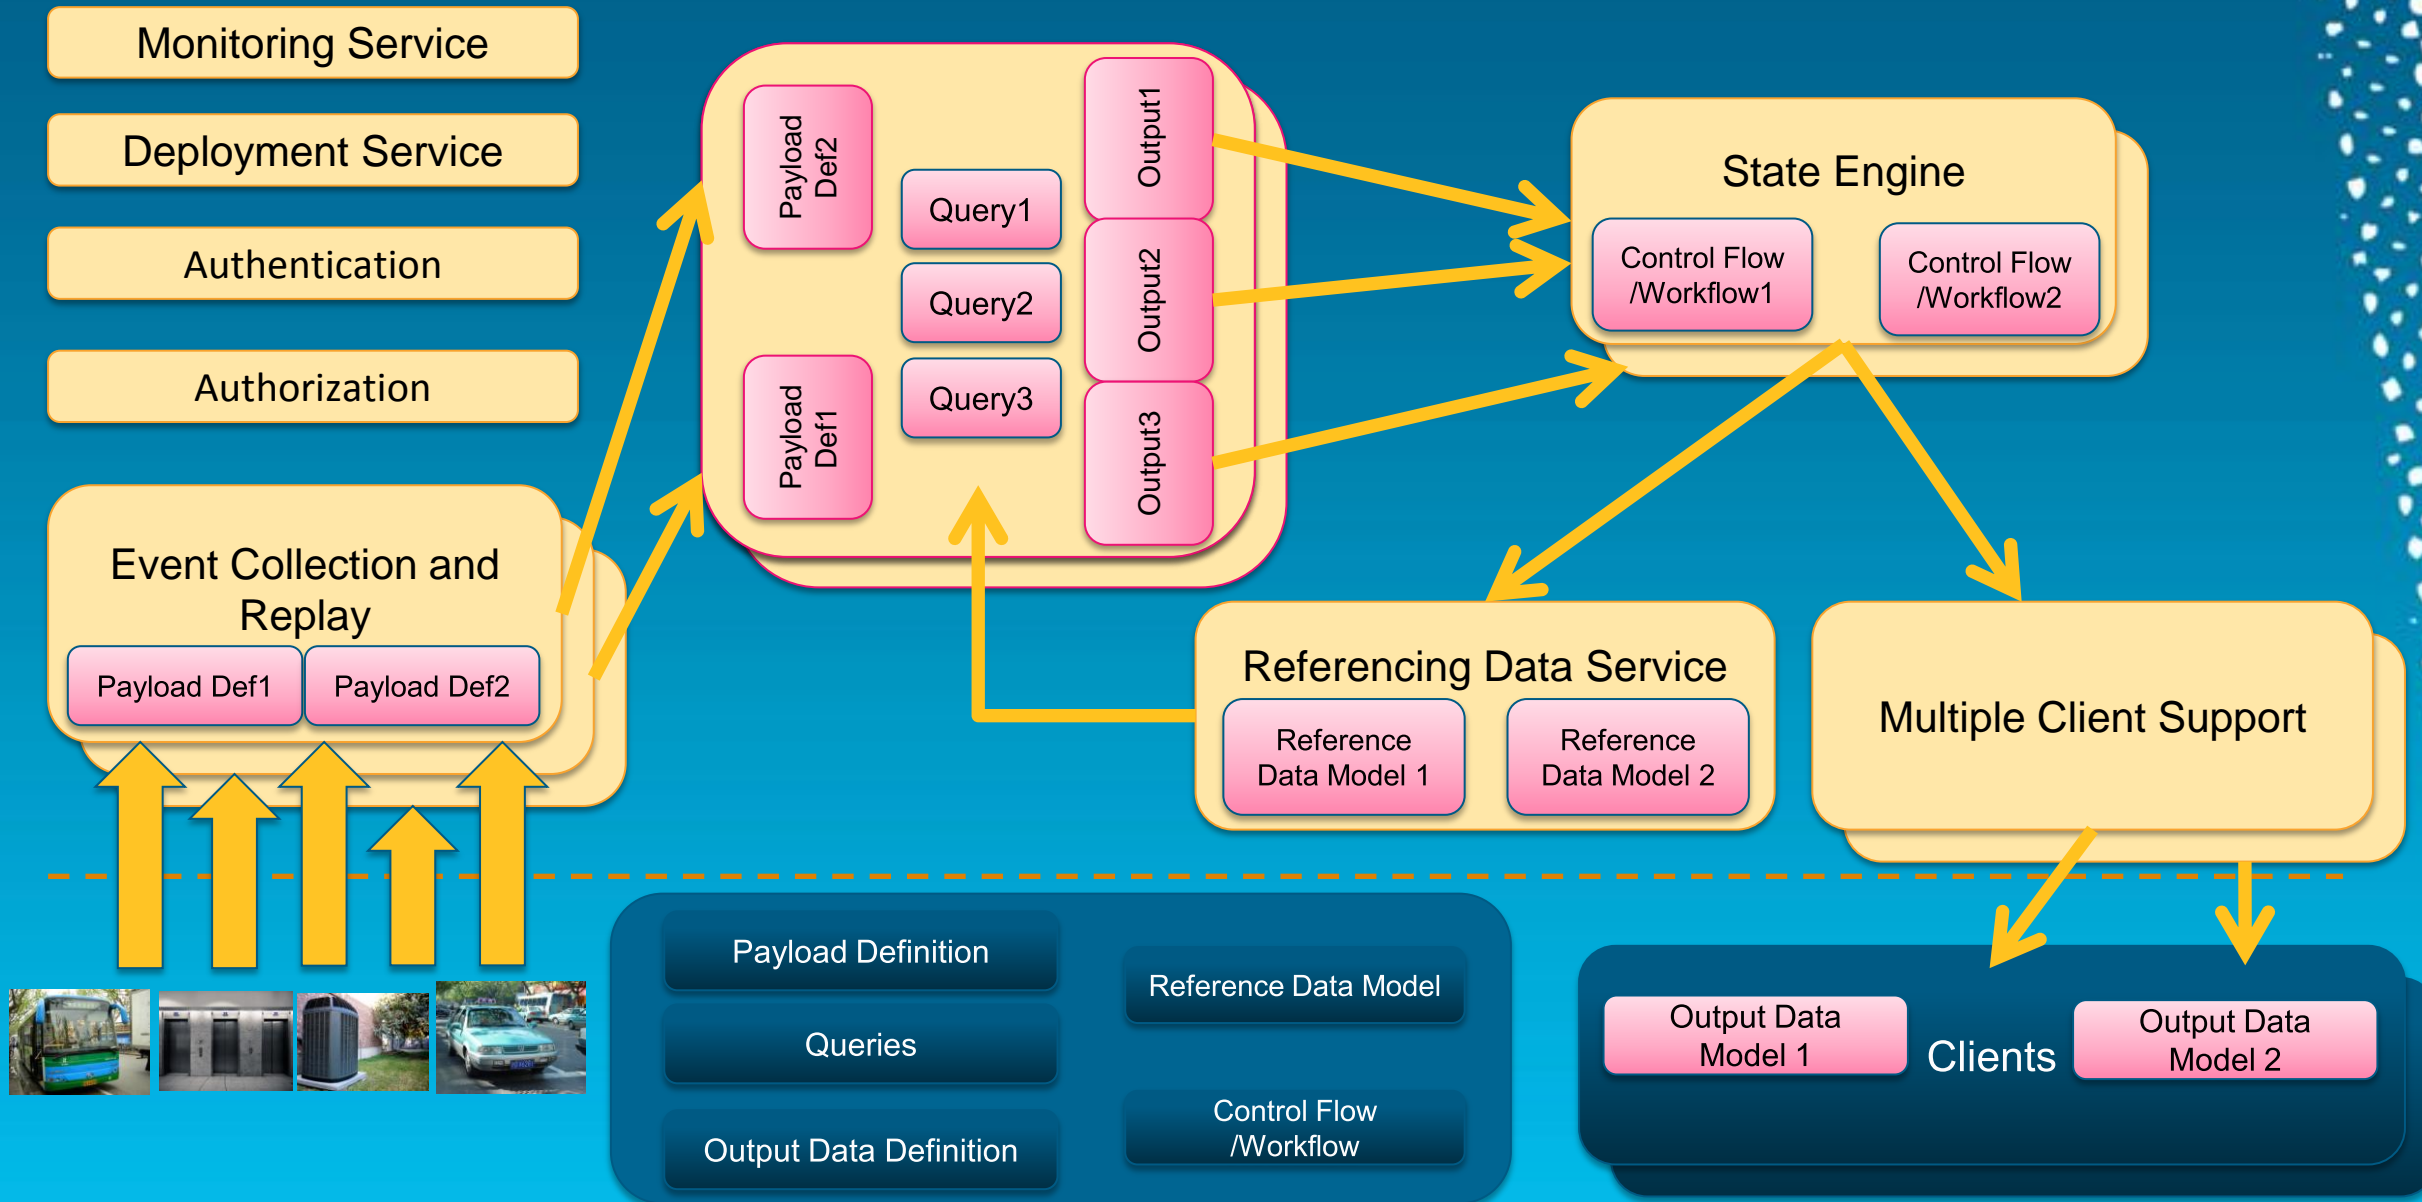
Turning “CitySense” into a Generalized Service



Turning "CitySense" into a Generalized Service



Turning "CitySense" into a Generalized Service



Open Problems for Research

- ▶ Streaming processing of unstructured data
- ▶ Multi-dimensional data: e.g. spatial + temporal + additional dimensions
- ▶ Dynamic deployment of event processing rules and re-balance or scale
- ▶ Management and monitoring, including registration, and retirement of massive numbers of devices in the system

Additional Resources

▶ Windows Azure Platform

Developer information - <http://www.windowsazure.com/en-us/develop/overview/>

Free trial - <http://www.windowsazure.com/en-us/pricing/free-trial/>

▶ StreamInsight (CEP)

<http://www.microsoft.com/sqlserver/en/us/solutions-technologies/business-intelligence/complex-event-processing.aspx>


▶ Big Data

<http://www.microsoft.com/sqlserver/en/us/solutions-technologies/business-intelligence/big-data-solution.aspx>

Microsoft Hadoop on Azure (CTP) - <http://connect.microsoft.com/SQLServer/Survey/Survey.aspx?SurveyID=13697>

▶ SQL Server & BI

SQL server 2012 2012 RC0 - <http://www.microsoft.com/sqlserver/en/us/future-editions/sql2012-editions.aspx>



Q & A

Microsoft Asia-Pacific Research and Development Group

- Largest, most comprehensive R&D base for Microsoft outside of US
- 4,000+ developers and researchers
- Beijing, Shanghai, Shenzhen, Taipei, Hong Kong, Tokyo, Seoul, ...
- Five Focus Areas:
 - Search Technology Center (STCA): Search, Ads
 - Server & Tools (STB): Private & Public Clouds
 - Office: Productivity & Communications
 - MACH: Hardware & Embedded
 - ATC: Emerging Market Solutions



Research



Incubation



Development



Ecosystem

Microsoft Asia-Pacific Research and Development Group

“Asia has become an important R&D base for the company. The expanded charter underscores Microsoft’s continued confidence in the talent pool and innovation capacity in the region.”

*Steve Ballmer, CEO
Microsoft Corporation*



Research



Incubation














Development



Ecosystem

Server & Tools Business China At A Glance

Business Platform	Server & Cloud	Management & Security	Commerce
 	  Winterop Engineering	  	   

China Cloud Innovation Center



Cloud Power

We Are Hiring!!

- Full-time employees
- Interns (summer or other times)
- Shanghai / Beijing

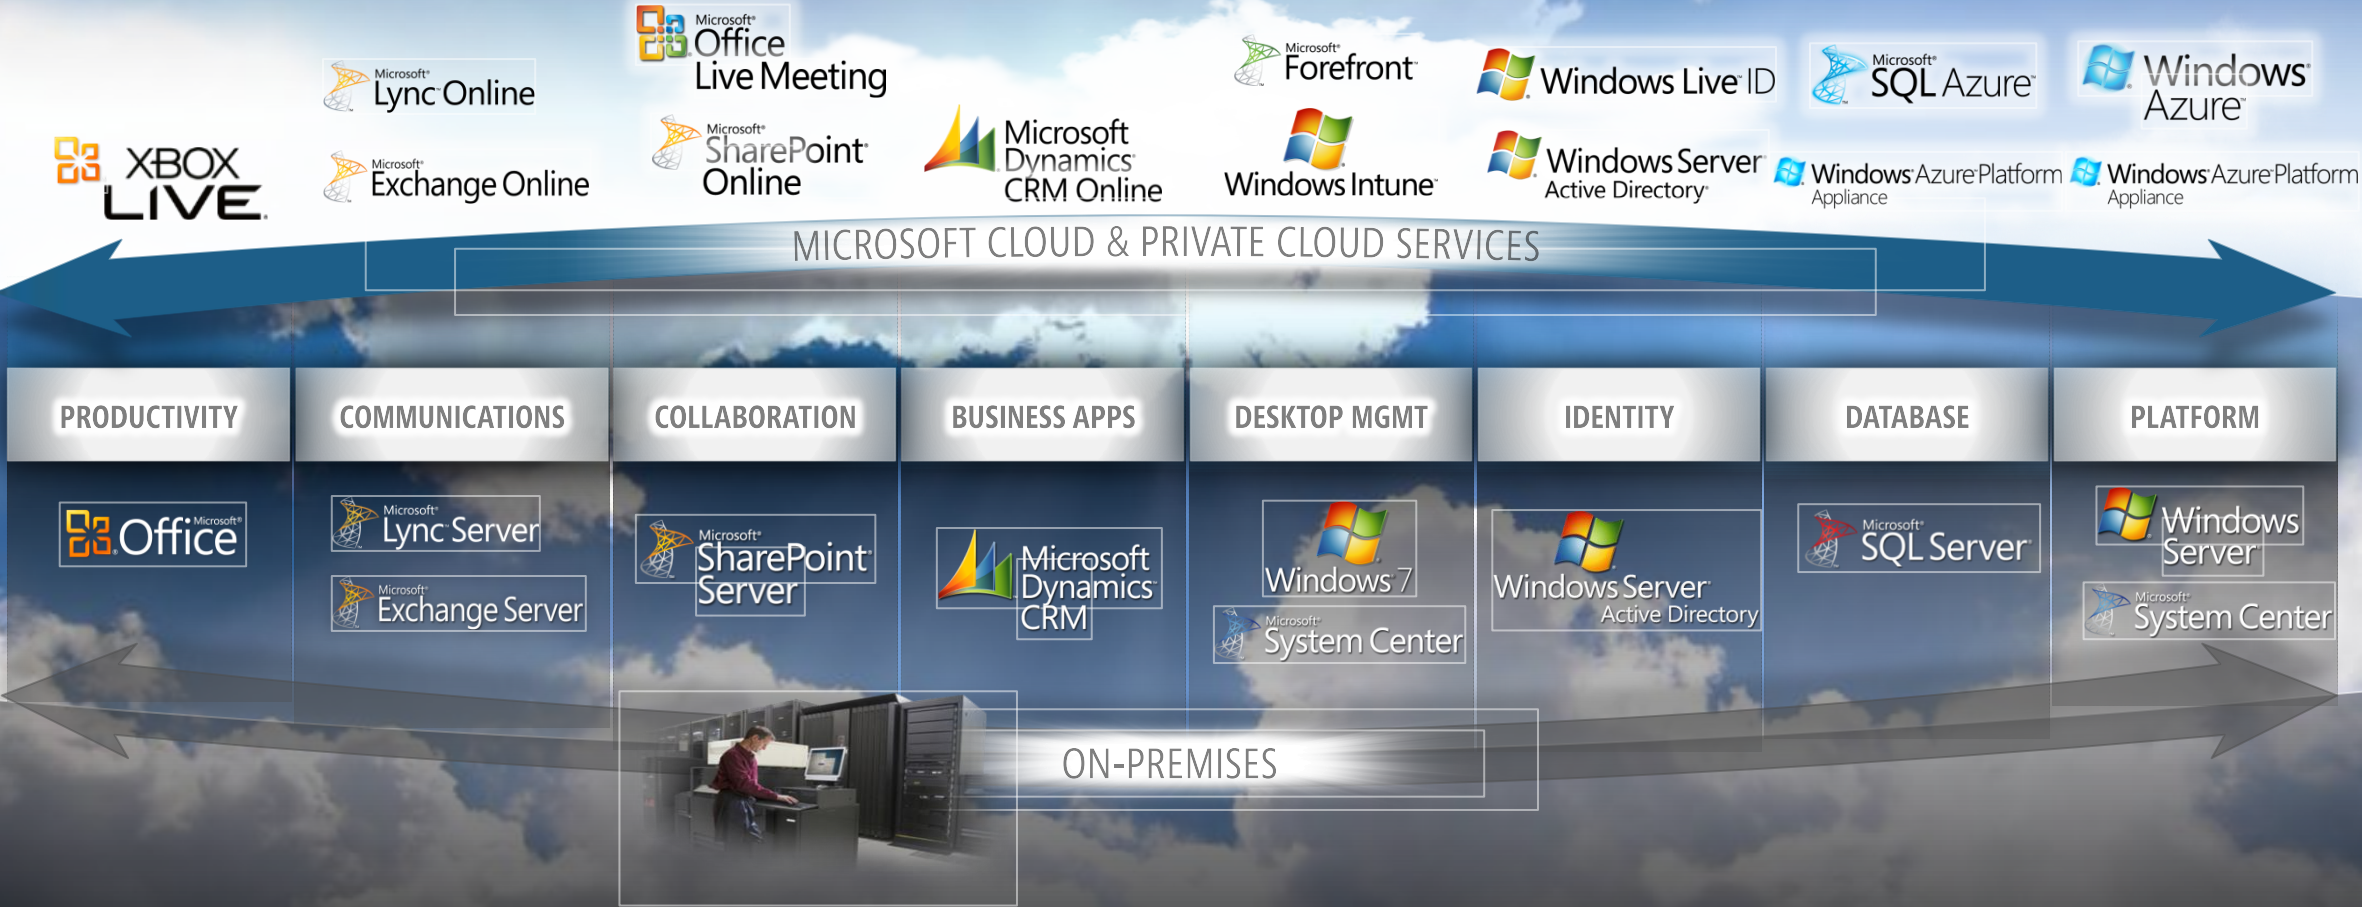
- If you're interested, send email to
 - Bess Zhou (bezhou@microsoft.com) or
 - Rita Zhuang (a-ritaz@microsoft.com) or
 - Prakash Sundaresan (prakas@microsoft.com)

Microsoft[®]

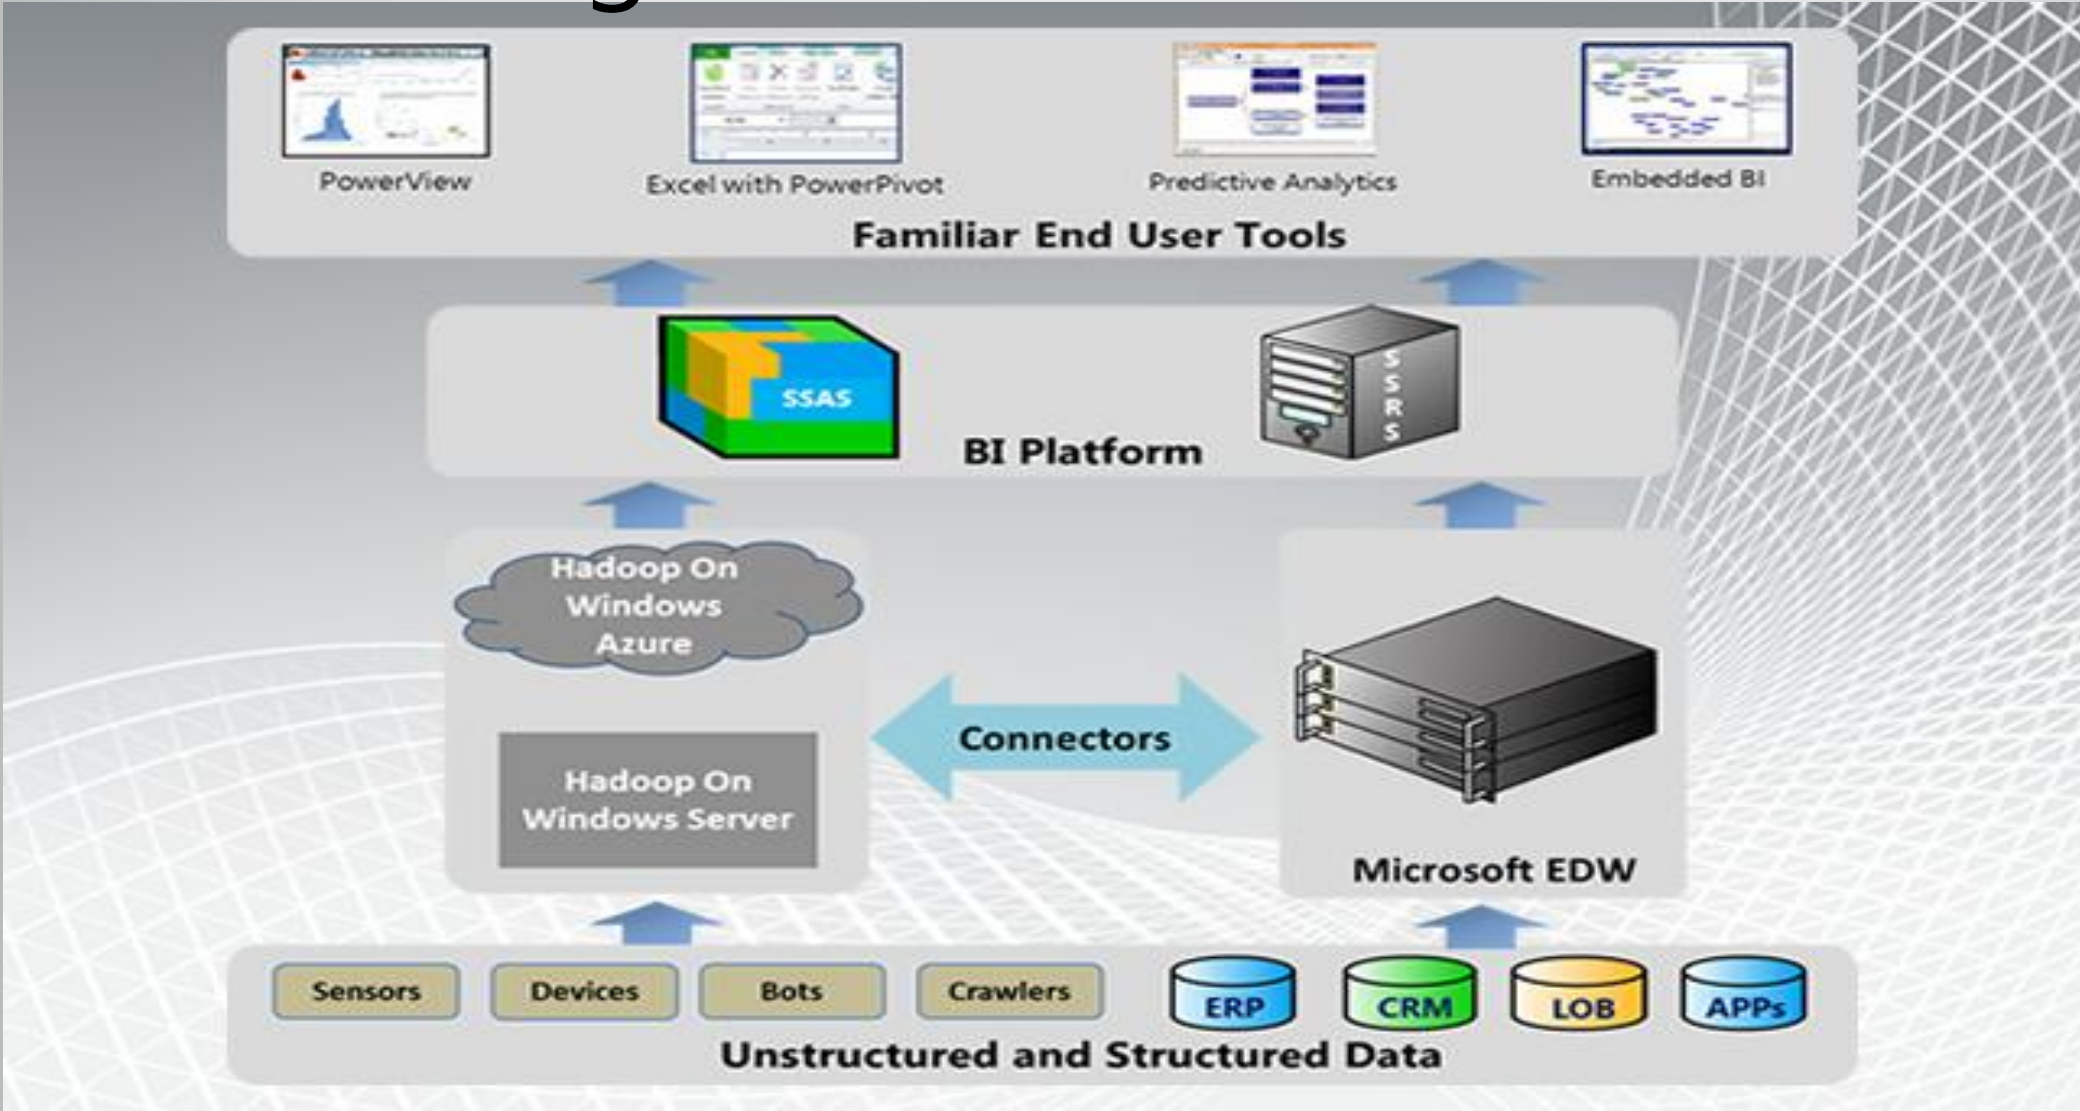
© 2011 Microsoft Corporation. All rights reserved. Microsoft, Windows and other product names are or may be registered trademarks and/or trademarks in the U.S. and/or other countries.

The information herein is for informational purposes only and represents the current view of Microsoft Corporation as of the date of this presentation. Because Microsoft must respond to changing market conditions, it should not be interpreted to be a commitment on the part of Microsoft, and Microsoft cannot guarantee the accuracy of any information provided after the date of this presentation. MICROSOFT MAKES NO WARRANTIES, EXPRESS, IMPLIED OR STATUTORY, AS TO THE INFORMATION IN THIS PRESENTATION.

Microsoft Commercial Cloud Services



Microsoft Big Data Solution



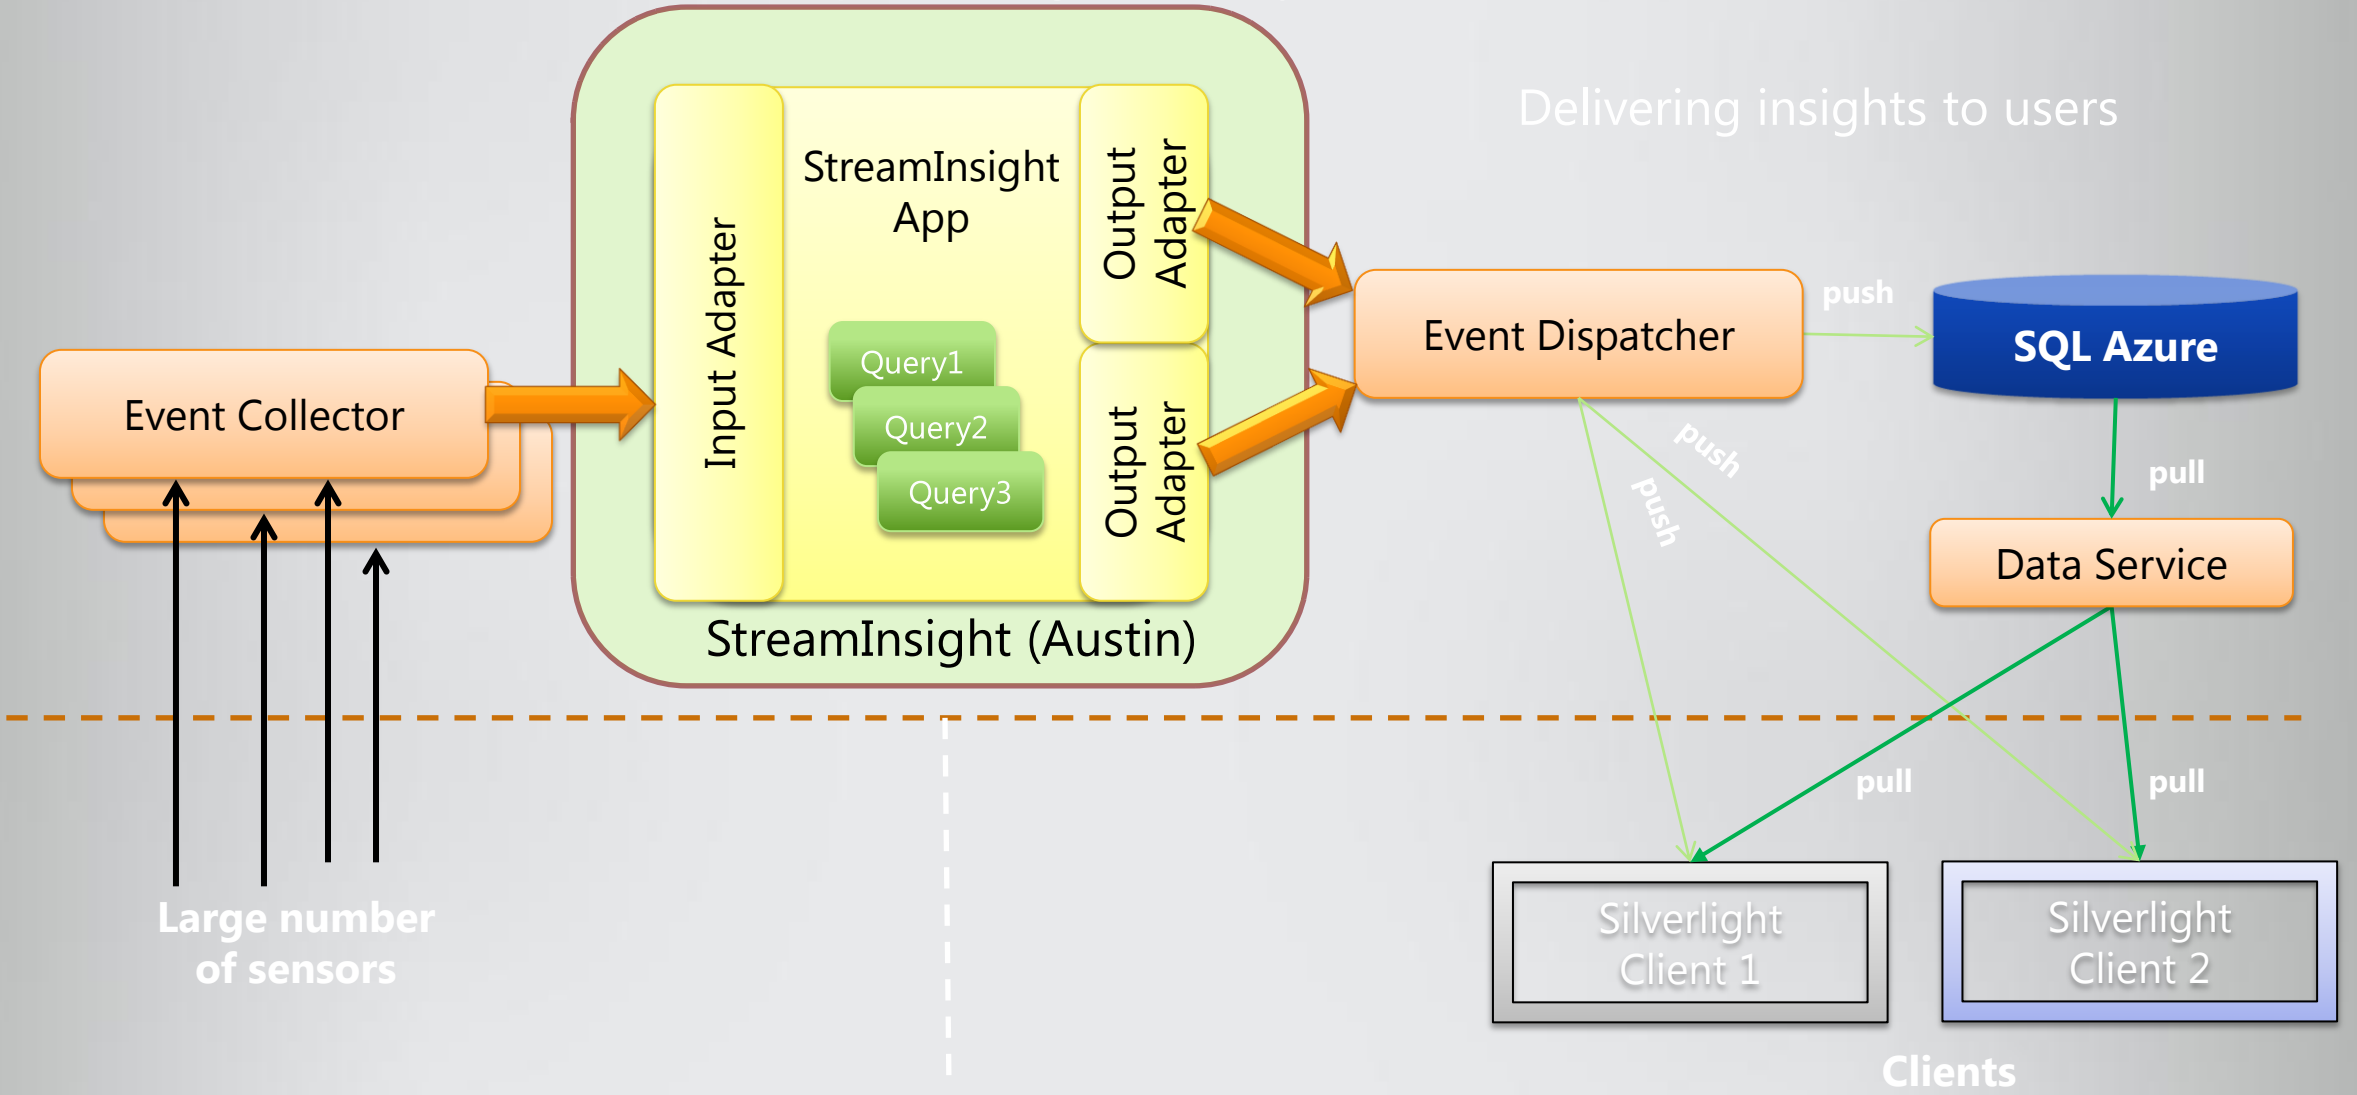
CitySense Architecture

Web Role
Worker Role

Real-time data processing and analytics


Delivering insights to users

Data Acquisition



Clients

Hadoop Architecture



MapReduce
Layer

HDFS
Layer

