


Leveraging Clouds for Small and Big Data

- 6 December 2011



> Featured Speakers



Leon Katsnelson

Program Director,
IM Cloud Computing and Emerging Technologies

IBM



Uri Budnik

Director,
ISV Partner Program

RightScale

Agenda

- DB2 and Cloud
- The Big Data Challenge
- IBM's Approach to Big Data
- Big Data on Cloud
- RightScale Cloud Management Platform

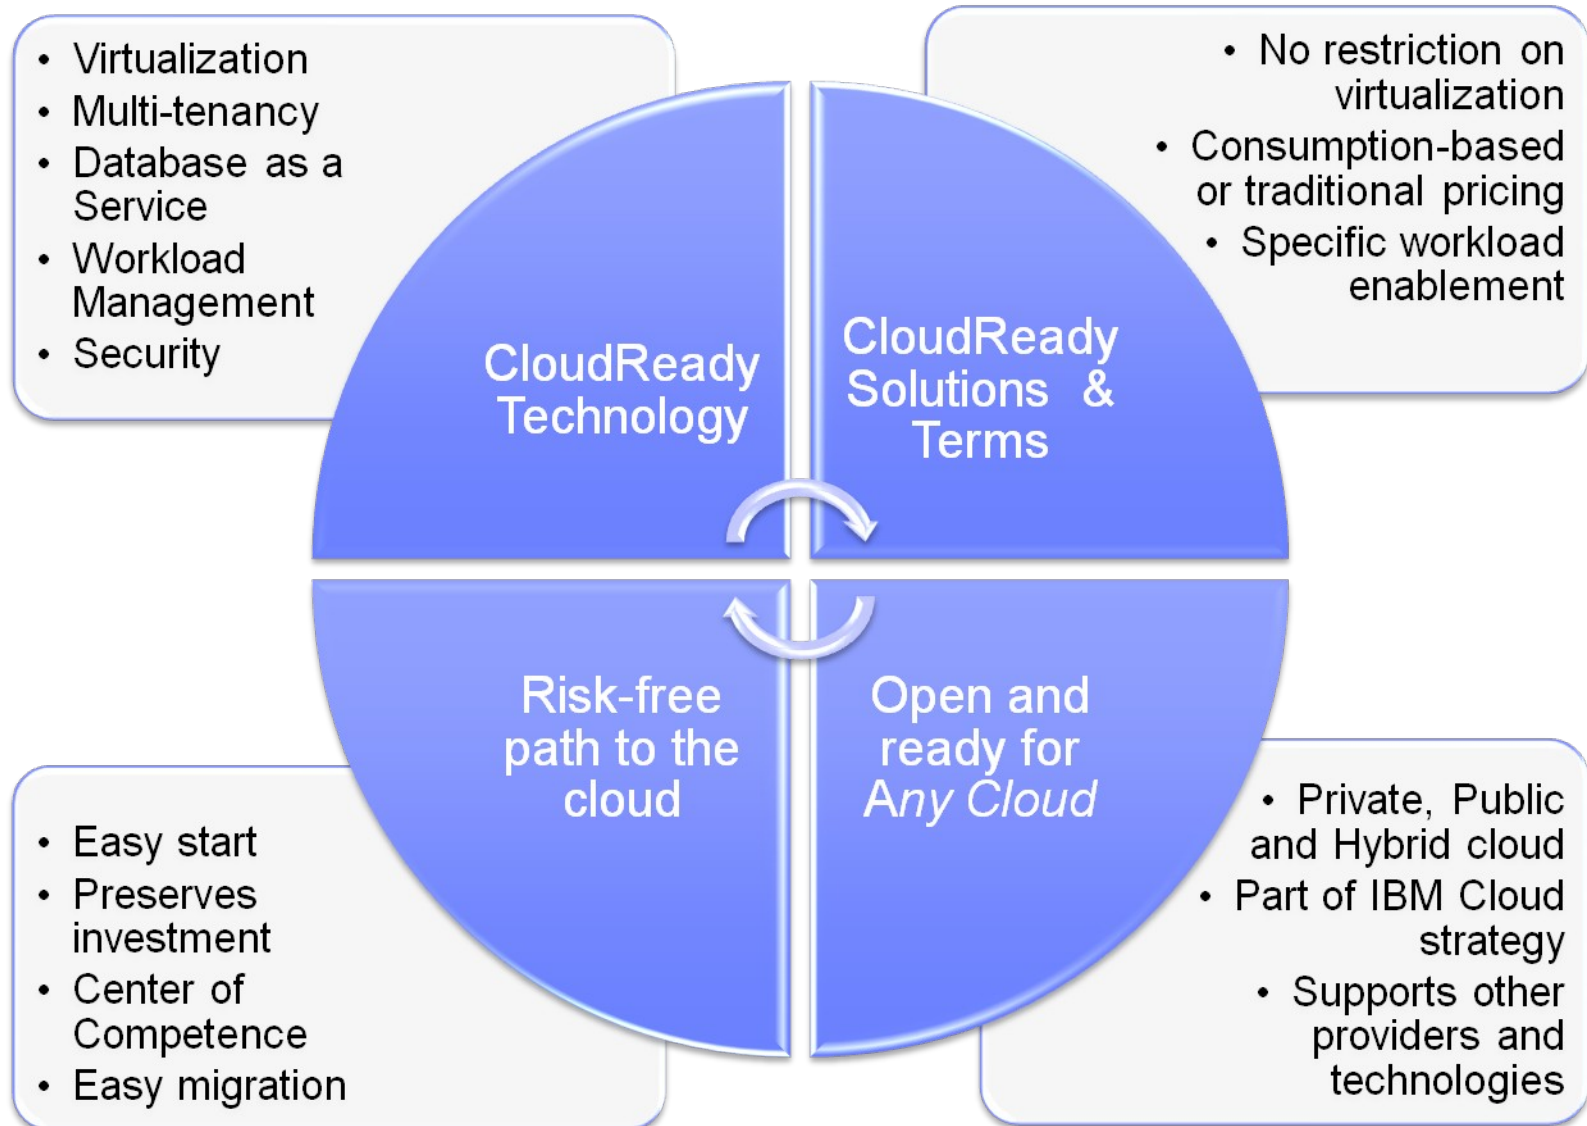
DB2 ON CLOUD

... many reasons clients consider Cloud Computing



DB2 is
CloudReady

DB2 Strategy for Cloud Computing



DB2: Ready for Any Cloud



IBM Workload Deployer



Free DB2-based Website for a year

- All-in-one instance (server):
 - DB2 Express-C
 - Web application server (Nginx, Ruby on Rails)
 - Content Management System (RadiantCMS)
- Runs on Amazon EC2 Micro instance:
 - 613MB of memory
 - Up to 2 EC2 Compute Units (for short periodic bursts)
 - EBS storage
 - 32-bit or 64-bit platform
- Free for a year for new AWS customers
- Tutorial
<http://programmingzen.com/2011/04/04/running-radiant-cms-on-db2-in-the-cloud/>

The screenshot shows a web browser window with the URL `ec2-50-19-50-143.compute-1.amazonaws.com`. The website has a header with a train image and the title "DB2 on Rails" with the tagline "The simplicity of Rails, the power of DB2". Navigation links for "Home", "Articles", and "Links" are visible.

The main content area features a "Welcome!" section with the Radiant CMS logo and text: "This is Radiant, a Ruby on Rails CMS. Radiant CMS is a powerful, yet simple to use, content management system. It's written in the popular Ruby on Rails web application framework and it fully supports DB2."

Below this is an article titled "DB2 on RightScale.com" posted by Administrator on November 27, 2009. The article text reads: "As I mentioned before, high degree of automation is the hallmark of Cloud Computing and is the attribute that delivers the self service aspect of the cloud that users appreciate so much. In addition to automating installation of DB2 by making it available as a set of packaged Amazon Machine Images (AMIs) DB2 is also available for deployment as a set of templates on RightScale.com."

Another article titled "What is DB2 pureScale?" is also shown, posted by Administrator on November 27, 2009. The text states: "Commercial databases, particularly DB2, excel at vertical scalability, which incidentally is one of the weakest aspects of open source RDBMS like MySQL. When it comes to horizontal scalability though, even commercial databases tend to struggle a little. Scalability is achieved, but a lot of work is required to configure complex setups involving masters, slaves, replication, failover servers, and so on. Until now."

On the right side of the page, there are several promotional boxes:

- A green box: "This is just an example of what your blog/CMS could look like. To really make it your own, login at /admin with the user **admin** and password **radiant**, in order to configure and customize your Radiant CMS. Be sure to read the documentation as well."
- An RSS feed icon with the text: "RSS Get the feed for our latest news. Here is where you'd convince your users to subscribe. You can change this description in the rss snippet within the admin section. And don't forget to customize your feed, by altering the RSS Feed page within the same admin."
- A download icon with the text: "FREE Download DB2 Express-C DB2 is free! DB2 Express-C is a 100% free version of DB2. It doesn't impose any limits on the number of users, connections, databases, or your data size. That's right, one of the world's leading datasevers is now available to you for free (and it's powering this very site!)"
- A wrench icon with the text: "Get FREE Tools Manage DB2 Free Tools Not only is DB2 Express-C free of charge, but you can also download IBM Data Studio, a powerful tool to administer and develop for DB2, free of charge as well."

Dogfooding

@your place, @your pace DB2 skills acquisition

Community operated property

10,500+ Registered Students

Runs on the cloud, Managed by RightSale

Leverages DB2 technologies e.g. HADR

Continuous availability since launch in January 2011

The screenshot shows the DB2 University website. At the top, there is a navigation bar with links for HOME, LEARN, DOWNLOAD, RESOURCES, and JOBS. The main header features the DB2 University logo with the tagline "Learn from the industry's best" and a "BETA" badge. There are buttons for "login" and "sign up", and a search bar for "search courses".

The main content area is divided into several sections:

- Welcome message by Raul Chong:** A video player showing a man sitting at a desk with a laptop and books.
- Why register?:** A list of benefits:
 - Easy and Affordable:** Learning DB2® has never been more affordable! Many courses are FREE!
 - Latest industry trends:** Acquire valuable skills and get updated about industry's latest trends right here. Today!
 - Learn from the Experts!** DB2 University offers education about DB2 and other technologies by the industry's best!
 - Learn at your Own Pace!** Find everything right here when you need it and from wherever you are.
- Video Course Overviews:** A list of courses with "go to course" links:
 - SQL Fundamentals I
 - DB2 Essential Training I
 - Data Studio Essential Training I
 - DB2 Essential Training II
- sign me up:** A button to register, accompanied by a "View sample completion document!" link and a thumbnail of a document.
- Study Made Easy! Here, on DB2 University!:** A promotional banner featuring a man at a laptop. Text includes "sign up now for a chance to receive: DISCOUNTED (50% off) IBM Certification Exam Vouchers FREE DB2 Books" and a "sign up" button.
- Student Testimonials:** A section titled "Balázs (USA)" with a "go to sign up" link. The testimonial text describes the training material as short, easy to digest, and includes transcripts and examples.

At the bottom right, there are links for "about us", "legal", "contact", and "bug reports".


DB2 Galileo Early Experience Program

Clients get hands-on experience with the upcoming version of DB2 in 30 minutes

- No need to purchase, rack and cable servers
- No operating system patching
- No need to install DB2
- Always use the latest drop i.e. no wasted time dealing with resolved issues

How:

- Apply for DB2 Galileo Early Experience at <http://freedb2.com/db2-early-experience-program-on-the-cloud/>
- Email IMcloud@ca.ibm.com to request \$500 credit
- Enjoy your own private virtual DB2 server in one of the data centers WW



DB2 on Cloud Resources

- IM Cloud Computing Center of Competence
 - IMcloud@ca.ibm.com
- DB2 on Cloud homepage
 - www.ibm.com/db2/cloud

THE “BIG DATA” CHALLENGE

What we hear from customers

- Lots of potentially valuable data is dormant or discarded due to size/performance considerations
- Large volume of unstructured or semi-structured data is not worth integrating fully (e.g. Tweets, logs, . . .)
- Not clear what should be analyzed (exploratory, iterative)
- Information distributed across multiple systems and/or Internet
- Some information has a short useful lifespan
- Volumes can be extremely high
- Analysis needed in the context of existing information (not stand alone)



Big Data Presents Big Opportunities

Extract insight from a high volume, variety and velocity of data in a timely and cost-effective manner



Variety: Manage and benefit from diverse data types and data structures

Velocity: Analyze streaming data and large volumes of persistent data

Volume: Scale from terabytes to zettabytes

Big Data Scenarios Span Many Industries



Multi-channel customer sentiment and experience analysis



Detect life-threatening conditions at hospitals in time to intervene



Predict weather patterns to plan optimal wind turbine usage, and optimize capital expenditure on asset placement



Make risk decisions based on real-time transactional data



Identify criminals and threats from disparate video, audio, and data feeds

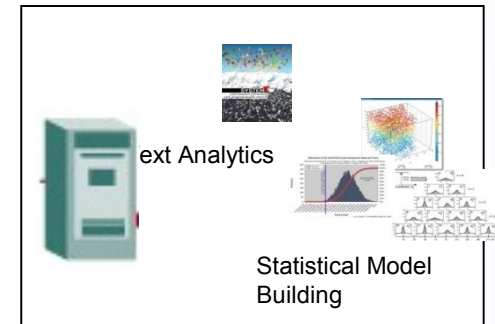
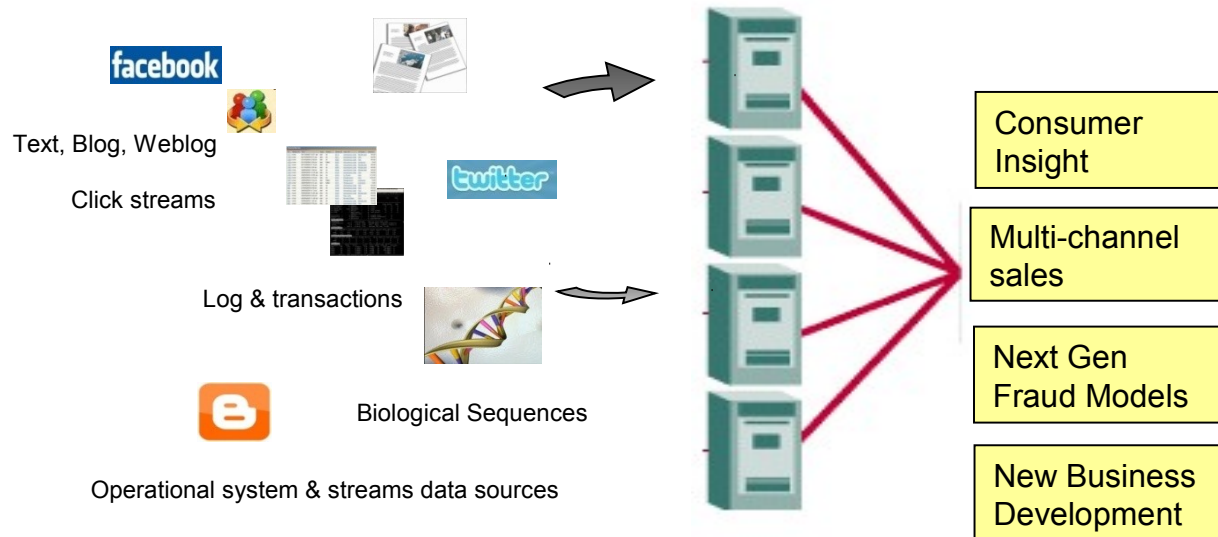
Customer Engagements

Use patterns

- Customer sentiment analysis (cross-sell, up-sell, campaign management)
- Integrated retail and web customer behavior modeling
- Predictive modeling (credit card fraud)
- System log analytics (reduce operational risk)

Common requirements

- Extract business insight from large volumes of raw data (often outside operational systems)
- Integrate with other existing software
- Ready for enterprise use

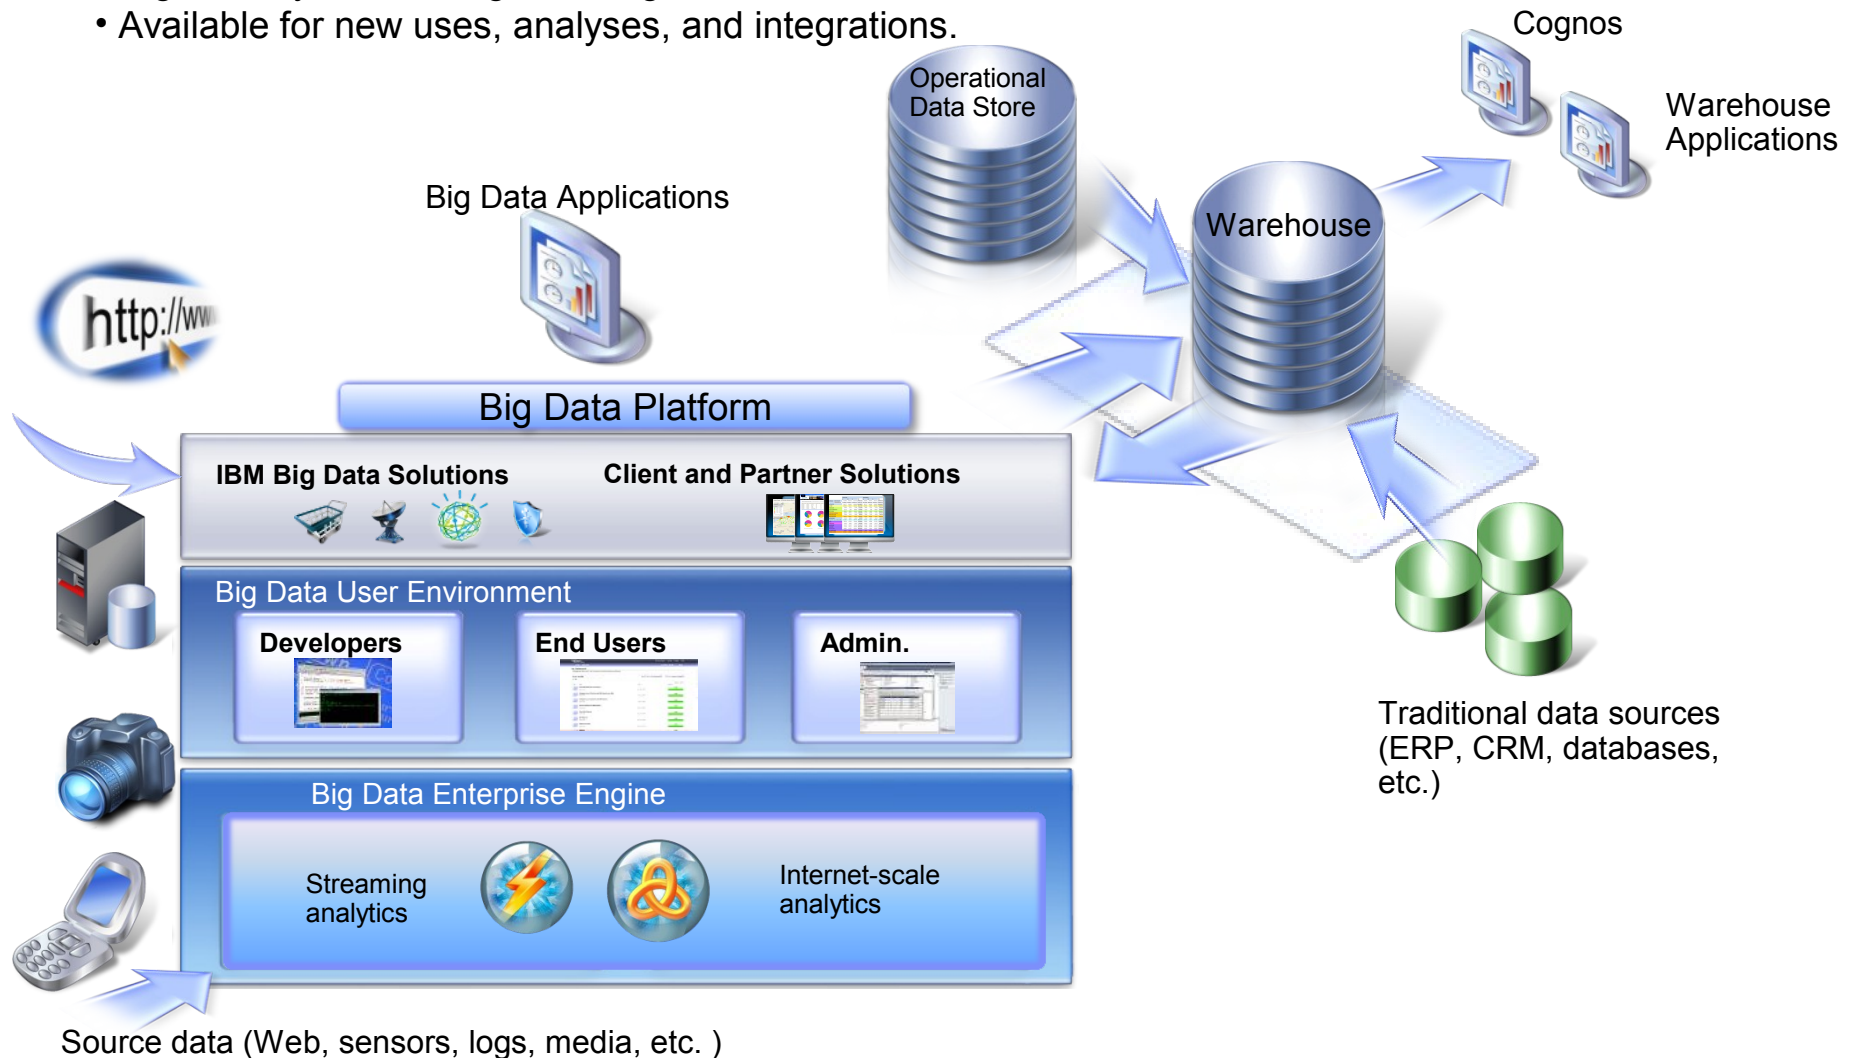


Big Data

IBM'S APPROACH

Big Data: an integral part of an enterprise data platform

- Manage Big Data from the instant it enters the enterprise
- High fidelity – no changes to original format
- Available for new uses, analyses, and integrations.



IBM's Platform Addresses Key Requirements

1. Platform for V³ – Variety, Velocity, Volume

- Variety - manage data & content “As Is”
- Handle any velocity - low-latency streams and large volume batch
- Volume - huge volumes of at-rest or streaming data

2. Analytics for V³

- Analyze Sources in their native format - text, data, rich content
- Analyze all of the data - not just a subset
- Dynamic analytics - automatic adjustments and actions

3. Ease of Use for Developers and Users

- Developer UIs, common languages & automatic optimization
- End-user UIs & visualization

4. Enterprise Class

- Failure tolerance, Security and Privacy
- Scale Economically

5. Extensive Integration Capabilities

- Integrate wide variety of sources
- Leverage enterprise integration technologies

Big Data Platform



Platform Vision



IBM Big Data Solutions



Client and Partner Solutions



Big Data Accelerators

Text Statistics Financial Geospatial Acoustic

Image/Video Mining Times Series Mathematical

Connectors

Applications

Blue Prints

Big Data Enterprise Engines

InfoSphere Streams

InfoSphere BigInsights

Productivity Tools & Optimization

Workload Management & Optimization

Provisioning

Workflow

Job Scheduling

Job Tracking

Data Ingestion

Management

Admin Tools

Configuration Manager

Activity Monitor

Identity & Access Mgmt

Data Protection

Rules / BPM

iLog & Lombardi

Data Warehouse
InfoSphere Warehouse

Warehouse Appliances
IBM & non-IBM

Master Data Mgmt
InfoSphere MDM

Database
DB2 & non-IBM

Content Analytics
ECM

Business Analytics
Cognos & SPSS

Marketing
Unica

Data Growth Management
InfoSphere Optim

INTEGRATION

Information Server

BigInsights Summary



- **BigInsights = analytical platform for persistent “Big Data”**
 - Based on open source & IBM technologies
 - Managed like a start-up Emphasis on deep customer engagements, product plan flexibility

- **Distinguishing characteristics**
 - Built-in analytics *Enhances business knowledge*
 - Enterprise software integration *Complements and extends existing capabilities*
 - Production-ready platform with tooling for analysts, developers, and administrators. . . . *Speeds time-to-value; simplifies development and maintenance*

- **IBM advantage**
 - Combination of software, hardware, services and advanced research

InfoSphere BigInsights

Platform for **volume, variety, velocity** -- V^3

- Enhanced Hadoop foundation

Analytics for V^3

- Text analytics & tooling

Usability

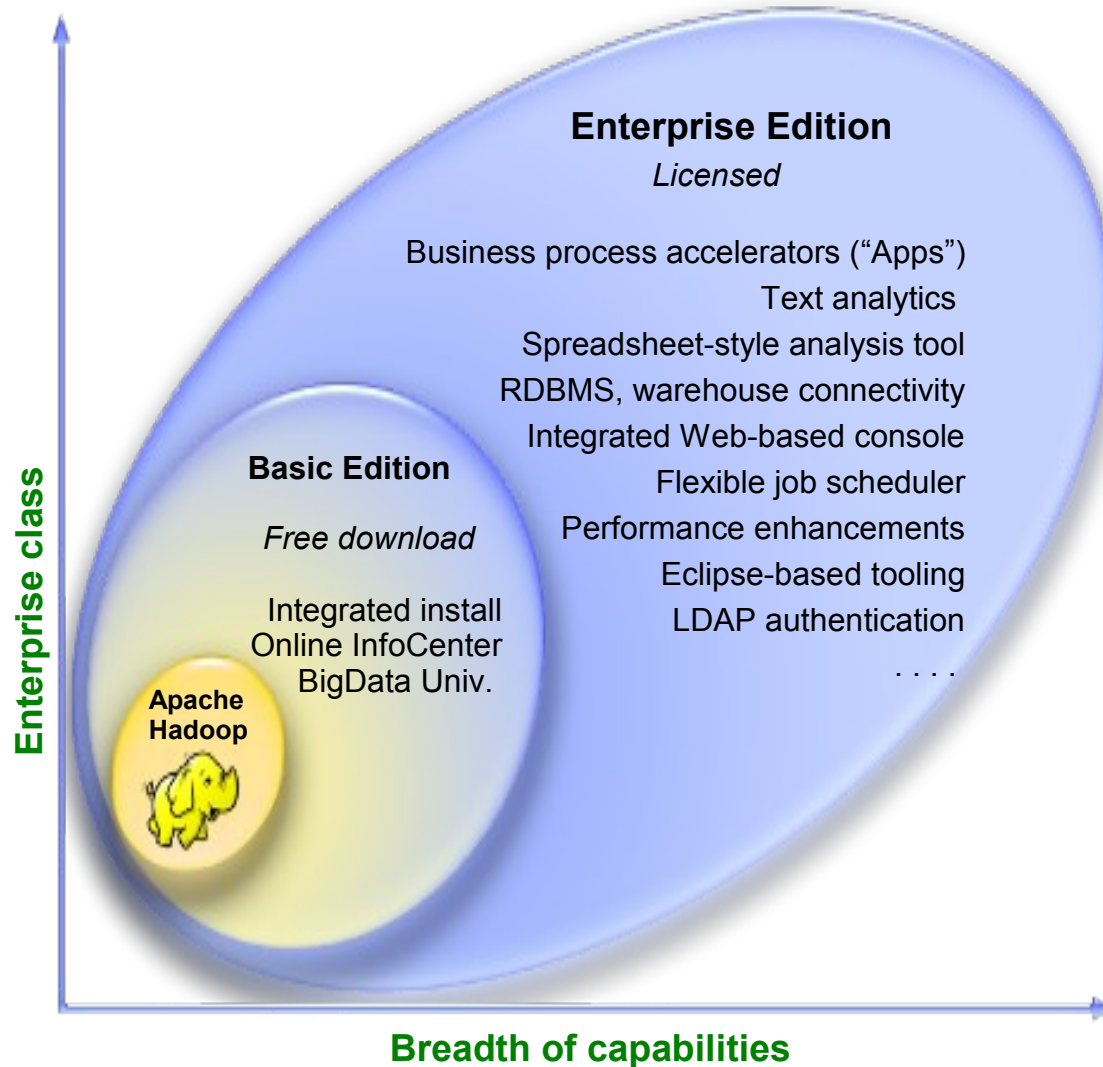
- Web console
- Integrated install
- Spreadsheet-style tool
- Ready-made “apps”

Enterprise Class

- Storage, security, cluster management

Integration

- Connectivity to DB2, Netezza, JDBC databases



Big Data
ON CLOUD

... Hadoop is great at getting the value out of large volumes of data

“94% of Hadoop users perform analytics on large volumes of data not possible before; 88% analyze data in greater detail; while 82% can now retain more of their data”

Ventana Research Benchmark Study “Hadoop and Information Management” 2011

... but it requires large scale compute clusters and lots of storage

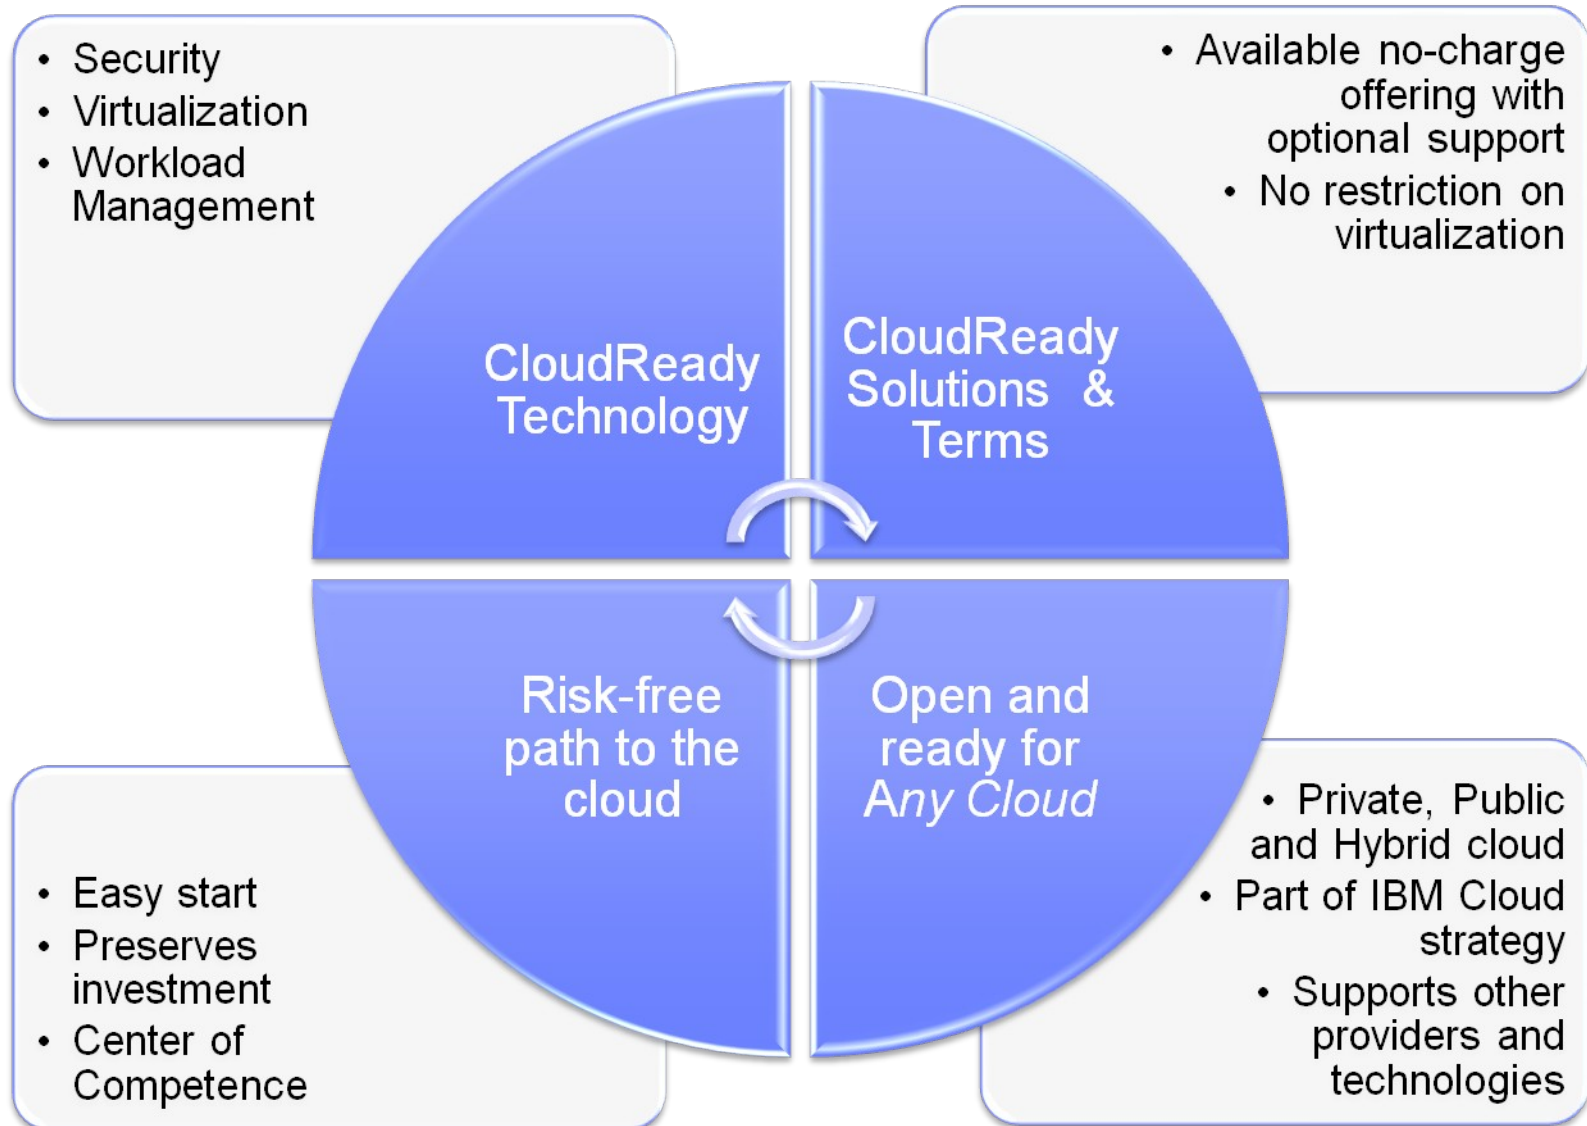
- Yahoo! Has over 40000 nodes running Hadoop managing 180-200 petabytes of data.
- Facebook has 2000 compute nodes with 20 Petabytes of data.
- Average cluster size is 200 servers but most clusters around 30 servers
- Most companies can't make upfront capital investment in compute, storage and networking resources:
 - Makes it hard to evaluate and pilot
 - Difficult to develop real hands on skill
 - Impractical for non continuous needs

VALUE: The Fourth **V** of the Big Data Story

What insight could you gain if you had full use of a 100-node Hadoop cluster for an hour?

What if one hour of this 100-node cluster would cost \$34?

BigInsights Strategy for Cloud Computing



IBM BigInsights: Ready for Any Cloud

IBMSmartCloud
Enterprise

RIGHT SCALE®



cloud.com
now a part of Citrix Systems

 **rackspace**®
HOSTING

EUCALYPTUS

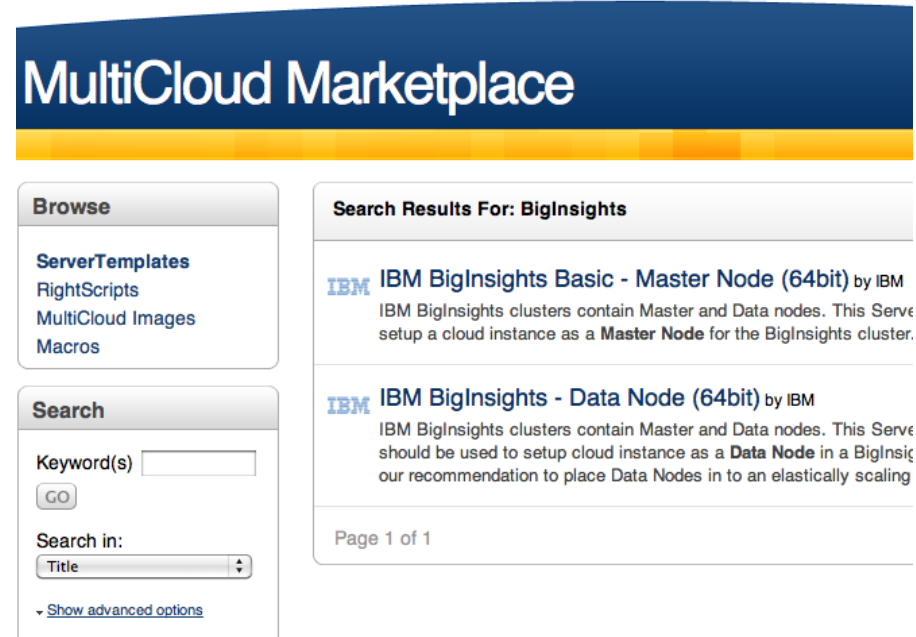
IBM BigInsights on Cloud

Hadoop for everyone ... no upfront investment

- Your own Hadoop cluster on the cloud in less than 30 minutes
- No need to buy install, patch, maintain hardware
- Deploy on Amazon, IBM, Rackspace or your private cloud infrastructure
- Pay as you go for your infrastructure starting at \$0.34/node/hour and only pay for what you actually use
- Use BigInsights Basic Edition at no charge, available low cost support. Seamlessly transition to BigInsights Enterprise Edition when ready

IBM BigInsights and RightScale.com Hadoop on Amazon, Rackspace, Private & Hybrid Cloud Infrastructure

- Your own Hadoop cluster on the cloud in less than 30 minutes
- Deploy to Amazon, Rackspace cloud DCs or on your private cloud
- Pay as you go for your infrastructure starting at \$0.34/node/hour
- Sophisticated cloud resource management by RightScale
- Use BigInsights Basic Edition ... seamlessly transition to BigInsights Enterprise Edition when ready
- Take a free course on BigDataUniversity.com



The screenshot displays the MultiCloud Marketplace interface. At the top, there is a blue header with the text "MultiCloud Marketplace" in white. Below the header, the page is divided into two main sections: "Browse" and "Search".

Browse

- ServerTemplates
- RightScripts
- MultiCloud Images
- Macros

Search

Keyword(s)

Search in:

[Show advanced options](#)

Search Results For: BigInsights

IBM IBM BigInsights Basic - Master Node (64bit) by IBM
IBM BigInsights clusters contain Master and Data nodes. This Server should be used to setup cloud instance as a **Master Node** for the BigInsights cluster.

IBM IBM BigInsights - Data Node (64bit) by IBM
IBM BigInsights clusters contain Master and Data nodes. This Server should be used to setup cloud instance as a **Data Node** in a BigInsights cluster. Our recommendation is to place Data Nodes in to an elastically scaling cluster.

Page 1 of 1

Available Now!

Deploy BigInsights on IBM SmartCloud Enterprise

- Your own Hadoop cluster on the IBM cloud in less than 30 minutes
- No need for hardware, install, patch, maintain
- Locate your Hadoop cluster in one of IBM Cloud data centers WW
- Low hourly charges starting at \$0.30/cluster/hour.
- Explore using BigInsights Basic seamlessly transition to BigInsights Enterprise when ready
- Take a [free course](#) on [BigDataUniversity.com](#)

The screenshot displays the IBM SmartCloud Enterprise user interface. At the top, there is a navigation bar with 'Overview', 'Control panel', 'Account', and 'Support'. Below this, a 'Welcome Leon Katsnelson' message is visible. The main content area features a 'Add an instance to get started' section with three steps: Step 1 (Click Add Instance and select an image), Step 2 (View image details and customize to your needs), and Step 3 (Watch your instances provisioning and start managing). To the right, a 'Summary' section lists 'Your account', 'Active instances', 'Active storage units', and 'Stored images'. Below the steps, there are three small thumbnail images showing different parts of the interface. At the bottom, a 'Compiled notifications' table is shown.

Type	Name	Notification	Date
	testSmall	Storage deleted	Jul 28, 2011 9:18:16 AM
	testMedium	Storage deleted	Jul 28, 2011 9:18:06 AM

Evaluating Big Data Technology

IBM IM Demo Cloud (in limited beta)

- IBM representative initiates the project, invites customer to participate.
- Types of projects: **Demos, Proof of Technology, Proof of Concept:**
- Also great for Beta, Product Introductions, Technology Previews
- Skip construction and go straight to evaluation. No need to get hardware, install/patch OS, BigInsights software, write MapReduce jobs, data sets
- BigInsights sandbox:
 - Dedicated cluster of virtual systems. Not a shared sandbox
 - Customer operates the systems as if local
 - Ability to bring large data sets
 - Data centers in NA, EU and AP

w3 IM Demo Cloud IBM.
About | Benefits | Testimonials | Contact | BluePages | Feedback
Login | Sign up

Demo, PoT, PoC environments in 10 minutes!

Demos, Proof of Technology, Proof of Concept are great for engaging clients and winning deals but they are time consuming and can put a delay in the sales cycle. Customers have to allocate equipment, install, patch and configure the operating system and install and configure our products. Tasks take the time away from getting to the final goal of exploring product features and functions.

IM Demo Cloud does away with traditional inefficient demos, PoT and PoC execution by leveraging the power of Cloud Computing. You don't have to wait for customers to build a demo/PoT/PoC environment. And you don't have to use slow VMWare images. IM Demo Cloud will let you stage a complete demonstration and evaluation environment for your client that can have 1 or 20 servers that demonstrates full capabilities of our offerings. You can guide your customer's evaluation and leave the system with them to continue to explore on their own.

[Register](#)

4 easy steps to your sales success

- 1. Register**
A quick and painless way for customer facing IBMers to register with IM Demo Cloud
- 2. Create**
Identify customer, opportunity and select demos/PoT you would want your customer to participate in
- 3. Nominate**
Nominate your customer contact so that we can create a demo environment and send an invitation to them
- 4. Conduct**
Conduct the demonstration /PoT with your customer and leave the environment for them to continue to explore at their convenience.

[Sign Up to Start Today](#)

BigInsights on the Cloud

Making Learning Hadoop Easy and Fun

- Flexible on-line delivery allows learning @your place and @your pace
- Free courses, free study materials.
- **Cloud-based sandbox for exercises – zero setup**
- 10500+ registered students.
- Hadoop Programming Challenge - 3 students sent to IOD 2011 in Las Vegas, all expenses paid!

The screenshot shows the BigData University website. At the top, there is a navigation bar with the logo "BigData University BETA" and the tagline "Learn from the industry's best". The navigation menu includes HOME, LEARN, DOWNLOAD, RESOURCES, JOBS, and LEARN Hadoop. There are also buttons for "login" and "sign up", and a search bar for "search courses".

The main content area features several sections:

- Why register?**
 - Easy and Affordable:** Learning Hadoop and other Big Data technologies has never been more affordable! Many courses are FREE!
 - Latest industry trends:** Acquire valuable skills and get updated about industry's latest trends right here. Today!
 - Learn from the Experts!** Big Data University offers education about Hadoop and other technologies by the industry's best!
 - Learn at your Own Pace!** Find everything right here when you need it and from wherever you are.
- Hadoop Fundamentals:** A yellow box with the Hadoop logo and text: "Start your career on Hadoop with this FREE course. Learn with hands-on exercises on a Hadoop cluster. Enroll now!"
- Study Made Easy!** A large image of server racks with text: "sign up now for a chance to receive: FREE Books" and a "SIGN UP" button.
- Student Testimonials:** A section titled "Balázs (USA)" with a testimonial about the training material and a "go to sign up" link.

At the bottom right, there are links for "about us", "legal", "contact", and "bug reports".

... in Summary

- Big Data analytics using technologies such as Hadoop can deliver great value but its appeal is somewhat diminished by the extensive capital requirements
- Cloud is a great way to sidestep the CAPEX challenges
- IBM InfoSphere BigInsights, enterprise-ready distribution of Hadoop is available for cloud deployment on Amazon, IBM SmartCloud Enterprise, Rackspace and on private clouds
- IBM helps clients leverage the power of Hadoop and ease of deployment of the cloud. Contact IM Cloud Computing Center of Competence IMCloud@ca.ibm.com.
- BigDataUniversity.com is a great place to start skills acquisition when budget, time or travel is an issue.

Resources

- Use IBM BigInsights on IBM SmartCloud Enterprise
 - <http://www.db2university.com/courses/course/view.php?id=310>
- Use IBM BigInsights on Amazon, Rackspace or your private cloud:
 - <http://www.db2university.com/courses/course/view.php?id=309>
- Learn Hadoop and other Big Data technologies
 - <http://BigDataUniversity.com>
- Get free help from IM Cloud Computing Center of Competence:
 - IMcloud@ca.ibm.com
- Download Free IBM BigInsights Basic Edition
 - <http://www-01.ibm.com/software/data/infosphere/biginsights/basic.html>
- Visit home page for IBM BigInsights on ibm.com:
 - <http://www-01.ibm.com/software/data/infosphere/biginsights/>

Cloud Management Platform
RIGHTSCALE

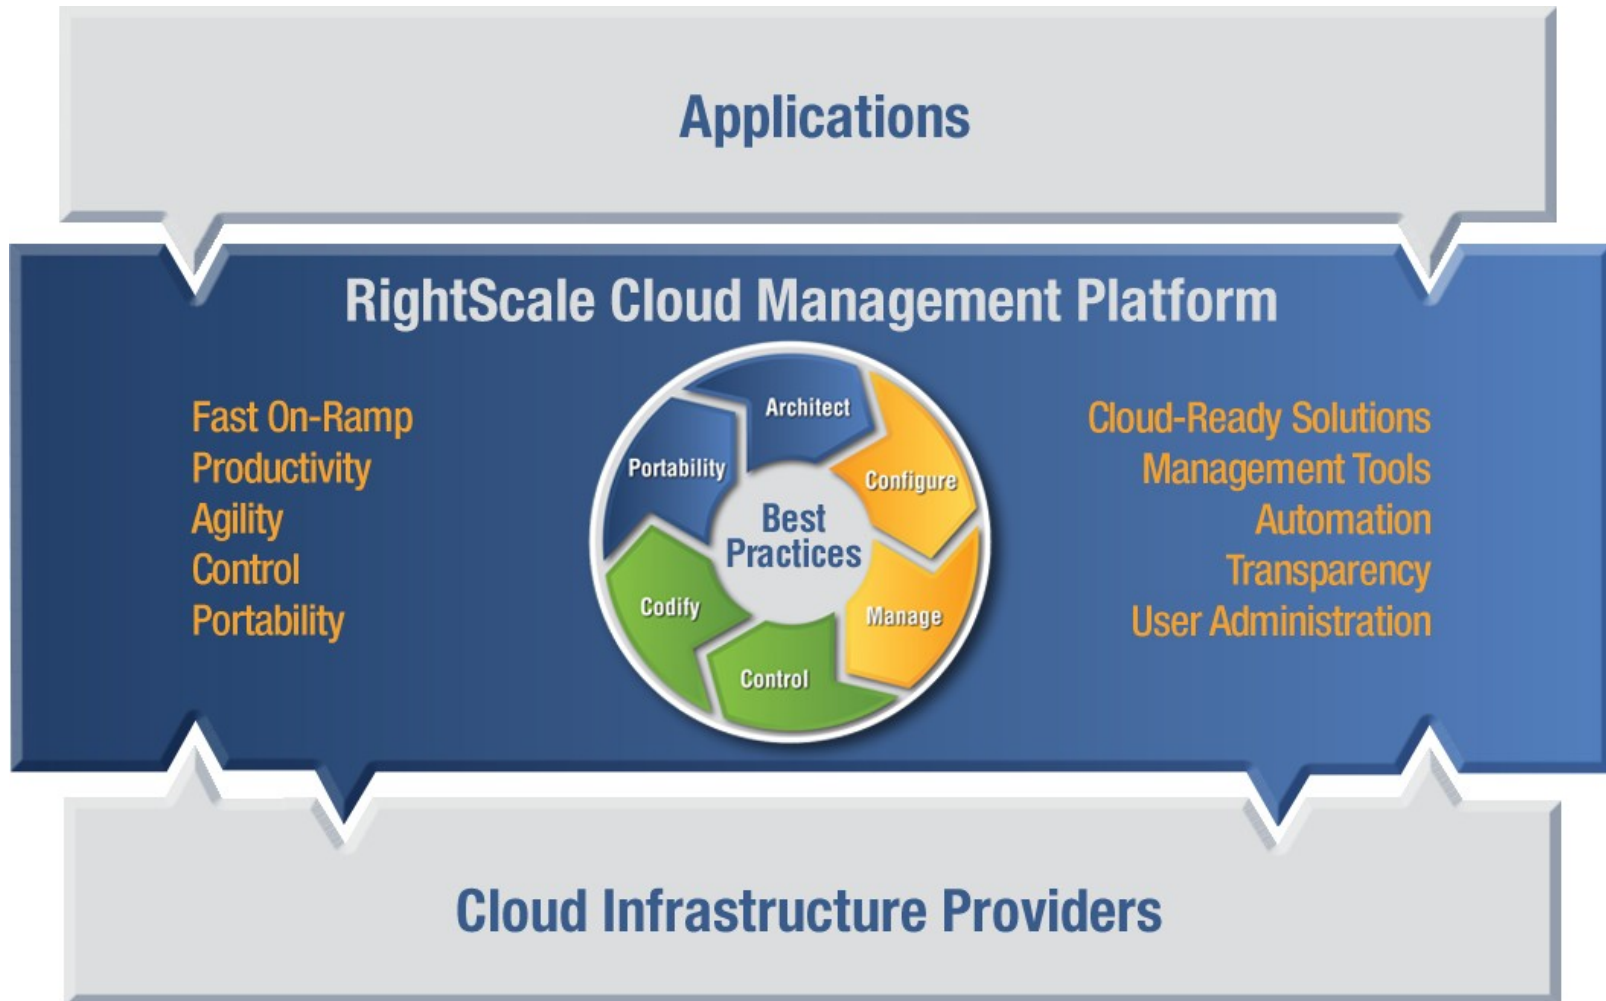
RightScale

Real Customers, Real Deployments, Real Benefits

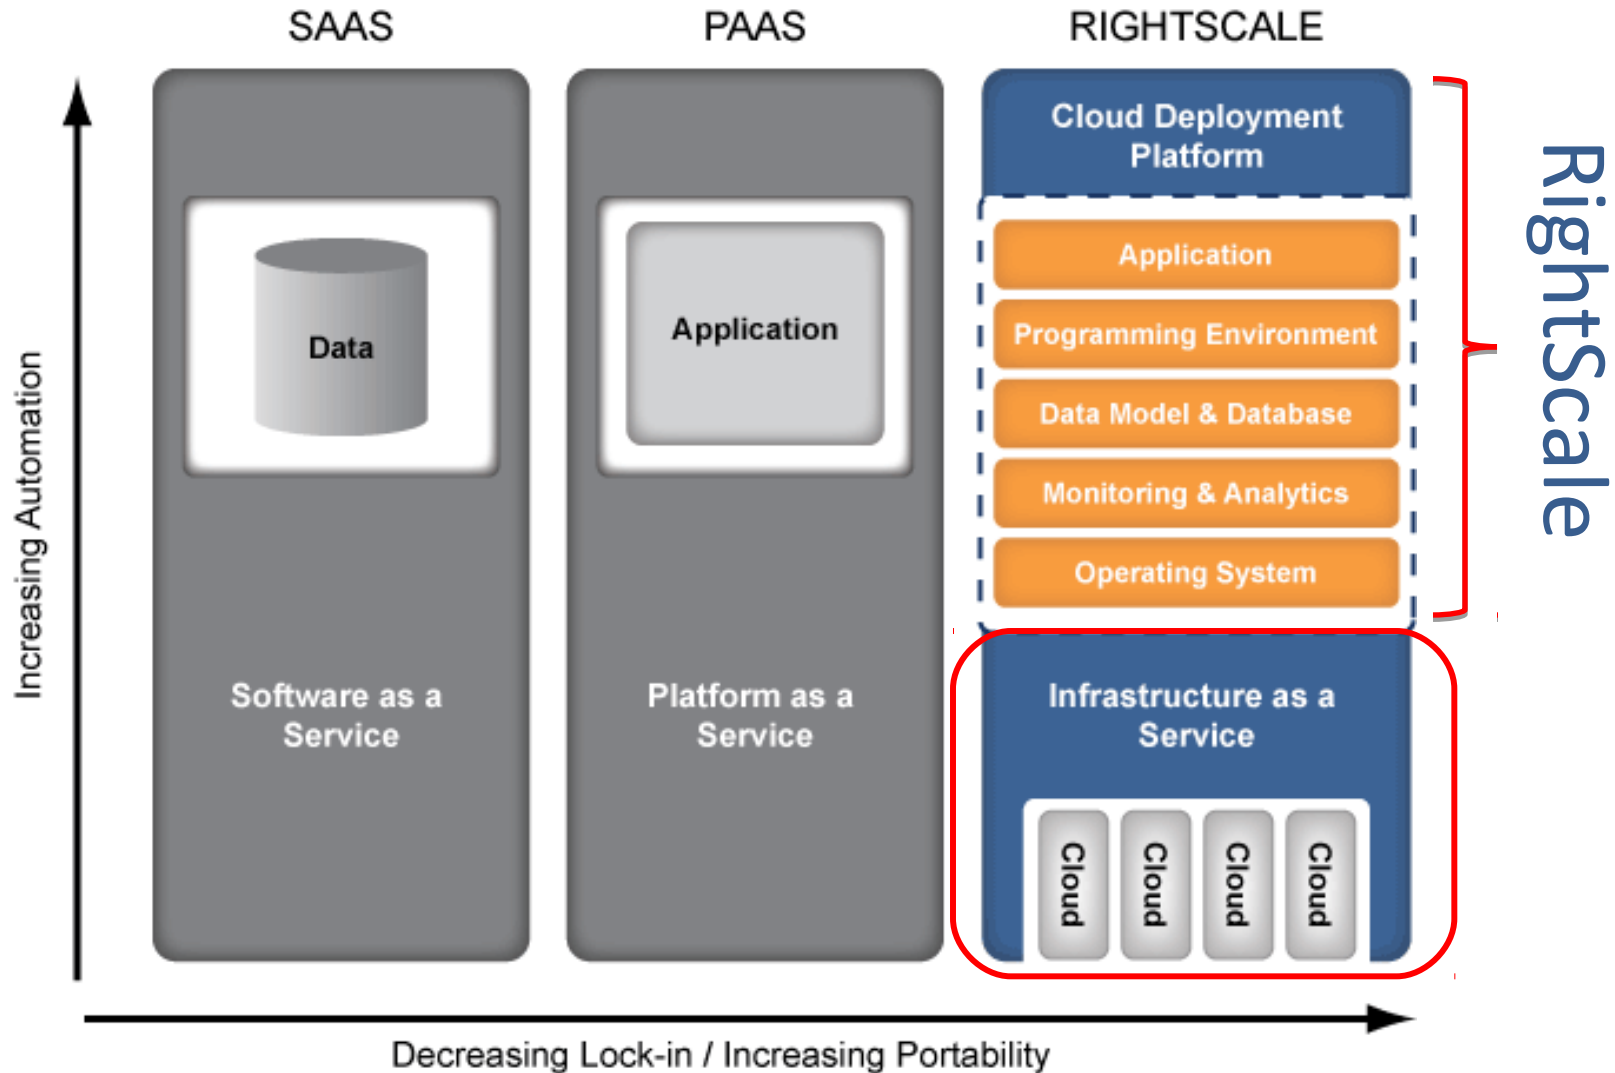
- Managed Cloud Deployments for 4 Years — globally
- More than 45,000 users; launched more than 3MM servers!
- Powering the largest production deployments on the cloud



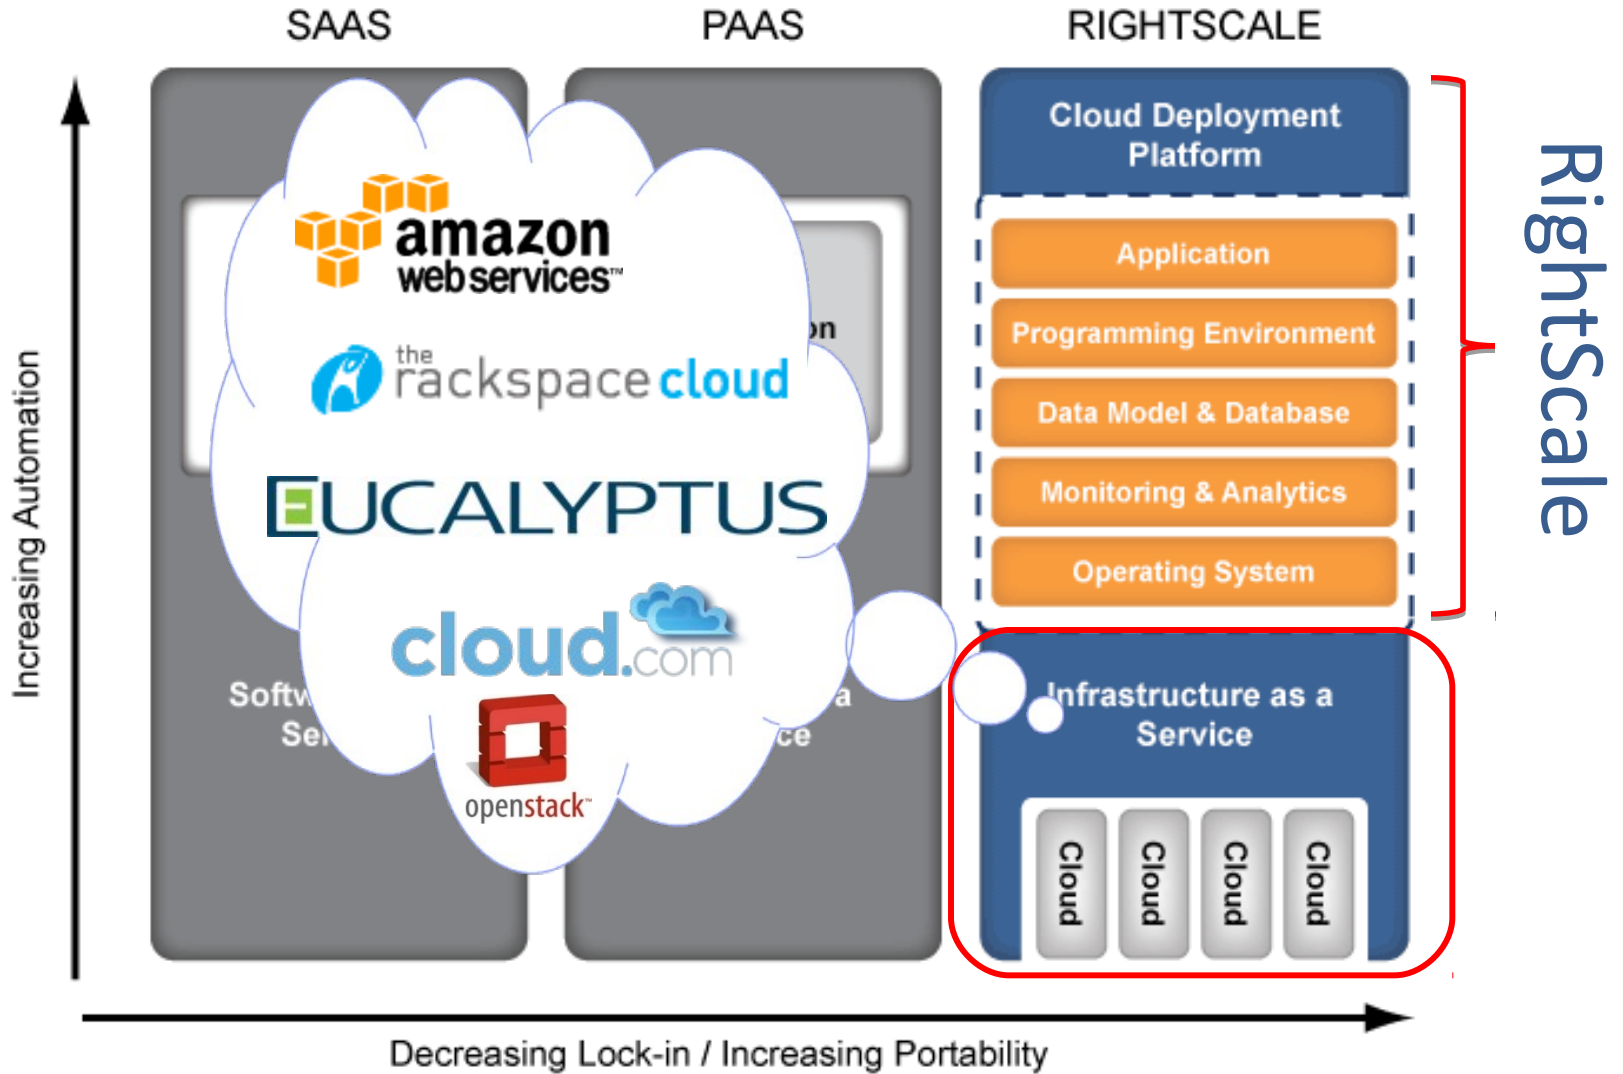
Complete Systems Management



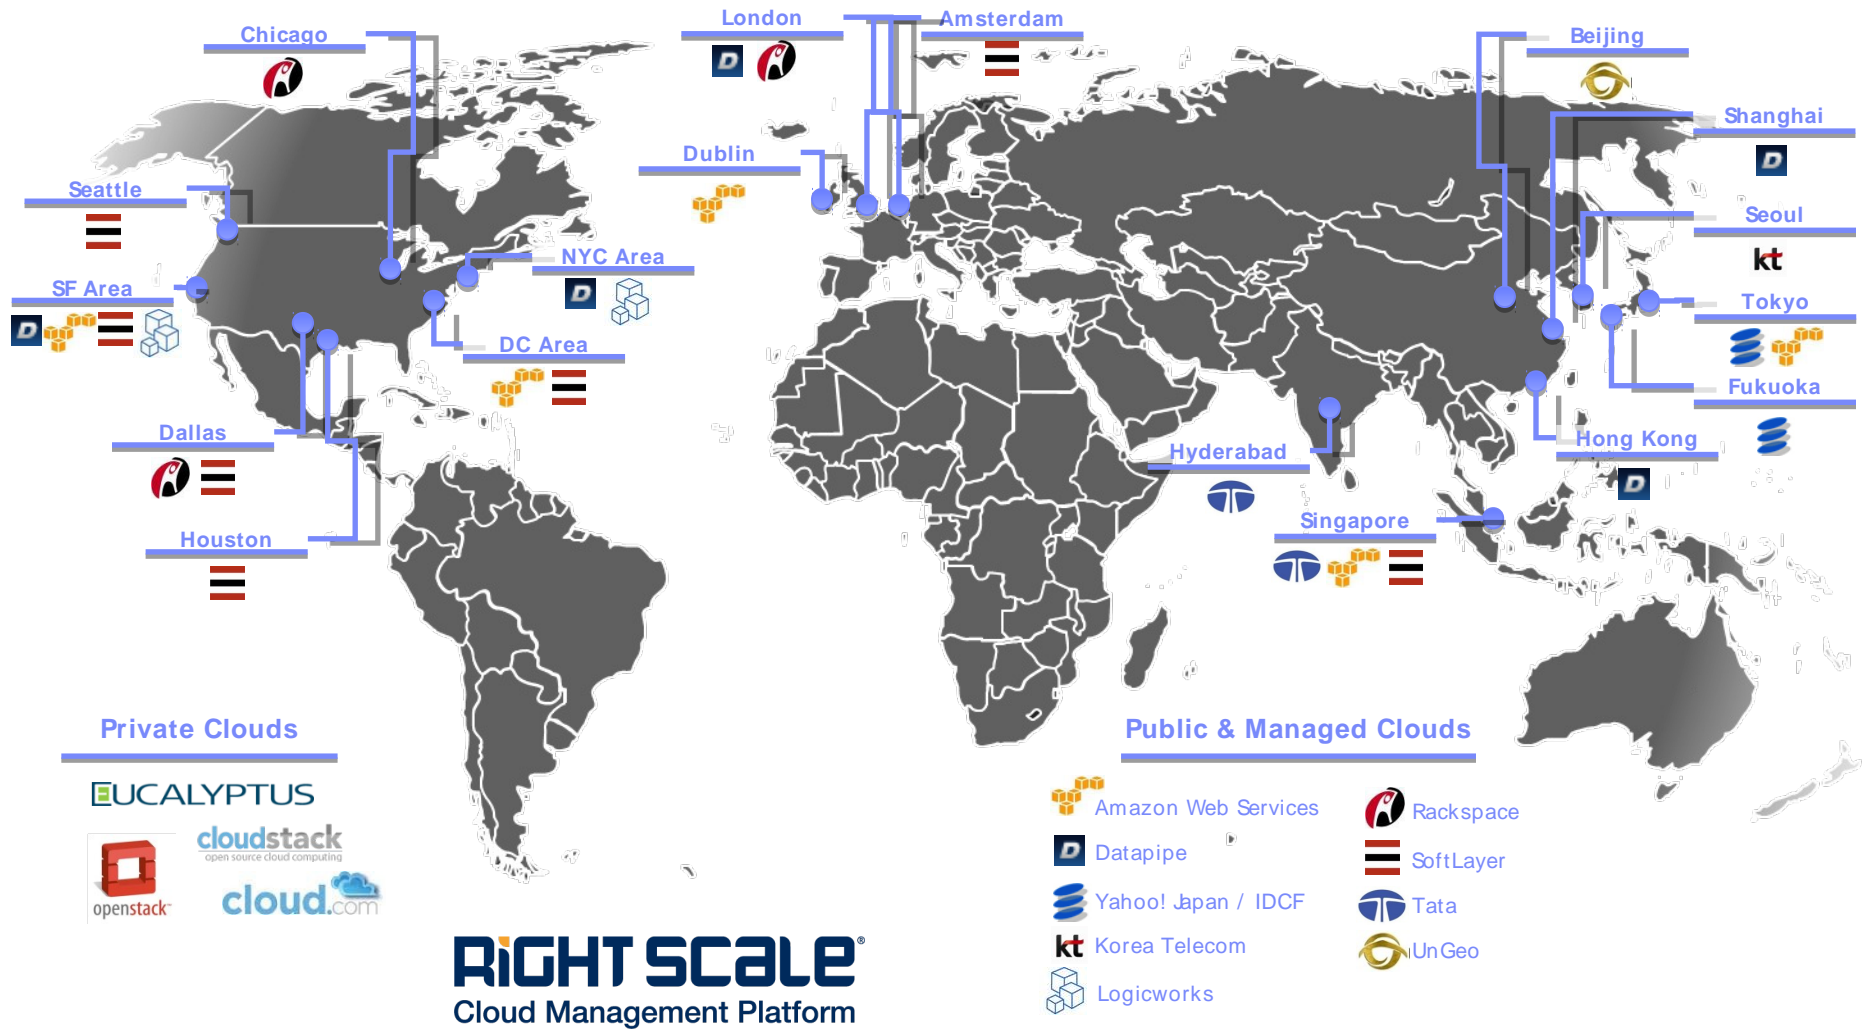
What do we Mean by Cloud Computing?



RightScale Manages IaaS Clouds



Take advantage of many resource pools

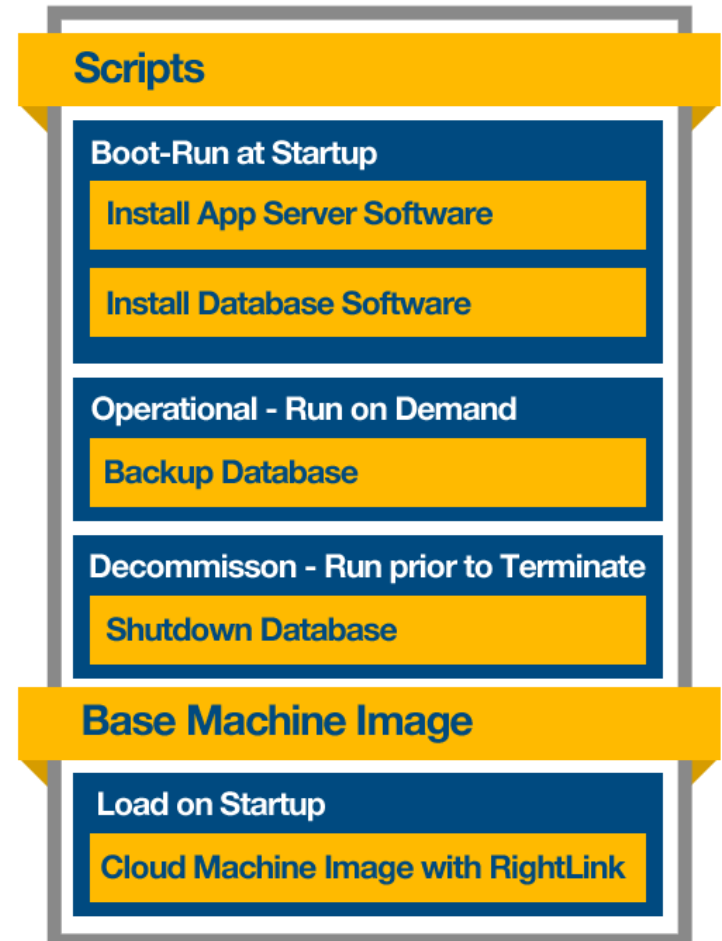


RIGHT SCALE
Cloud Management Platform

ServerTemplates: Reproducible servers on demand

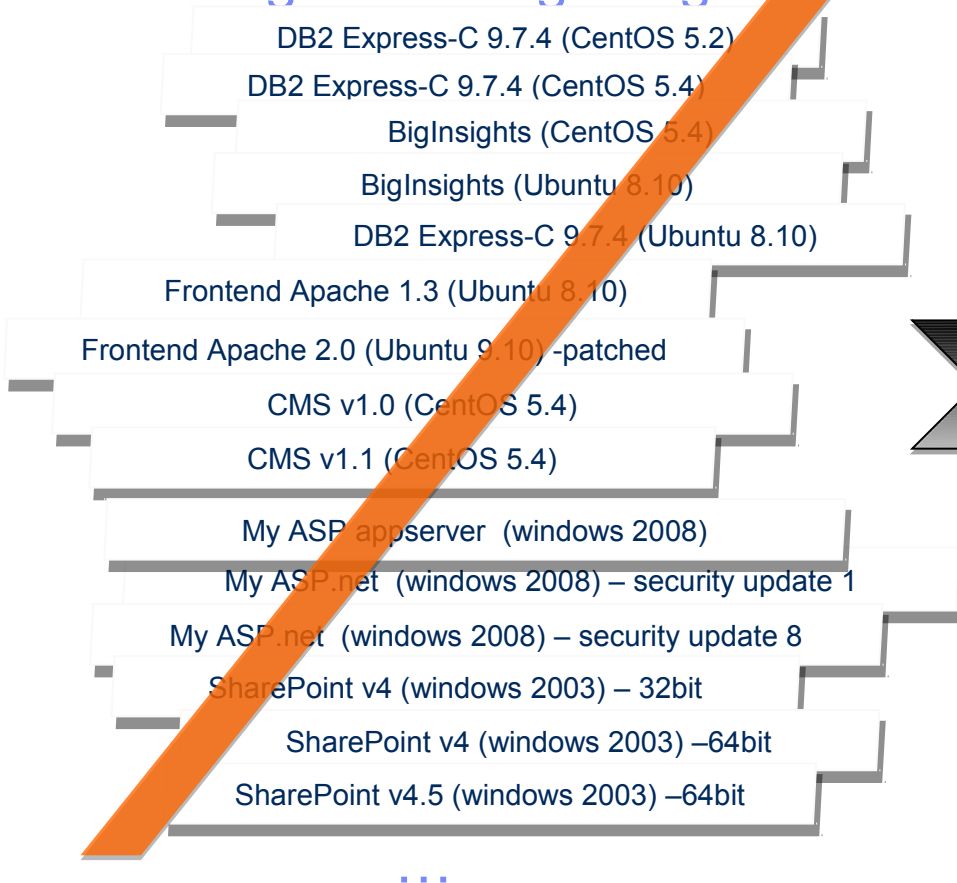
Server Template

- Dynamic configuration
- Abstract role and behavior from cloud infrastructure
- Predictable deployment
- Cloud agnostic / portable
- *Object-oriented programming for sysadmins*

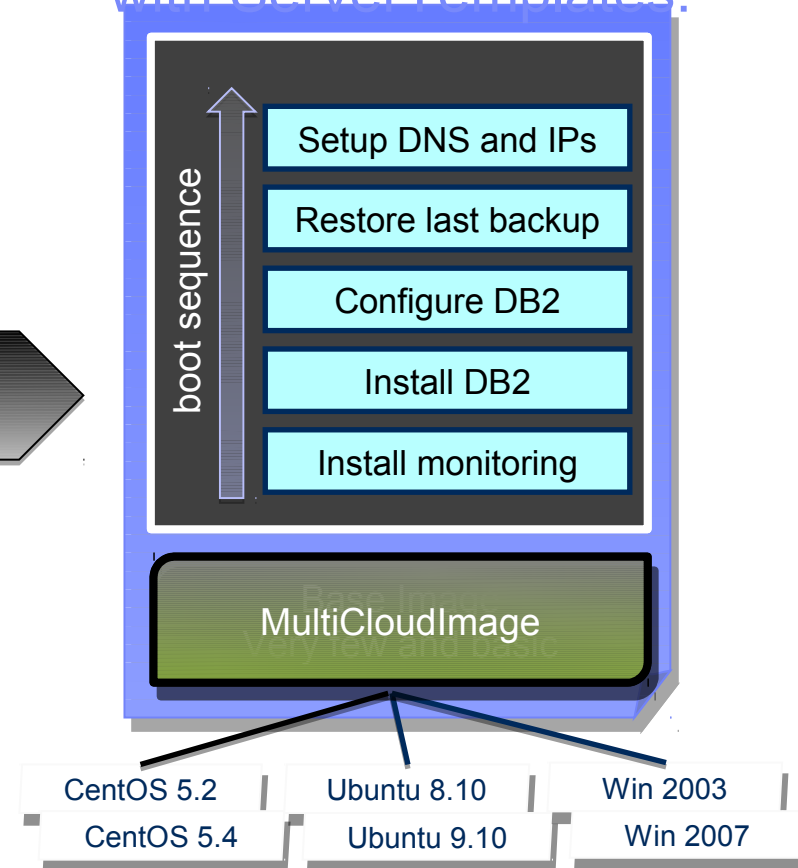


ServerTemplates: Reproducible servers on demand

Configuring servers through bundling images:



Configuring servers with ServerTemplates:



ServerTemplates – Reproducible servers on demand

ServerTemplates
are like Playlists



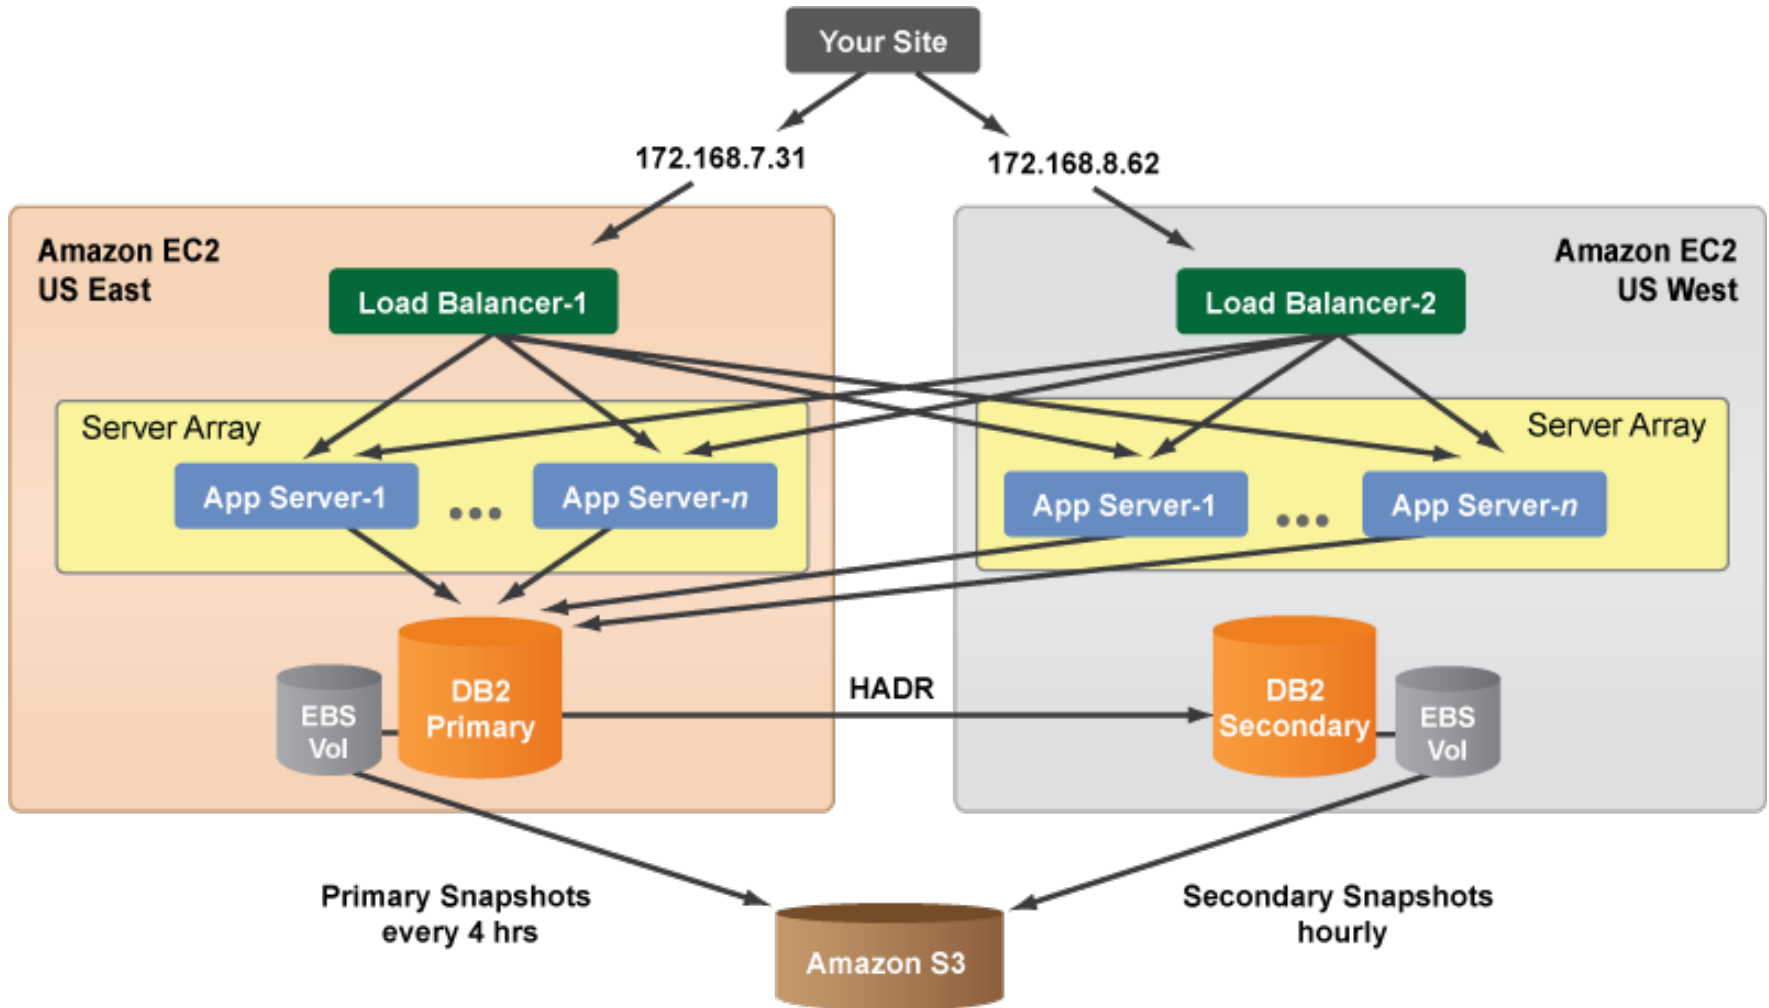
VS.

VMs and Standard Cloud Images
are like burned CDs



Integrated approach that puts together all the parts needed to architect single & multi-server deployments

Not just single servers — complete environments



RightScale Runs in Your Web Browser

The screenshot displays the RightScale dashboard in a web browser. The interface includes a navigation menu with options like 'Manage', 'Design', 'Clouds', 'Reports', and 'Settings'. A search bar and user profile information are visible in the top right. The main content area is divided into several widgets:

- Left Sidebar:** Lists system events such as 'Frontend-01 operational', 'boot completed: SYS Timezone', 'Managed login enabled', 'BigInsights Master Node terminated', 'Decommission requested', 'RAX PHP App #15 completed: app::do_loadbalan', and 'PHP App 1 completed: app::do_loadbalan'.
- Dashboard Header:** Shows 'Overview', 'Deployments', 'Monitoring', and 'Alerts' tabs, along with an 'Add Widget' button.
- Cloud Credential Status:** A table showing the status of services like Amazon Web Services (Online) and Rackspace (cloud controller and cloud account).
- Server State Pie Chart:** A 3D pie chart titled 'Active Server Statistics' showing 18.0 Operational servers (dark orange) and 5.0 Stranded servers (light orange).
- Deployment Budget Estimate:** A table listing various deployment environments and their associated metrics (MTD*) and server counts.

Service	Credential Status
Amazon Web Services	Online
Rackspace	cloud controller
Rackspace	cloud account

Deployment	MTD*	Servers
App Test	-	...10x
Default	-	...56x
DotNet IIS	-	...6x + [in arrays]
DotNet IIS-PERF TEST	-	
Elasticity Web Demo	-	
IBM BigInsights	-	
Load Test Group	-	
Nokia Siemens Networks	-	
Rackspace PHP App 3-Tier	-	... + [in arrays]
Roll Your Own ServerTemplates Demo	-	
Softlayer sandbox	-	
TEST HANA	-	

State	Count
Operational	18.0
Stranded	5.0

Organize Servers in Deployments

RIGHT SCALE® Manage ▾ Design ▾ Clouds ▾ Reports ▾ Settings ▾ Search

Size **Sm** Lg ◀ Hide

BOOKMARKS Edit + Add

QUICKMONITORING Edit + Add

EVENTS Tools ▾

Frontend-01 (3) ▾
 RightScript: 'SYS Monitoring in... <5s
 Managed login enabled <1m
 booting <1m
 RAX PHP App #15 (1) ▾
 completed: app::do_loadbalan 6m
 BigInsights Master Node (1) ▾
 Decommission requested 6m
 PHP App 1 (1) ▾
 completed: app::do_loadbalan ✘
 MySQL 5.1 DB 1 (1) ▾
 completed: db::do_appservers 8m
 dotnet-lb01 (1) ▾
 completed: Email action 8m

Dashboard

Overview | Deployments | Monitoring | Alerts

Add Widget

[October 27th Dashboard Release](#) | [October 7th ServerTemplate Release](#) | Customize the [widgets](#) below!

Cloud Credential Status

Service	Credential Status
Amazon Web Services	● Online
Rackspace	● cloud controller ● cloud account

Deployment Budget Estimate

Deployment

- [App Test](#)
- [Default](#)
- [DotNet IIS](#)
- [DotNet IIS-PERF TEST](#)
- [Elasticity Web Demo](#)
- [IBM BigInsights](#)
- [Load Test Group](#)
- [Nokia Siemens Networks](#)
- [Rackspace PHP App 3-Tier](#)
- [Roll Your Own ServerTemplates Demo](#)
- [Softlayer sandbox](#)

Server State Pie Chart

Active Server Statistics

Stranded

- 17.0 (Operational)
- 5.0 (Stranded)

Organize Servers in Deployments

Reports ▾ Settings ▾ Search Production ▾ Support ▾ Logout

Help ▾

Monitoring Alerts

October 7th ServerTemplate Release | Customize the [widgets](#) below!

Credential Status

- Online
- cloud controller
- cloud account

Deployment	MTD*	Servers
App Test	-	...10x
Default	-	...56x
DotNet IIS	-	...6x + [in arrays]
DotNet IIS-PERF TEST	-	
Elasticity Web Demo	-	
IBM BigInsights	-	
Load Test Group	-	
Nokia Siemens Networks	-	

Server Statistics

Manage Systems, not Servers

Deployment **IBM BigInsights** Help

Clone Archive Delete Add Server Add Array

An IBM BigInsights Cluster Deployment

Info	Servers	Inputs	Monitoring	Scripts	Volumes	Alerts	Audit Entries	History	Changes	Server Defaults
------	----------------	--------	------------	---------	---------	--------	---------------	---------	---------	-----------------

Filter by

Results 1 - 6 of 6 Servers

Deployment is unlocked (Locking prevents deletion of servers)

Deployment Usage

	Current Runrate usage	Month-to-Date usage	Projection usage	Prior Month usage
AWS US-West	\$2.28/hour (\$54.72/day)	--	\$1,655.67	--
Total	\$2.28/hour (\$54.72/day)	--	\$1,655.67	--

Nickname	ServerTemplate	Cloud	CPU	Zone	Runtime	Actions
BigInsights Master Node	Hadoop - IBM BigInsights Basic Edition (Master Node) [rev 11]	AWS US-West	<div style="width: 100%; height: 15px; background-color: green;"></div>	us-west-1c	4 hours	
Data Node 1	Hadoop - IBM BigInsights Basic Edition (Data Node) [rev 9]	AWS US-West	<div style="width: 100%; height: 15px; background-color: green;"></div>	us-west-1c	3 hours	
Data Node 2	Hadoop - IBM BigInsights Basic Edition (Data Node) [rev 9]	AWS US-West	<div style="width: 100%; height: 15px; background-color: green;"></div>	us-west-1c	3 hours	
Data Node 3	Hadoop - IBM BigInsights Basic Edition (Data Node) [rev 9]	AWS US-West	<div style="width: 100%; height: 15px; background-color: green;"></div>	us-west-1c	3 hours	
Data Node 4	Hadoop - IBM BigInsights Basic Edition (Data Node) [rev 9]	AWS US-West	<div style="width: 100%; height: 15px; background-color: green;"></div>	us-west-1c	3 hours	
Data Node 5	Hadoop - IBM BigInsights Basic Edition (Data Node) [rev 9]	AWS US-West	<div style="width: 100%; height: 15px; background-color: green;"></div>	us-west-1c	3 hours	

MultiCloud Marketplace ServerTemplates

Browse

ServerTemplates

RightScripts

MultiCloud Images

Macros

Search

Keyword(s)

GO

Search in:

Title ▾

▾ [Show advanced options](#)

Categories

All

Shared

11H1 Compatible

All-In-One

Application Server

Search Results For: ibm

Sort By Newest First ▾

- | | | |
|--|--|---|
| | <p>Hadoop - IBM BigInsights Basic Edition (Data Node)
by IBM</p> <p>A core component of IBM's platform for Big Data, IBM InfoSphere BigInsights is inspired by, and is compatible with, open source Apache Hadoop and used to store, manage, and gain insights from Internet-scale data at rest.</p> <p>InfoSphere BigInsights Basic Edition v1.9</p> |
★ ★ ★ ★
No Ratings
Nov 22, 2011 |
| | <p>Hadoop - IBM BigInsights Basic Edition (Master Node)
by IBM</p> <p>A core component of IBM's platform for Big Data, IBM InfoSphere BigInsights is inspired by, and is compatible with, open source Apache Hadoop and used to store, manage, and gain insights from Internet-scale data at rest.</p> <p>InfoSphere BigInsights Basic Edition v1.9</p> |
★ ★ ★ ★
No Ratings
Nov 18, 2011 |
| | <p>IBM DB2 Express-C 9.7.4 by IBM</p> <p>IBM DB2 Express-C is the "no-limits" free edition of the enterprise-class IBM database server. It features advanced technologies such as pureXML, a hybrid relational/XML database engine. It's easy-to-use, includes self-managing capabilities, and offers enterprise-proven performance and scalability.</p> |
★ ★ ★ ★
2 Ratings
Nov 15, 2011 |
| | <p>BigData University - IBM BigInsights Basic Edition (Master Node) by IBM</p> <p>A core component of IBM's platform for Big Data, IBM InfoSphere BigInsights is inspired by, and is compatible with, open source Apache Hadoop and used to store, manage, and gain insights from Internet-scale data at rest.</p> |
★ ★ ★ ★
No Ratings
Sep 22, 2011 |

MultiCloud Marketplace Macros

Browse

ServerTemplates
RightScripts
MultiCloud Images
Macros

Search

Keyword(s)

GO

Search in:

Publisher ▾

▾ [Show advanced options](#)









Categories

All

Shared

Search Results For: ibm

Sort By Newest First ▾

- | | | |
|--|--|--|
|  | Create BigInsights Cluster Deployment by IBM | Oct 14, 2011  |
| A core component of IBM's platform for Big Data, IBM InfoSphere BigInsights is inspired by, and is compatible with, open source Apache Hadoop and used to store, manage, and gain insights from Internet-scale data at rest. | | |
|  | Create BigInsights Deployment - BigData University
by IBM | Sep 22, 2011  |
| A core component of IBM's platform for Big Data, IBM InfoSphere BigInsights is inspired by, and is compatible with, open source Apache Hadoop and used to store, manage, and gain insights from Internet-scale data at rest.
<small>InfoSphere BigInsights Basic Edition, IBM entry-level offering, is available.</small> | | |
|  | Create BigInsights (Hadoop) Cluster Deployment - Simple two node setup by IBM | Jul 14, 2011  |
| A core component of IBM's platform for Big Data, IBM InfoSphere BigInsights is inspired by, and is compatible with, open source Apache Hadoop and used to store, manage, and gain insights from Internet-scale data at rest. | | |
|  | Setup Highly Available DB2 database cluster (BETA 1)
by IBM | Apr 12, 2011  |
| Create a cluster of 2 DB2 servers in a high availability configuration (HADR). | | |

Free Training at Your Own Pace

- RightScale is a Sponsor of Big Data University
- Free Course on Hadoop in the Cloud Includes:
 - RightScale
 - IBM BigInsights
 - Amazon EC2 \$25 credits (limited quantity)
- Go to: BigDataUniversity.com
- Course Code: BD005EN
- Direct Link: <http://www.bigdatauniversity.com/courses/course/view.php?id=309>

Resources – Learn More

- Web Resources:
 - RightScale.com/webinars
 - RightScale.com/whitepapers
 - Vimeo.com/RightScale
 - IBM.com/smartercomputing
 - BigDataUniversity.com
- Follow presenters on Twitter
 - @uribudnik
 - @katsnelson

Contact Us

RightScale

sales@rightscale.com

1-866-720-0208

IBM IM Cloud Team

IMcloud@ca.ibm.com

> Questions



Additional / Backup Slides

Two Really “Hot” Areas: Cloud and Data

- “Big Data will earn its place as the next ‘must have’ competency in 2012” (IDC)
- “Through 2015, more than 85 percent of Fortune 500 organizations will fail to effectively exploit big data for competitive advantage” (Gartner).
- “2012 is likely to be a busy year for Big Data-driven mergers and acquisitions” (IDC)
- “[In 2012,] 80% of new commercial enterprise apps will be deployed on cloud platforms” (IDC)
- IP traffic over data center networks will reach 4.8 zettabytes a year by 2015, and cloud computing will account for one-third of it, or 1.6 zettabytes. (CISCO)
- “Amazon Web Services [will] exceed \$1 billion in cloud services business in 2012 with Google’s Enterprise business to follow within 18 months” (IDC).

Merging the Traditional and Big Data Approaches

Traditional Approach

Structured & Repeatable Analysis

Business Users

Determine what question to ask



IT

Structures the data to answer that question




Monthly sales reports
Profitability analysis
Customer surveys

Big Data Approach

Iterative & Exploratory Analysis

IT

Delivers a platform to enable creative discovery



Business

Explores what questions could be asked



Brand sentiment
Product strategy
Maximum asset utilization

BigInsights Content

Function	Version	Basic Edition	Enterprise Edition
Integrated Install*		Inc	Inc
Hadoop (including common utilities, HDFS, MapReduce framework)	0.20.2	Inc	Inc
Jaql (programming / query language)	0.5.2	Inc	Inc
Pig (programming / query language)*	0.8.1	Inc	Inc
Flume (data collection/aggregation)	0.9.1	Inc	Inc
Hive (data summarization/querying)*	0.7.1	Inc	Inc
Lucene (text search)*	3.3.0	Inc	Inc
Zookeeper (process coordination)	3.3.3	Inc	Inc
Avro (data serialization)	1.5.1	Inc	Inc
HBase (real time read/write)*	0.90.4	Inc	Inc
Oozie (workflow/ job orchestration)*	2.3.1	Inc	Inc
Online documentation*		Inc	Inc
Capability to integrate with JDBC sources through general-purpose Jaql module		Inc	Inc
Capability to integrate with DB2, InfoSphere Warehouse (DB2 UDF samples to submit jobs, and read results from BigInsights)		Inc	Inc

BigInsights Content (cont'd)

Function	Version	Basic Edition	Enterprise Edition
Capability to integrate with R (Jaql module to invoke R statistical capabilities from BigInsights)		n/a	Inc
Capability to integrate with Netezza, DB2 LUW with DPF from Jaql		n/a	Inc
LDAP authentication and additional security features*		n/a	Inc
Integrated Web Console*		n/a	Inc
Integrated workflow capabilities and flexible job scheduler		n/a	Inc
Platform performance enhancements (Adaptive MapReduce, efficient processing of compressed text files, large-scale text indexing, etc.)*		n/a	Inc
Text analytics		n/a	Inc
Eclipse plugins for text analytic development, Jaql, Hive, Java*		n/a	Inc
Ready-made “apps” for data import/export, Web crawl, Boardreader., etc. *		n/a	Inc
Web-based application catalog*		n/a	Inc
Spreadsheet-like analytical tool *		n/a	Inc
IBM support		Opt	Inc
Unlimited storage		n/a	Inc

BigInsights: Value Beyond Open Source



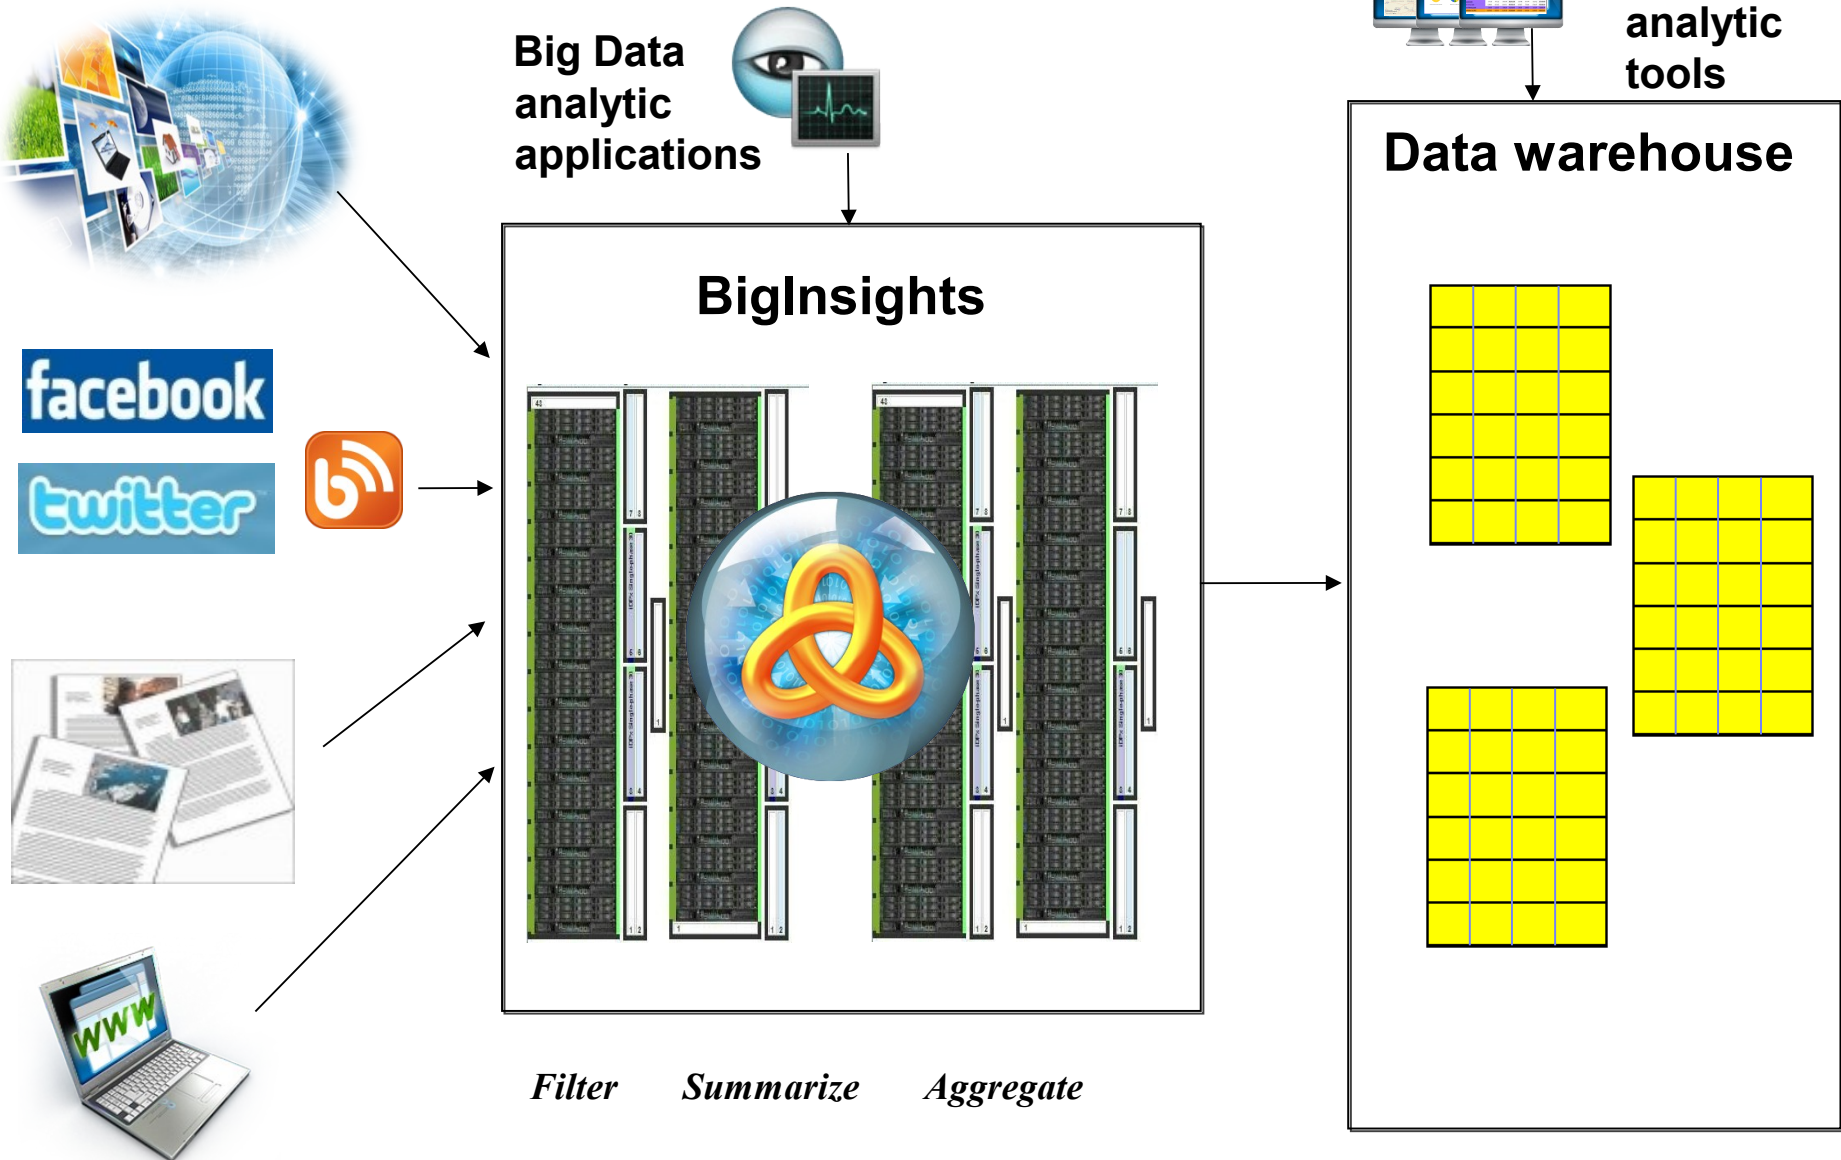
■ Technical differentiators

- Built-in analytics
 - Text processing engine, annotators, Eclipse tooling
 - Interface to project R (statistical platform)
- Enterprise software integration (DBMS, warehouse)
- Spreadsheet-style analytical tool for analysts
- Ready-made business process accelerators
- Integrated installation of supported open source and IBM components
- Web Console for administration and application access
- Platform enrichment: additional security, performance features, . . .
- Standard IBM licensing agreement and world-class support
- More to come in future releases!

■ Business benefits

- Quicker time-to-value due to IBM technology and support
- Reduced operational risk
- Enhanced business knowledge with flexible analytical platform
- Leverages and complements existing software assets

BigInsights and the data warehouse



BigInsights and the data warehouse

