Hadoop Patterns & Practices

11.30.2012

George Chu (**朱金生**)
Sr. Director of Engineering
Hadoop, Cloud Services & Mobile

YAHOO!®

Hadoop Today

About Yahoo!

- Aspire to make the world's daily habits more inspiring & entertaining
- Focus on building highly personalized experiences that connect people to what matters most to them
- Connect advertisers and partners with the audiences who build their business

Email – 35M hours per day, 190M user engagements per day

News & **Information** – Personalized news, sports, finance, weather, etc.

Photos – Reimaging Flickr to make it faster, more beautiful & social

Search – Innovating search across platforms (Direct Display, Axis)

Personalization – 13M different versions of the Homepage tailored for users' distinct interests

Across platforms – Online, mobile, TV, second screen

Trusted & **transparent**

Personalized user experiences help advertisers target the people that matter

Global scale

Innovative platforms (mobile, ad networks, video, media, partnerships, etc.)

Hadoop at Yahoo!

Hadoop is behind every click at **Yahoo!**, turning data into insights and making content and ads relevant for our consumers. You will quickly learn from this presentation about how **Yahoo!** is leveraging the cloud, scaling the core, and expanding the ecosystem.

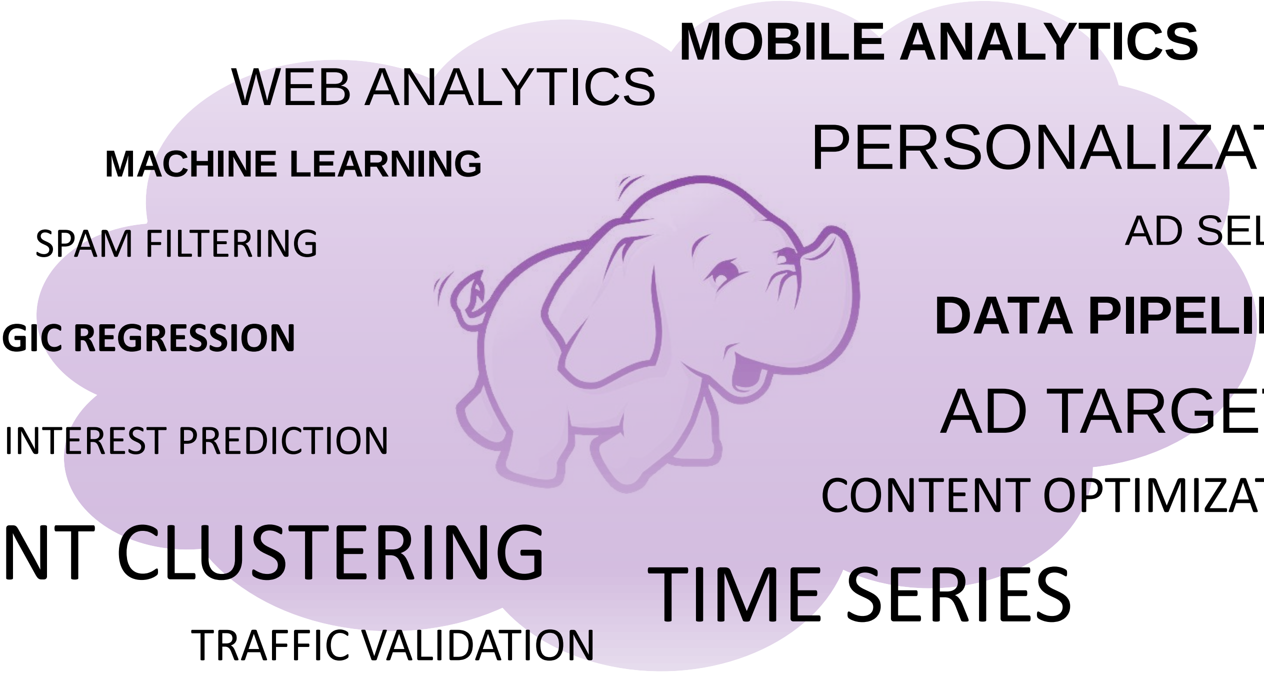

MOBILE ANALYTICS

WEB ANALYTICS

PERSONALIZATION

MACHINE LEARNING

AD SELECTION

SPAM FILTERING

DATA PIPELINES

LOGIC REGRESSION

AD TARGETING

USER INTEREST PREDICTION

CONTENT OPTIMIZATION

CONTENT CLUSTERING

TIME SERIES

TRAFFIC VALIDATION

Yahoo!'s commitments to Hadoop & community

Leveraging the Cloud

Yahoo! operates one of the world's largest private cloud infrastructure, handling **15B page visits per month** & **100B events per day** from more than **700M unique monthly users**. Content handled by the Yahoo! Cloud has grown to **more than 200PBs**, with **50TBs of additional data collected daily**.

The Hadoop project is an integral part of Yahoo!'s cloud infrastructure, and is at the heart of many of Yahoo!'s important business processes like Yahoo!'s next generation display advertising system, which **optimizes forecasting and pricing for over 24B ads served daily.**

Hadoop at Yahoo! works in concert with Yahoo!'s other cloud services such as **Edge Services** (built with standard Yahoo! Technologies like **Yahoo! Traffic Server), Data Servicing Containers, Distributed Structured** & **Unstructured Storage Services**, and **Data Highway** (Yahoo!'s event collection & delivery platform) for collecting, storing, processing, managing and analyzing data.

Scaling the Core

Yahoo!'s technical leadership has taken Hadoop from a science project to a mainstream big data technology serving thousands of companies around the world. Yahoo! continues to be a key contributor across all areas of Hadoop, including **Hadoop 0.23** with next generation **MapReduce** and **HDFS federation**.

Yahoo! currently operates the largest production deployment of Hadoop clusters in the world made up of **over 42,000 servers** which process **10.8M jobs per month** and **140PBs of data**, stretching the scale and stability limits of core Hadoop MapReduce and HDFS.

Like **GridMix** (benchmark for Hadoop clusters) and **Vaidya** (rule-based performance diagnostics), Yahoo! continues to develop key tools that improve the overall stability and usability of Hadoop. Tools in the pipeline include **Groundhog** (Pig record and playback for regression testing), **QuAREH** (replay workloads and model utilization), **C3** (compute capacity calculator), and **Anarchy Ape** (fault injection into Hadoop clusters).

Expanding the Ecosystem

Yahoo! continues to make the Hadoop ecosystem stronger, working closely with key collaborators in the Hadoop community and helping to drive more users and contributors to Hadoop. Yahoo! remains significantly invested in code, resources, and adoption of technology to further ensure a strong and vibrant Hadoop community.

Yahoo! has contributed a majority of the current Hadoop code and other related projects such as Pig (language to express data transformation), **Oozie** (workflow scheduling & coordination system), **Zookeeper** (centralized service for highly reliable distributed coordination), and **HCatalog** (unified table & schema management).

Yahoo! has also adopted other Hadoop stack components such as **Hive** and **HBase** from the open source community to solve additional use-cases, and looks forward to extending the capabilities of these components and contributing the development back to the community.

Hadoop Operational Statistics | May, 2012

Number of (centrally-managed) nodes	**42,000**
Cluster types	**Production** **Research** (Ad-hoc usage) **Sandbox** (Release validation) **Innovation** (Dev, QE, Benchmarking) **Data Loading**
Maximum nodes per cluster	**4,000**
HDFS	**>350** petabytes
Compute Slots (Map and Reduce Slots)	**>9M** slot hours available / day
Jobs submitted	**>10M** / month
Number of monthly unique users (submitting jobs in a month)	**>1000** / month
Average daily unique users (submitting jobs each day)	**>300** / day

Hadoop usage at Yahoo!

Search Assist

Problem

Related concepts appear close together in text corpus to assist users with search term

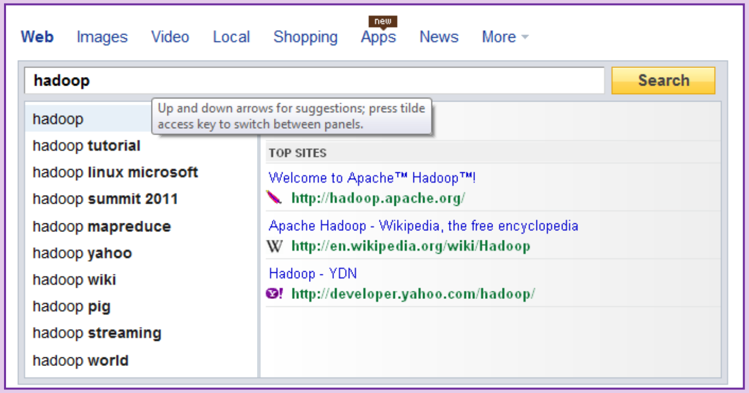

hadoop

Up and down arrows for suggestions; press tilde access key to switch between panels.

- hadoop
- hadoop tutorial
- hadoop linux microsoft
- hadoop summit 2011
- hadoop mapreduce
- hadoop yahoo
- hadoop wiki
- hadoop pig
- hadoop streaming
- hadoop world

TOP SITES

Welcome to Apache™ Hadoop™!
http://hadoop.apache.org/

Apache Hadoop - Wikipedia, the free encyclopedia
http://en.wikipedia.org/wiki/Hadoop

Hadoop - YDN
http://developer.yahoo.com/hadoop/

Solution

Hadoop helps Yahoo! process 1B web pages of about 10K bytes each (10 TB of input data) to create the output list of related words

Behavioral Targeting

Problem

Quickly make complex decisions to serve the right ads to the right customer by targeting billions of impressions per day across one of the largest ad networks in the world

Solution

Hadoop helps Yahoo! process declared data and recent activity to segment users and determine the right ads to serve in milliseconds

Hadoop usage at Yahoo! | Continued

Mail Anti-Spam

Problem
Yahoo! Mail delivers 5.6 billion email a day across 300 million mailboxes. Users want to see emails from friends and family in the inboxes, from the people who matter the most… not from spammers and phishers

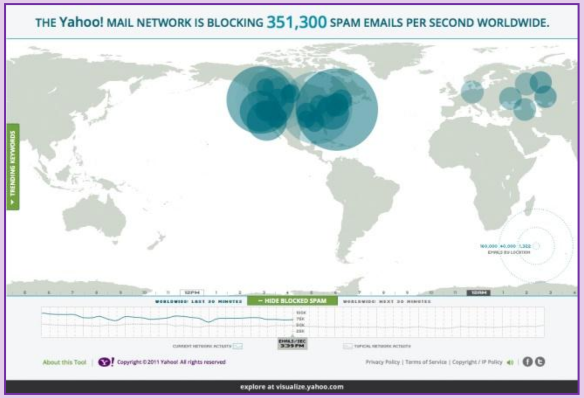
THE Yahoo! MAIL NETWORK IS BLOCKING 351,300 SPAM EMAILS PER SECOND WORLDWIDE.

explore at visualize.yahoo.com

Solution
Hadoop helps Yahoo! block 20.5 billion spam emails per day through machine learning on the grid. SpamGuard in conjunction with Hadoop has reduced spam by 60%

Membership Anti-Abuse

Problem
Membership processes 2.22M new registrations (127 M logins) every day! Abuse taints metrics, and parsing out abusive vs. legitimate user is an ongoing challenge

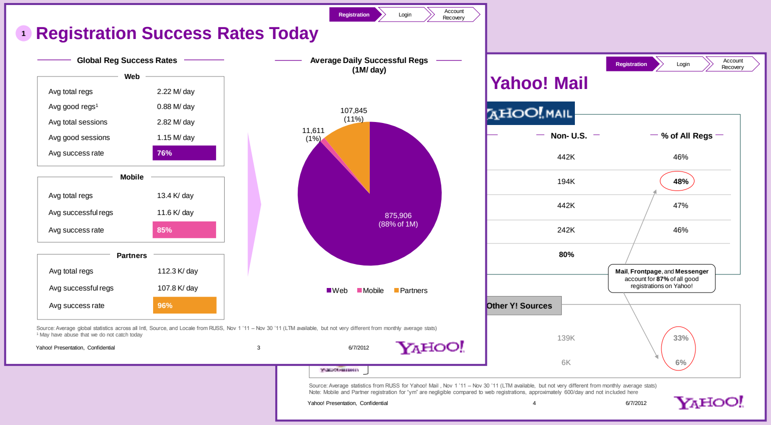

1 Registration Success Rates Today

Global Reg Success Rates

Web	
Avg total regs	2.22 M/ day
Avg good regs[1]	0.88 M/ day
Avg total sessions	2.82 M/ day
Avg good sessions	1.15 M/ day
Avg success rate	76%

Mobile	
Avg total regs	13.4 K/ day
Avg successful regs	11.6 K/ day
Avg success rate	85%

Partners	
Avg total regs	112.3 K/ day
Avg successful regs	107.8 K/ day
Avg success rate	96%

Average Daily Successful Regs (1M/ day)

Source: Average global statistics across all Intl, Source, and Locale from RUSS, Nov 1 '11 – Nov 30 '11 (LTM available, but not very different from monthly average stats)
[1] May have abuse that we do not catch today

Yahoo! Presentation, Confidential 3 6/7/2012

Yahoo! Mail

YAHOO! MAIL

Non- U.S.	% of All Regs
442K	46%
194K	48%
442K	47%
242K	46%
80%	

Mail, Frontpage, and Messenger account for 87% of all good registrations on Yahoo!

Other Y! Sources

139K	33%
6K	6%

Source: Average statistics from RUSS for Yahoo! Mail , Nov 1 '11 – Nov 30 '11 (LTM available, but not very different from monthly average stats)
Note: Mobile and Partner registration for "ym" are negligible compared to web registrations, approximately 600/day and not included here

Yahoo! Presentation, Confidential 4 6/7/2012

Solution
Hadoop helps Yahoo! detect abusive registrations through machine learning on the grid

Hadoop usage at Yahoo! | Continued

Content Agility

Problem

Properties had siloed approaches for CMS, front-end development and editorial. A common solution was needed for the entire content network to bring agility to Yahoo! properties

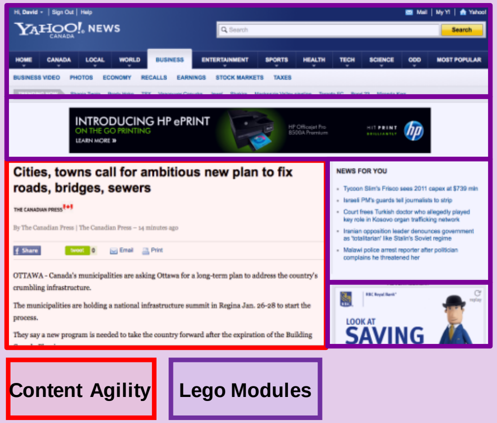

Hi, David | Sign Out | Help   Mail | My Y! | Yahoo!

YAHOO! NEWS CANADA   Search

HOME CANADA LOCAL WORLD BUSINESS ENTERTAINMENT SPORTS HEALTH TECH SCIENCE ODD MOST POPULAR

BUSINESS VIDEO PHOTOS ECONOMY RECALLS EARNINGS STOCK MARKETS TAXES

INTRODUCING HP ePRINT ON THE GO PRINTING LEARN MORE »

Cities, towns call for ambitious new plan to fix roads, bridges, sewers

THE CANADIAN PRESS

By The Canadian Press | The Canadian Press – 14 minutes ago

Share Tweet 0 Email Print

OTTAWA - Canada's municipalities are asking Ottawa for a long-term plan to address the country's crumbling infrastructure.

The municipalities are holding a national infrastructure summit in Regina Jan. 26-28 to start the process.

They say a new program is needed to take the country forward after the expiration of the Building

NEWS FOR YOU

- Tycoon Slim's Frisco sees 2011 capex at $739 mln
- Israeli PM's guards tell journalists to strip
- Court frees Turkish doctor who allegedly played key role in Kosovo organ trafficking network
- Iranian opposition leader denounces government as 'totalitarian' like Stalin's Soviet regime
- Malawi police arrest reporter after politician complains he threatened her

LOOK AT SAVING

Content Agility   Lego Modules

Solution

Leverage Content Agility as the single, grid-based, highly scalable CMS. Lego provides reusable UI modules and shared tools

Personalization

Problem

Increase engagement by showing the right content to users with input from science & human editors

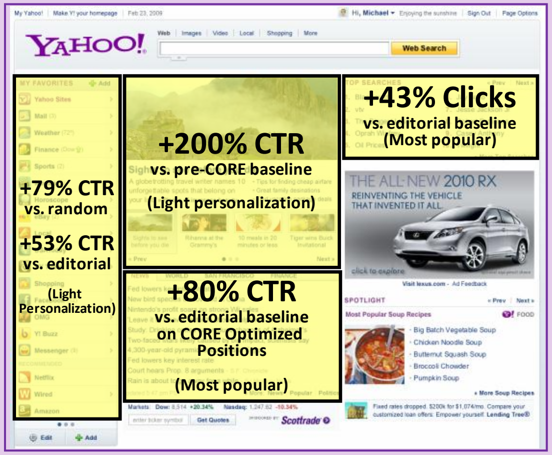

+200% CTR vs. pre-CORE baseline (Light personalization)

+43% Clicks vs. editorial baseline (Most popular)

+79% CTR vs. random

+53% CTR vs. editorial

(Light Personalization)

+80% CTR vs. editorial baseline on CORE Optimized Positions

(Most popular)

YAHOO!
Web Search
THE ALL-NEW 2010 RX
REINVENTING THE VEHICLE THAT INVENTED IT ALL
SPOTLIGHT
Most Popular Soup Recipes
Scottrade

Solution

Personalization requires a real-time feedback loop across properties, leveraging user interests, intent, and context to optimize user engagement. Hadoop/HBase is leveraged for modeling (item/ user) and async processing, Hive for analytics

Hadoop Tomorrow

Trends to address

T1	**Data ingestion** ▪ Data is doubling every 18 months (50-60% annual growth) with greater access from multiple touch points ▪ Analysis has become more important than storage and retrieval with increasing information velocity, volume, and variety
T2	**Thriving open-source community** ▪ Strong community contributing to Hadoop from Yahoo!, eBay, Facebook, LinkedIn, Twitter, Hortonworks, Cloudera, etc. ▪ Seen as the preferred solution for big data analytics with several interchangeable components e.g. Cassandra/ HBase
T3	**New use-cases** ▪ From increasing sales and user engagement to managing risk and fraud, detecting stock market patterns, predicting mortgage default rates to civil infrastructure and telco churn management ▪ From back-office offline analytics to customer-facing 24x7 production systems
T4	**Traditional vendor's dilemma – Extend or Fork** ▪ Consolidating advanced SQL and NoSQL space (EMC, IBM, HP, and Terradata making their moves with Greenplum, Netezza, Vertica, and Aster Data) ▪ Hadoop-based distribution in the hopes of pushing appliance and service sales (Greenplum HD Data Computing Appliance, IBM InfoSphere BigInsights, Informatica 9.1 Big Data Integration platform) ▪ No winning solution or approach to processing big data yet (transactional or non-transactional for processing and management)
T5	**Integrated & real-time processing** ▪ Need for an integrated customer solution that can handle transactions, processing, search, and analytics across the entire data set ▪ Increased interest in knowing what is happening right now vs. offline processing implies solution support for simultaneous read and write with real-time response

Source: McKinsey Quarterly, Gartner Research, Yahoo! Research

Paradigm shifts | Cost/Behavioral angle

Technological

Compute

- The ability to effectively process multi-petabytes of data across thousands of inexpensive computing resources

Economic

Data

- The cost of data acquisition has practically gone to zero
- The cost of data storage is approaching zero

Social & Mobile

Latency

- Consumer’s increasing comfort in pushing user-generated content for sharing
- Consumer’s increasing demand for pulling both public and private information, anywhere, anytime, on every form-factor

Paradigm shifts | Technology angle

	Scenario	**Characteristic**
Compute	**Offload**	Off critical resources
	Load-balance	Better resource utilization
	Batch	Faster time to completion
	Speculate	Real-time, immersive experience
Data	**Acquire**	Public or proprietary
	Cure	Transformative steps prior to compute
	Aggregate	Compose structured & unstructured
	Version	Storage management
Latency	**Cache**	Access anywhere, anytime, on any form-factor
	Stage	Move large-scale data across storage tiers
	Distribute	Locality of reference
	Transact	Stateful, resuming where you've left off

New value-props

- Run both compute- (math & statistical) as well as data-intensive workloads

HPC vs. Hadoop | Architectural comparison

HPC

Computational
(Supercomputing)
C++/Java Programming
MPI
(Numeric Solutions to Partial Differential Equations)
OpenMP
(Image Processing, Rendering, Transcoding)
SOA
(Black-Scholes, Monte Carlos, Correlation Matrices)

On-premise Enterprise Data
(in Multi-Terabyte range)
Structured, Floating Points

Closely-coupled Architecture
(Scale to 1000 nodes)
Fine-grain Scheduling
Multiple Network Topologies
(Compute Node secure isolation)
Low Latency
(InfiniBand)
Mission Critical High Availability
Enterprise-grade
(Security, Versioning, Servicing)

- *Run MR workloads*
- *Build a pathway to data in cloud*
- *Scale to a higher node count*
- *Add GP-GPU support*

Hadoop

Data Intensive
(NoSQL querying)
Java/Pig Scripting
Map-Reduce
(Web Analytics, Ads Targeting, Spam Filtering, Web Traffic Validation)

Web Data in Cloud
(in Multi-Petabyte range)
Unstructured, Semi-numerical

Loosely-coupled Architecture
(Scale to 4000 nodes)
Coarse-grain Scheduling
Single Network Topology
(RPC wire protocol)
Higher Latency
No High Availability

- *Add fine-grain scheduling*
- *Run non-MR workloads*
- *Be enterprise-grade*
- *Add GP-GPU support?*

HPC vs. Hadoop | HPC workload characterization

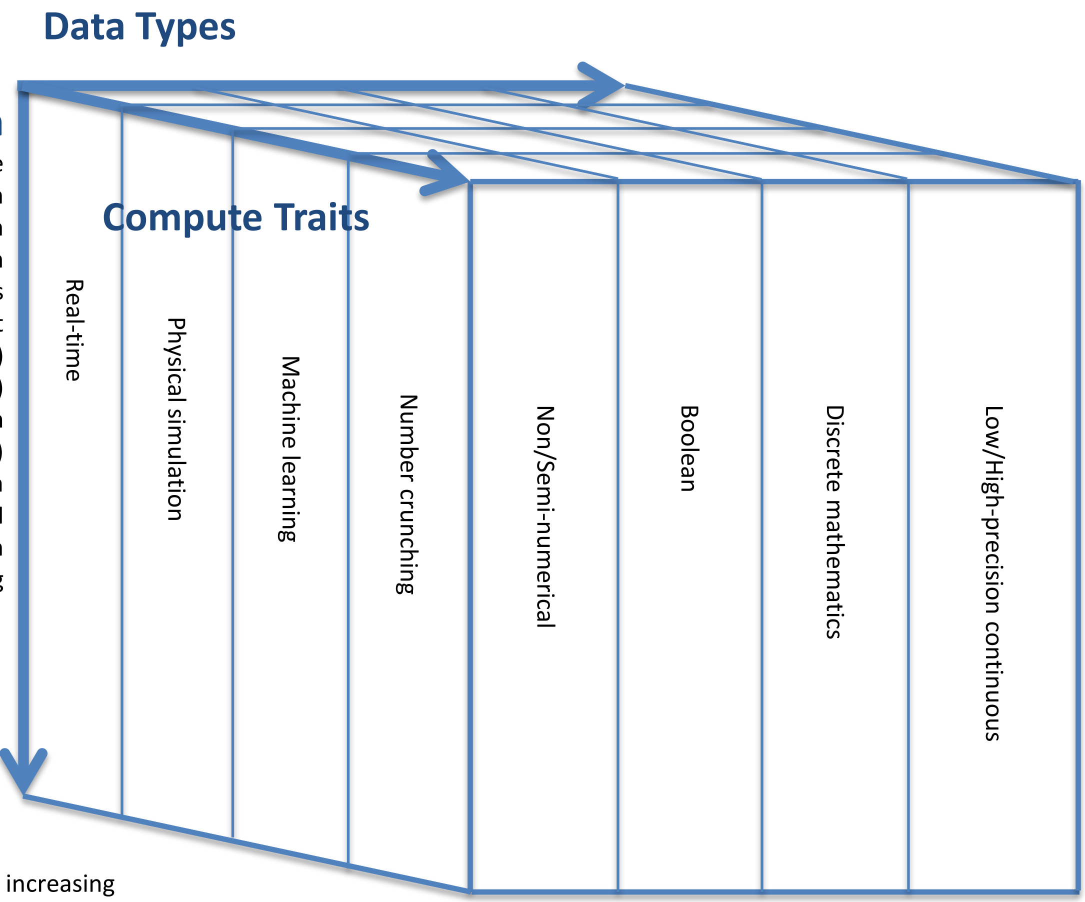

Data Parallelism

Traditional, transacted database
Economic simulation
Circuit simulation
Game-tree exploration
Particle movement under fields
Lagrangian fluid dynamics (fluid element decomposition)
Finite-element stress-strain analysis (crash test)
N-body problem
Euleran fluid dynamics (spatial decomposition)
Dense matrix multiplication & factorization
SETI
Integer factorization
Gene matching

Arrows indicate a direction toward generally increasing suitability or scalability to cloud migration

Putting it all together | High-level data workflow

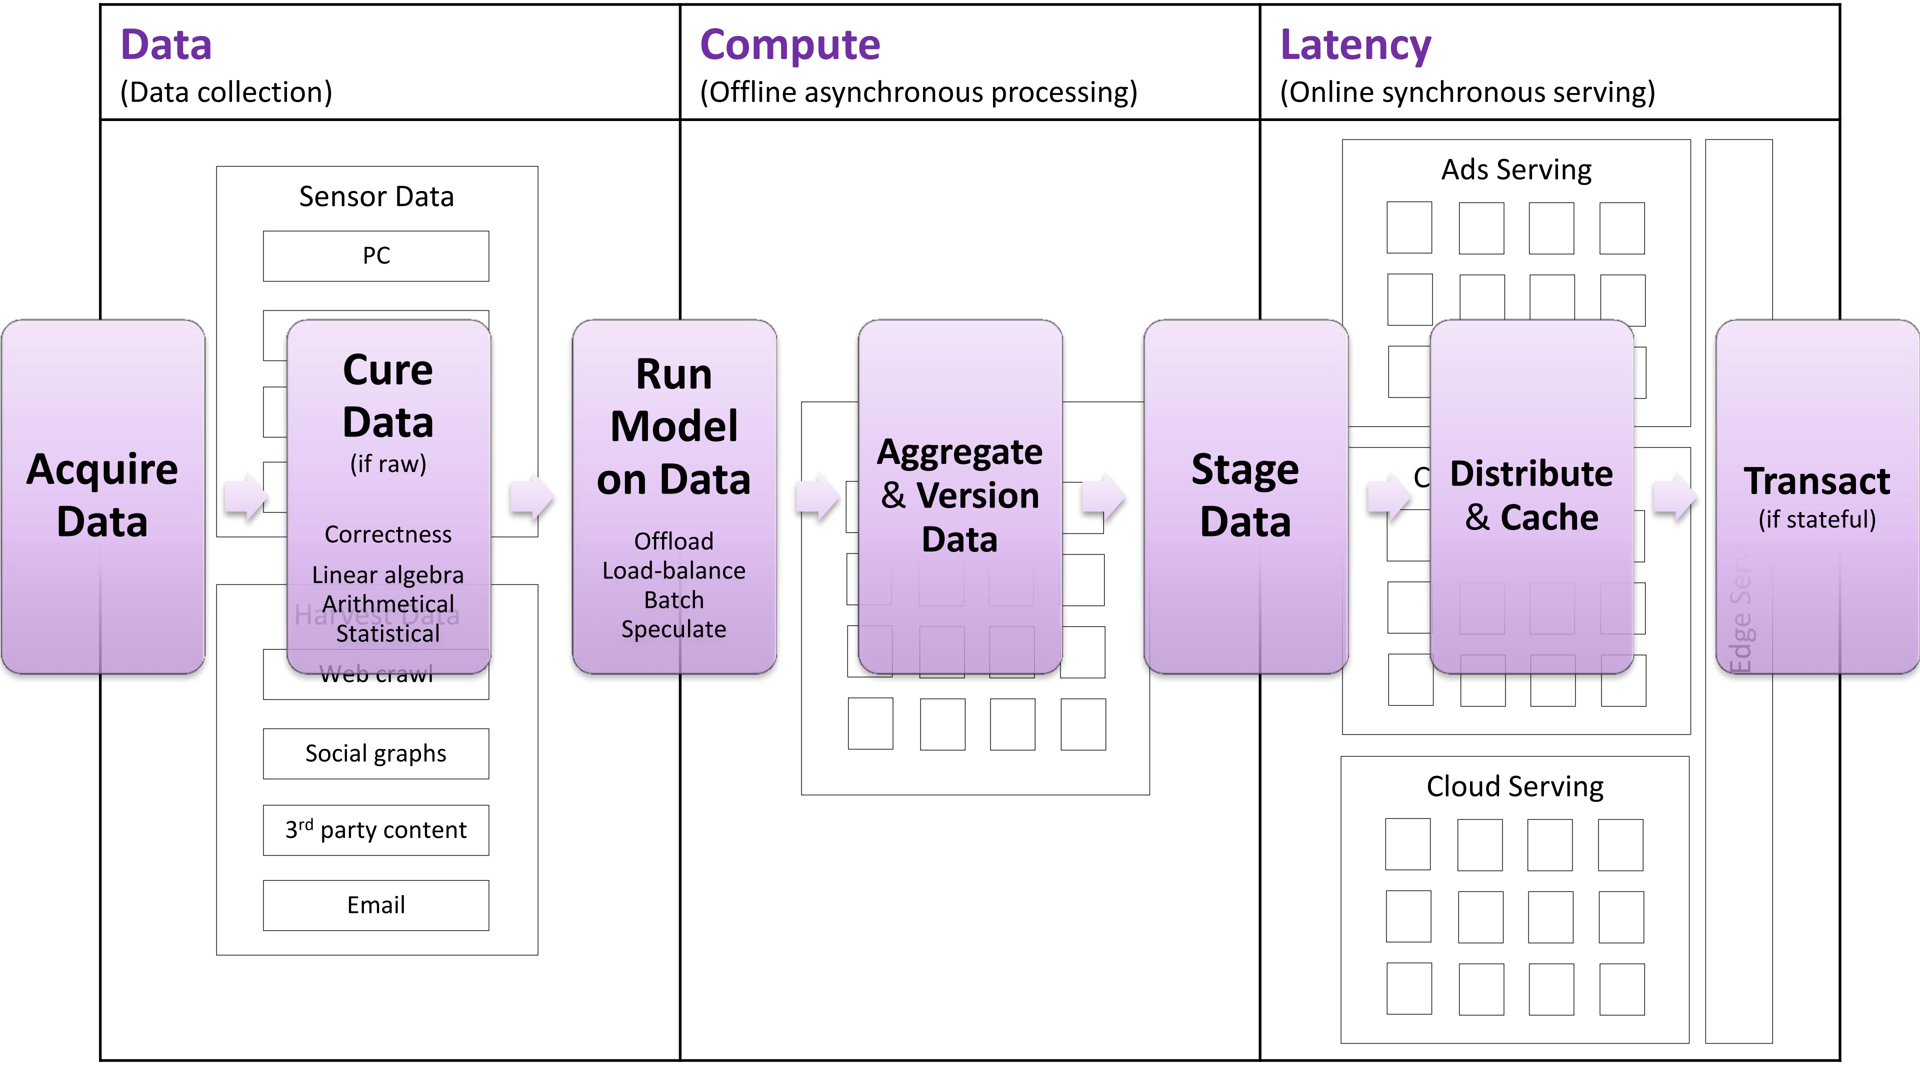

Data
(Data collection)

Compute
(Offline asynchronous processing)

Latency
(Online synchronous serving)