


The Intel Science & Technology Center for Cloud Computing

Center Vision and Research Overview

August 3, 2011

Intel Science and Technology Centers (ISTCs)



- ISTCs funded for 3+2 years and span multiple institutions
- Encourage collaboration among the best researchers in the field
- Four Intel funded researchers per center work on-campus
- Encourage collaboration between Intel and academia
- Public domain IP and open source software increase impact

ISTCs Announced To Date



Visual Computing
Stanford University
January 2011



Secure Computing
UC Berkeley
June 2011



Cloud Computing
Carnegie Mellon University
August 2011



Embedded Computing
Carnegie Mellon University
August 2011

...with more on the horizon

The Intel Science and Technology Center for Cloud Computing

Intel PI



Phil Gibbons
Intel Labs
Principal Researcher

Academic PI



Greg Ganger, CMU
Stephen J. Jatras Professor of
Electrical and Computer Engineering



Clouds on the Horizon: by 2015...

More Users



>1 Billion More
Netizen's¹

More Devices



>15 Billion
Connected Devices²

More Data




>1 Zetabyte
Internet Traffic³

More Energy?

Today's Technology Would Require
Building 45 New Coal Power Plants to
Support 2015 IT Infrastructure⁴

More Costs?


BAIN & COMPANY
IT will spend ~\$2T on
deployment & operations thru
2015⁵ unless smarter infrastructure radically
simplifies management of virtualized
environments.

Internet and device expansion drives need for cloud innovation

1. IDC "The Internet Reaches Late Adolescence" Dec 2009, extrapolation by Intel for 2015

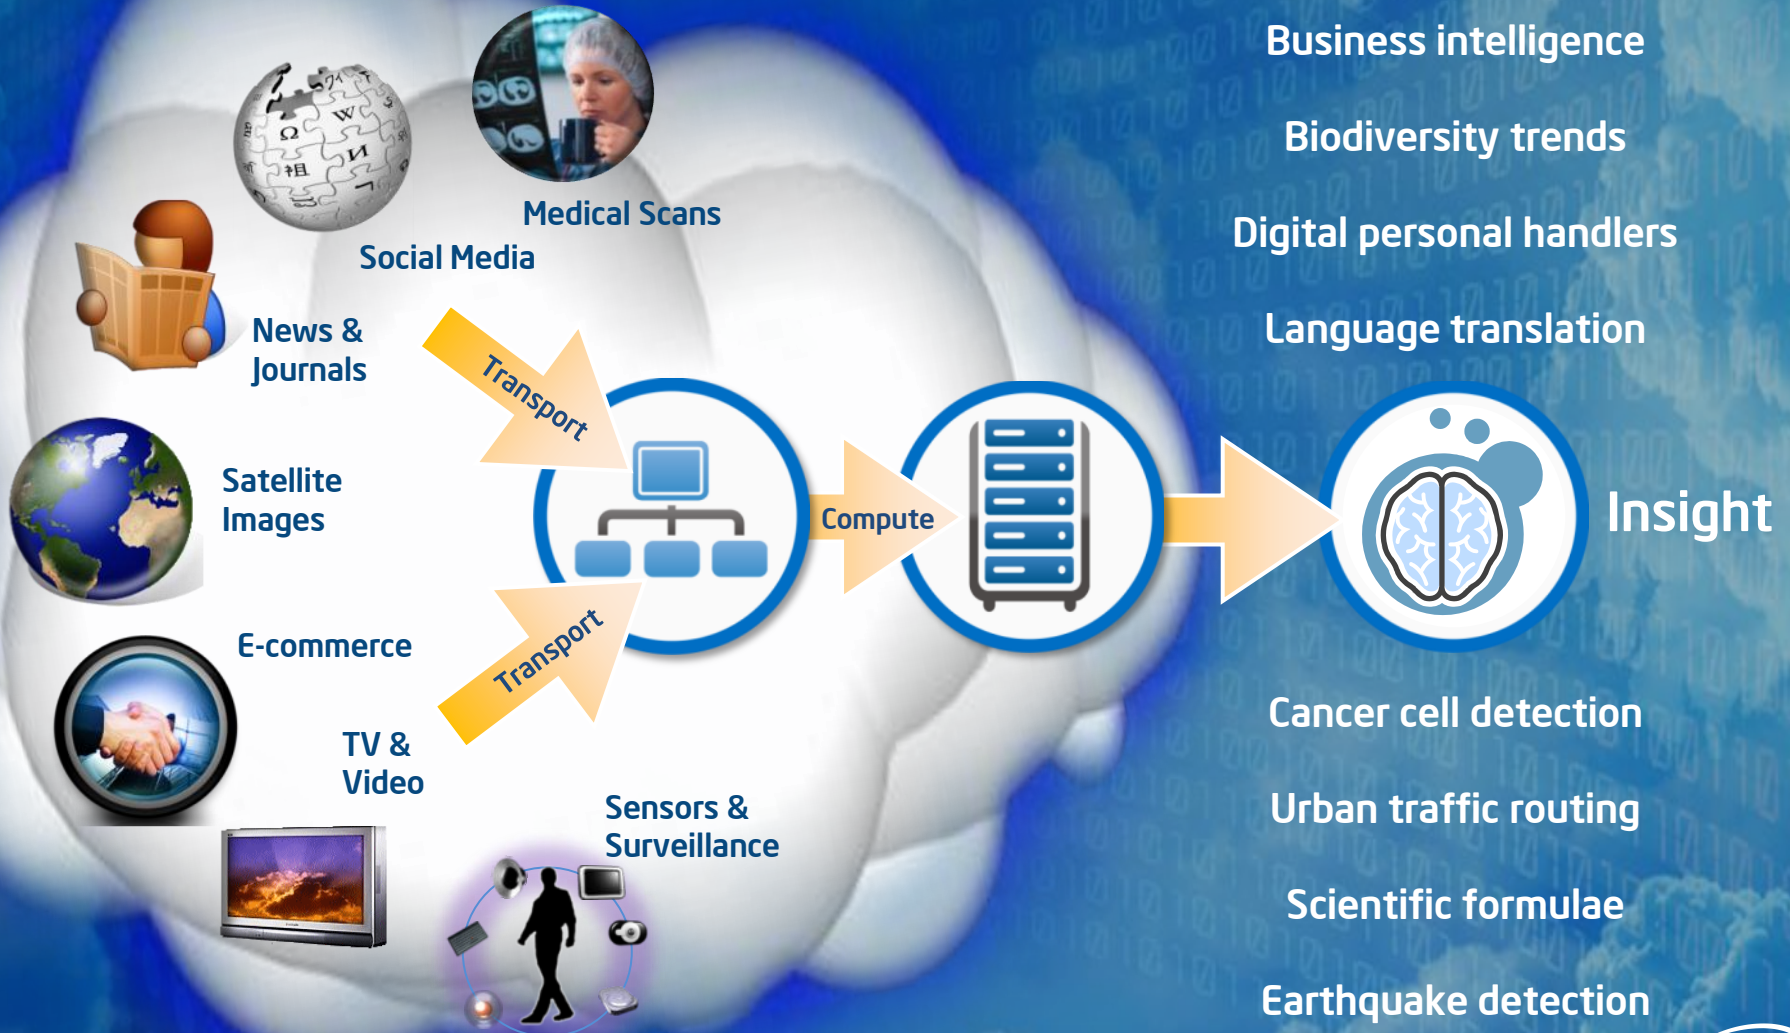
2. ECG "Worldwide Device Estimates Year 2020 - Intel One Smart Network Work" forecast

3. http://www.cisco.com/assets/sdc_content_elements/networking_solutions/service_provider/visual_networking_ip_traffic_chart.html extrapolated to 2015


4. Intel projection. See ISTC-CC whitepaper at intel.com/go/istc

5. Intel commission analysis with Bain Consulting. See ISTC-CC whitepaper at intel.com/go/istc

Opportunity: Channel Info into Insight



Intel Cloud 2015 Vision



To achieve vision and beyond, make the cloud more:

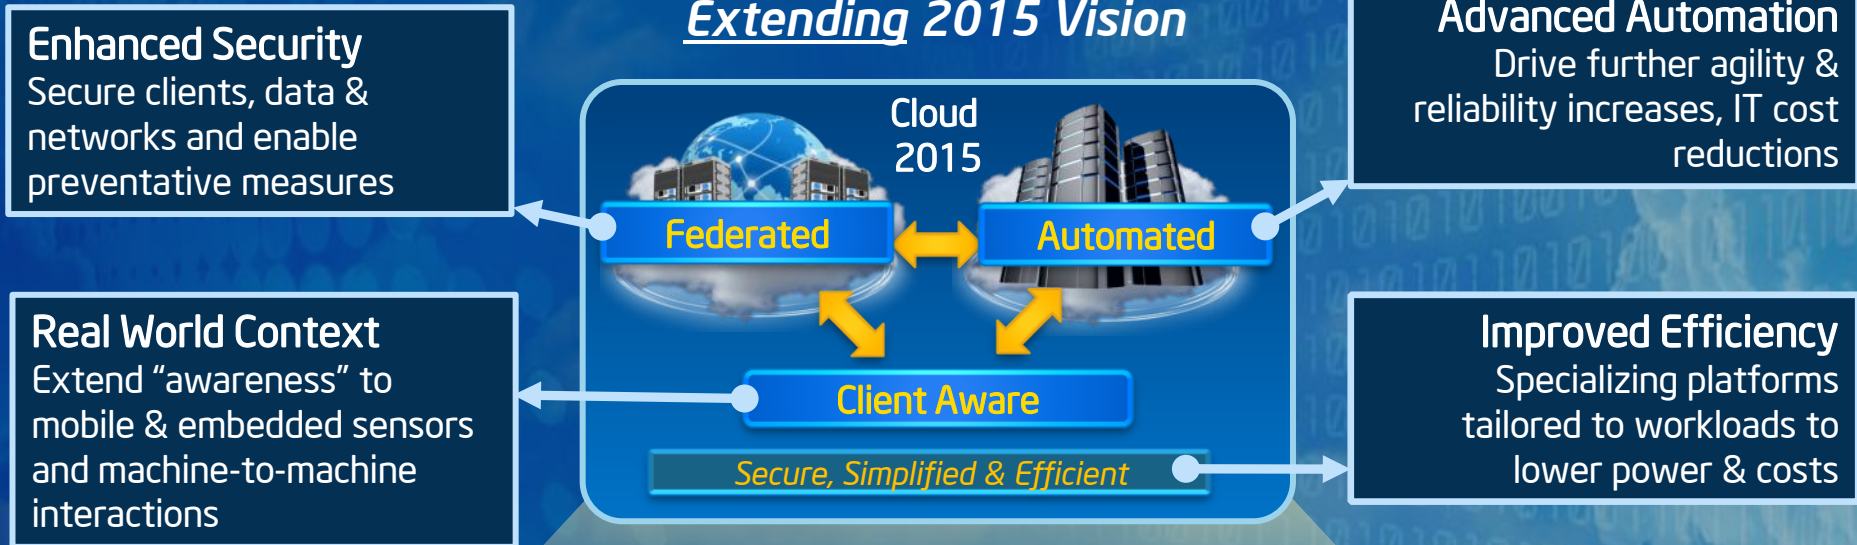
Secure
Reduce risk, increase compliance, manage hybrid clouds

Efficient
Optimizing technologies to decrease energy, human and physical asset consumption

Simplified
Simplify data center operations to reduce cables, complexity and cost

Beyond Cloud 2015: R&D for 2020+

Extending 2015 Vision



Expanding to broader demands for the future cloud

"Big Data" Insights

Derive business, scientific, & social insight from global knowledge

Distributed

Enabling compute to move wherever it is needed and apps that span cloud, client and edge.

Choosing Carnegie Mellon

EXPERTISE

- **Prof. Greg Ganger:** Director of Parallel Data Laboratory, chosen to testify before congressional hearing on the benefits and risks of moving federal IT to the cloud.
- **Prof. Garth Gibson:** IEEE R.B. Johnson Award recipient for seminal work on RAID, leader in PRObE effort
- **Prof. David Andersen:** Leader in use of lightweight nodes for improving data center energy efficiency

EXPERIENCE

- **Data Center Observatory:** >450 nodes for research
- **Open Cirrus:** network of 15 institutions including Carnegie Mellon, Georgia Tech, and Intel

ENVIRONMENT

- **Notable collaborative research efforts:** integration of optical networking in the data center, power-proportional cluster file systems, and data center energy-efficiency improvement
- **10 year history of successful Intel-CMU collaboration**



Data Center Observatory

ISTC-CC Research Agenda

RESEARCH PILLAR

TARGET BENEFIT

Specialization

- Dramatic efficiency to lower power & increase productivity

Automation

- Agility, efficiency & robustness increases, IT cost reductions

Big Data

- Business, scientific, and social insight from huge bodies of data

To the Edge

- Richer interactive experiences via smarter devices & edge-locality

Security


- Secure clients, data & networks and preventative measures

Newest ISTC for Cloud Computing


Recent ISTC for Secure Computing

Specialization

Lightweight
nodes



Many-
core



Phase-change
Memory (PCM)

- Specialization is fundamental to efficiency
 - No single platform best for all application types
 - Called **division of labor** in sociology (see also, bees)
- Cloud computing must embrace specialization
 - As well as consequent heterogeneity and change-over-time
 - Stark contrast to common data center practice
- Sample research activities
 - Specializing via lightweight nodes (e.g., for key-value stores)
 - Specializing via many-core (e.g., for visualization)
 - Exploring impact of new technologies (e.g., PCM)
 - Programming models for adapting to specializations

Automation

- Automation is crucial to cloud reaching potential
 - Needed yesterday, but cloud makes it worse
 - Larger scale, more varied app mixes, specialization,...
- Sample research activities:
 - New techniques for diagnosing problems
 - Including new instrumentation, algorithms, and approaches
 - More robust software upgrade management
 - Via runtime correctness checking at scale
 - Resource allocation and scheduling, in face of...
 - Multi-faceted goals: availability, efficiency, predictability, ...
 - Diverse mixes of workloads (type, duration, priority)
 - Mixes of specialized computing platforms



Big Data



**Customer
Database**

~600 TB

**HD
Internet
Video**



12 EB/yr



**Particle
Physics**

300 EB/yr

- Extracting insights from large datasets
 - A.k.a. “analytics” or “data-intensive computing” (DISC)
 - Becoming critical in nearly every domain
 - and likely to dominate cloud data centers of the future
- Need new programming/execution models
 - For productivity, efficiency, and agility
- Sample research activities
 - New frameworks for advanced machine learning
 - Tools for debugging DISC programs (inc. performance)
 - Usable high-ingest cloud data management

10^6
MEGA

10^9
GIGA


10^{12}
TERA

10^{15}
PETA

10^{18}
EXA

To the Edge

- Edge devices will participate in cloud activities
 - Serving as bridge to physical world (sense/actuate)
 - Enhancing interactivity despite location / connectivity
- Need new programming/execution models
 - For adaptive cloud + edge cooperation
- Sample research activities
 - Adaptive cloud-assisted mobile computation
 - Mitigating reliance on limited uplink bandwidth
 - Decentralized, edge-local cloud architectures



Collaboration, outreach & shared testbeds

- ISTC-CC will be open and collaborative
 - Extensive collaborations among initial participants
 - engaging others in vibrant, growing community
- Connecting with others on broader agenda
 - Ex: future cloud apps (e.g., visual and embedded ISTCs)
 - Ex: cloud security (e.g., ISTC-SC and CyLab)
- Shared testbeds play a crucial role
 - Engages cloud application researchers
 - Enables capture of usage patterns, case studies, etc.
 - Example: OpenCirrus (x3)

Summary

- ISTC investments extend and expand Intel's Cloud vision w/ insight from academic thought leaders
- Innovating in automation & specialization to improve efficiency, reduce power, lower costs
- Breaking down client/cloud/edge barriers to make the cloud more pervasive, responsive, and useful
- Improving 'Big Data' analysis to spur business, scientific, and social breakthroughs