

Epilogue

Monday, May 03, 2010
7:15 AM

Epilogue

What we know.

What we don't know.

What does the future hold?

We know how to...

Monday, May 03, 2010
7:25 AM

We know how to use the cloud to:

- hide implementation details.
- promote code portability.
- provide elastic performance and scalability.
- build robust services.
- implement business policy.
- support business process.
- reassign risks.
- harness parallel computing (Pig, Bloom).

Cloud challenges

Monday, May 03, 2010
7:30 AM

Cloud challenges

risks and guarantees: "public cloud privacy" is an oxymoron.

legacy business systems: cost of adoption is high.

vendor lock-in: one's code commits one to a specific cloud architecture.

green computing is a moving target: the more one widens one's view, the more the optimal answer changes.

Other aspects

Monday, May 03, 2010
7:32 AM

Other aspects we didn't cover:

Private clouds:

Eucalyptus/Xen

VMWare

Other cloud metaphors:

Microsoft Azure: clouds as a gateway to supercomputing.

The X10 programming language: efficient multi-core programming.

Lessons learned

Monday, May 03, 2010
7:44 AM

Lessons learned

One can gain ability by **giving up control** of parts of a program and treating them as opaque services.

One must choreograph use of these services with **partial knowledge** rather than complete understanding.

One can gain flexibility by expressing programs as **compositions of services**.

A theory of clouds governs one's choices.

Giving up control

Monday, May 03, 2010
7:36 AM

Giving up control

There is power in **giving up control** of parts of a program.

One has to program with incomplete information, but

One is not responsible for understanding that part

Examples

Distributed data stores (AppEngine).

Distributed processing (Pig).

Policy-based access control (SalesForce).

But behavior must be **observable**:

Strong or weak consistency.

Weak or strong concurrency.

Theory of clouds

Monday, May 02, 2011
10:44 AM

The (nascent) theory of clouds

The CAP theorem: one cannot have all of **consistency (C)**, **high availability (A)**, and **partitionability (P)** against failures in a single cloud.

The CALM principle: consistency arises from **logical monotonicity**.

These arise from one another

The world needs class AP,
so we have to do without C,
so we **use logical monotonicity to do without C.**

Power awareness

Monday, May 03, 2010
7:49 AM

Power awareness

Power awareness brings into question the most basic assumptions of service-oriented architectures.

SLAs are **not** a definitive expression of performance requirements.

The way we specify requirements has a **dire effect** on energy efficiency.

The dream of clouds

Monday, May 02, 2011
11:01 AM

The initial dream of clouds

Make computing and storage **someone else's problem.**

Treat computational cycles and storage as a **public utility** like water or electricity.

Be able to **pay for what you need** (rather than running your own infrastructure).

Exploit **economy of scale** to provide better service at lower cost.

This course has been about...

Monday, May 02, 2011

11:18 AM

This course has been about...

That dream, specifically.

What we can do about it.

What we can't do about it.

Barley's law

Monday, May 02, 2011
11:04 AM

Steven Barley: *When you make a technological change, you never ever get what you expect. You always get something else.*

The dream has changed

Monday, May 02, 2011
11:07 AM

The dream has changed

Certain parts of that initial dream are impossible
Still infeasible to store private data in public
clouds.

Most large corporations are "building their own".
"The internet isn't fast enough" for an app-
centered world (Google).

Certain parts of a new dream were previously
inconceivable

Ubiquitous (public) information gathering and
correlation.

The slow demise of personal privacy.

A new definition of personal accountability and
citizenship.

"The socialization of knowledge."

In other words

Monday, May 02, 2011
11:20 AM

In this course,

The cloud has proven unsuitable for storing private information,
and has, at the same time, shown potential for making many things that are not currently public easily available to the public.

Heading toward a singularity

Monday, May 02, 2011

11:21 AM

Charles Stross, Ray Kurzweil, Verner Vinge: computing is moving us toward a "technological singularity", after which the nature of the world cannot be predicted.

In their view, this singularity arises from **artificial intelligence**.

But in my shorter-term view, there is another singularity:

The **socialization of knowledge**.

The **demise of personal privacy**.

The **rise of personal security**.

that is much nearer than that.

From data to knowledge

Monday, May 03, 2010

8:03 AM

From data to knowledge

Data: just bits

Information: has some concept of meaning.

Knowledge: has some concept of causality, universality, and repeatability.

Due to John Strassner.

A quiet social revolution

Monday, May 03, 2010
12:20 PM

A quiet social revolution

Google allows their scientists to post to **blogs** instead of publishing papers, and let **reference counts** substitute for citations!

Reason: regular scientific publishing impedes progress!

Google tools advocate a **socialist** approach to knowledge:

There is no undo in Google Wave; if you call someone a jerk, and commit the change, it's recorded forever.

The average young person is much less concerned about **personal privacy** than someone my age...
... but much more concerned about **security!**

Privacy versus security

Monday, May 02, 2011

11:33 AM

Privacy **versus** security

One often hears "privacy" and "security" in the same sentence, but

Actually, the terms are somewhat in opposition in a social view of knowledge.

A clear distinction

Privacy is whether people **know what you do** or not.

Security is whether people **can change the record** of what you do, or not.

Clouds provide near-perfect security, at the cost of (perhaps seriously) inferior privacy.

If nothing is private, then
security violations can be traced.
and security can be perfected.

If anything is private, then
security problems of that aspect cannot be traced.

My prediction

Monday, May 02, 2011
11:39 AM

In the long term

Security is going to win.

Privacy is going to lose.

The concept of citizenship itself will change.

Economy of scale will win.

The short-term future of cloud computing

Monday, May 02, 2011
12:03 PM

My short-term predictions of the future of cloud computing

Resource management and SLA's: **solved.**

Green computing: **solved.**

Security: **solved.**

Privacy: **a non-issue.**

Frontiers of cloud computing

Monday, May 02, 2011

12:02 PM

So, what are the real frontiers of cloud computing?

Locating **knowledge**

The problem of **ontology**

The problem of **provenance**

Identity and location

Monday, May 03, 2010
12:08 PM

Identity and location

What the cloud really does is to **decouple data identity from location.**

Then how do we "locate" data?

One might say one searches via keywords, except that....

Ontology

Monday, May 03, 2010
8:00 AM

Ontology

For the first time, we have the ability to maintain **global** scientific databases

But we now face the **problem of ontology**.

Simply stated: data that mean the same thing can have **different names**.

The problem of ontology

Monday, May 03, 2010
12:11 PM

The problem of ontology

Data has a taxonomy: a description of what fields mean in context.

But two different data tables can have an ontology: how meanings correspond.

Example: compare

Name	ID	Phone
FullName	Number	Cell

An ontology relates data from two tables:

Is Name==FullName?

Is ID==Number?

Is Phone==Cell?

Not really.

Provenance

Monday, May 02, 2011
12:09 PM

Provenance

The provenance of a conclusion is the answer to the question: "why do you think that?"

Provenance-aware filesystems: keep track of what was done to create a file.

The real challenges...

Monday, May 03, 2010
10:56 AM

We need new concepts of...

Privacy.

Security.

Source/provenance.

Reification/reputation.

to truly use clouds to their full potential.

These concepts are as **social** as they are **technical**.

A vision of the future

Monday, May 03, 2010
12:15 PM

My vision of the future

Anyone can study global warming on a global scale by tying into the global database of climatological data.

Anyone can use supercomputing to simulate an experiment or test a hypothesis -- without even knowing it!

High-school students can use planetary-scale computing!

Anyone can mine social data for references of interest on a global scale -- even continuously.

This is the ultimate **socialism of knowledge** (and you note that I did not say "socialization"!)

But, actually...

Monday, May 03, 2010
12:19 PM

But, actually...

The capabilities are already there.

What is lacking is the **social** component...

...which has lagged years (or perhaps centuries)
behind technology...