

Clouds on the Horizon: Evolving Teaching, Learning and Technology in Higher Education

Dr. Charles Severance
University of Michigan School of Information
October 6, 2009




Outline

- Looking at Google for inspiration
- Google App Engine
- Cloud Style High Performance Computing
- Cloud Intelligence
- Cloud Identity
- Cloud Learning
- A ToDo List of Tasks for IT and CS

Google

Cape Town and George...

- Having an IMAP server in George is a pain
- Gmail works no matter where he travels
- It is the *same* undersea cables, etc..
- Google is simply better than the Internet/Web

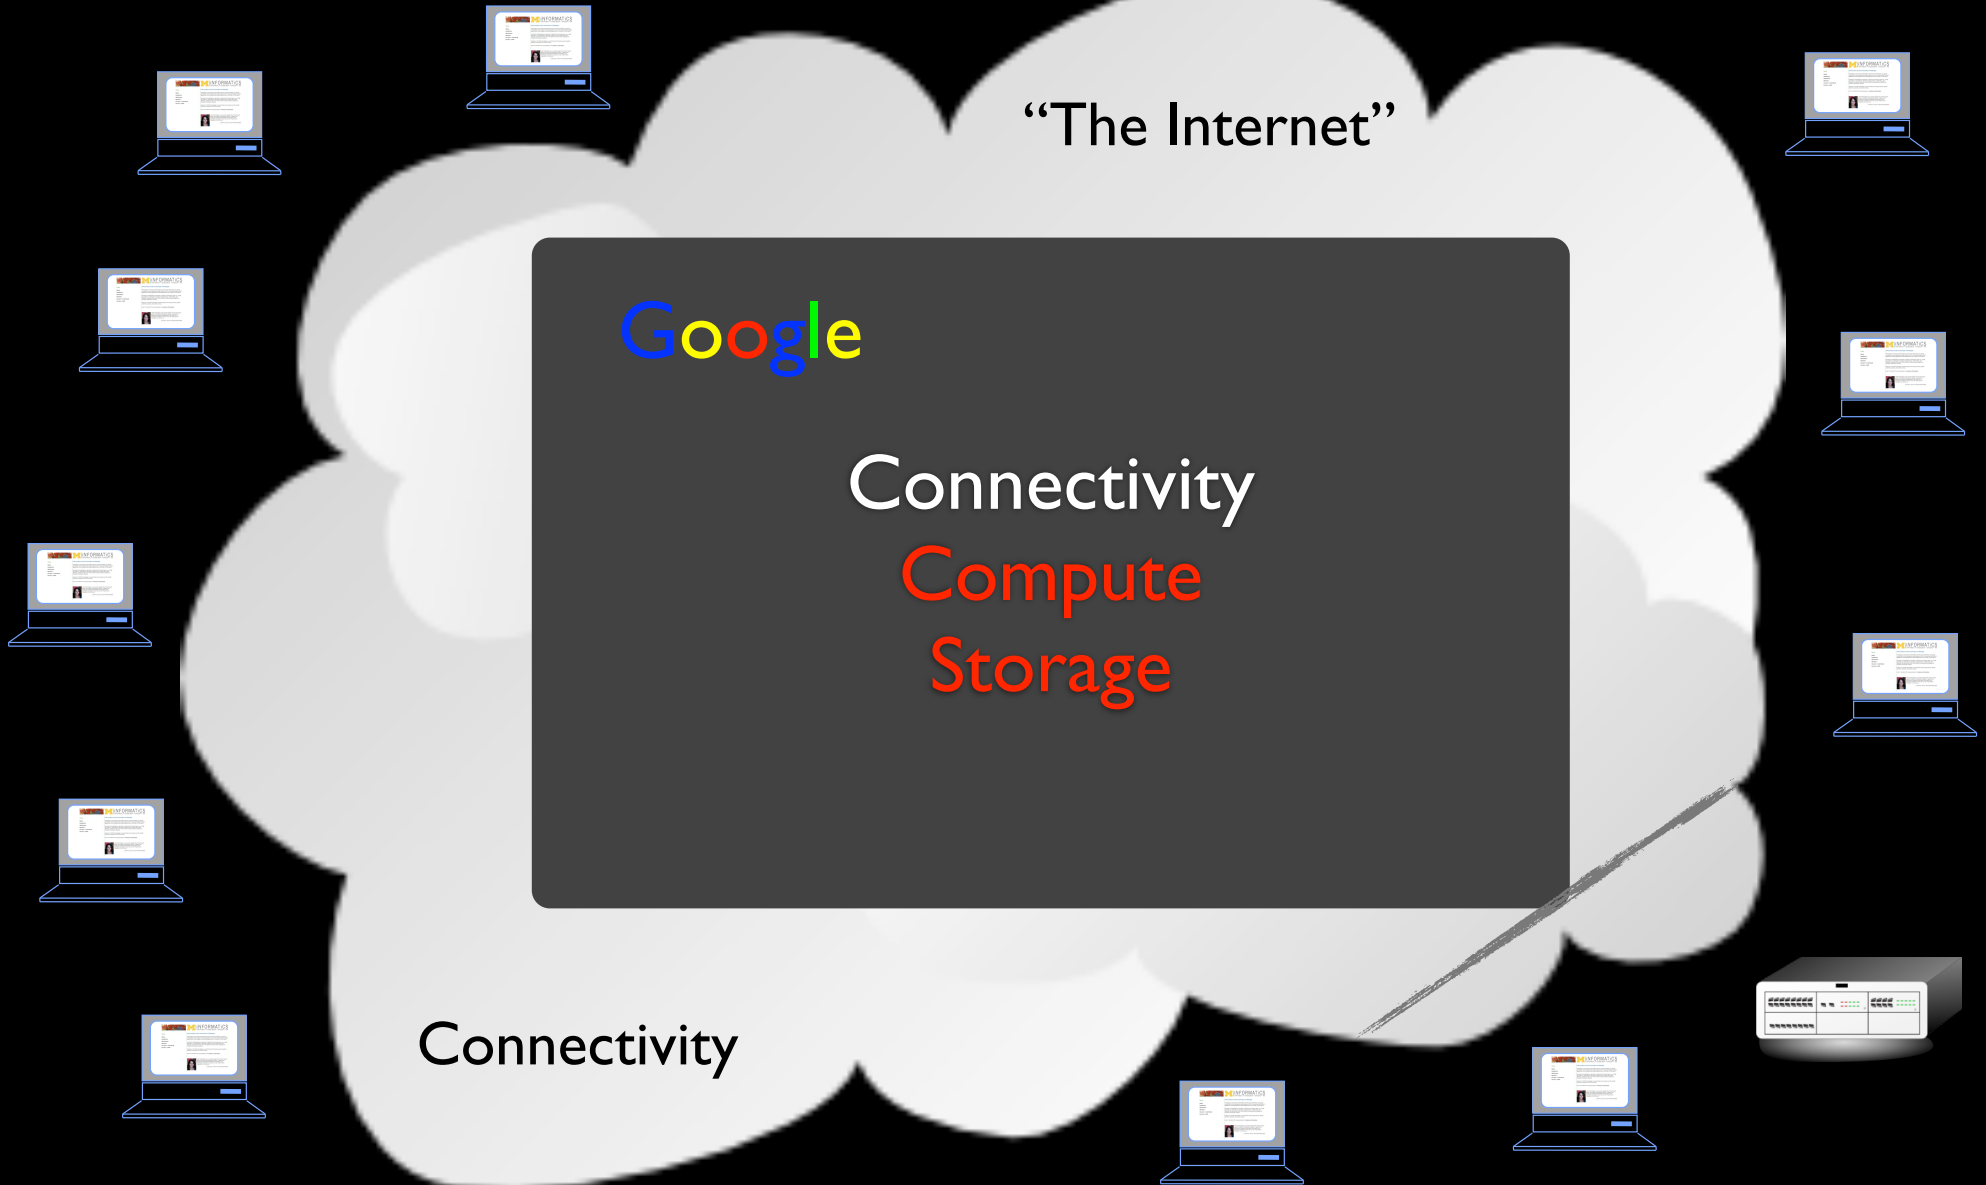


“The Internet”


Google

Connectivity
Compute
Storage

Connectivity



Google I/O 2008



Google I/O 2008 Keynote

- Google I/O '08 Keynote by Marissa Mayer
- Usability / User Experience / User Testing / Architecture / Philosophy



<http://www.youtube.com/watch?v=6x0cAzQ7PVs>




<http://www.youtube.com/watch?v=6x0cAzQ7PVs>

Lessons

- The cloud is wide - we can touch 1000 servers in 0.1 seconds
- For things that seem “intelligent” 0.2 seconds is fast enough - as long as you can do a lot of them
- Lots of spread-out storage and a fast scan is important
- Data - Information - Knowledge - starts with data and the ability to look through that data quickly

Scalable Infrastructure



<http://www.youtube.com/watch?v=zRwPSFpLX8I>

Infrastructure

- The only sustainable scalability is when you scale with inexpensive, green solutions
- Tape Backup is a rate limiting factor - so we need something creative
- Disaster recovery - “Of course!”

Google App Engine



- Expose Google's worldwide Infrastructure to us as developers
- Make the web better
- Be the first widely used "cloud" environment - beat Amazon, Microsoft, and Yahoo!

Campfire One: Introducing Google App Engine (pt. 1)



<http://www.youtube.com/watch?v=3Ztr-HhWXIc>
<http://www.youtube.com/watch?v=oTFL7FPLnXY>

Google App Engine



- When you write a Google Application Engine Application - you are running in the Google Cloud
- Just like you were a Google Developer
- You don't know where you are running or if one copy of a thousand copies of you are running
- Google hosts small applications for *free* - larger applications pay by usage




Hmmm...

- Clouds are an emerging new technical “form”
- Google is the first “prototype”
- Dive in...
- Teach a class
- Write a book
- Change the world?



Cloud-Style HPC



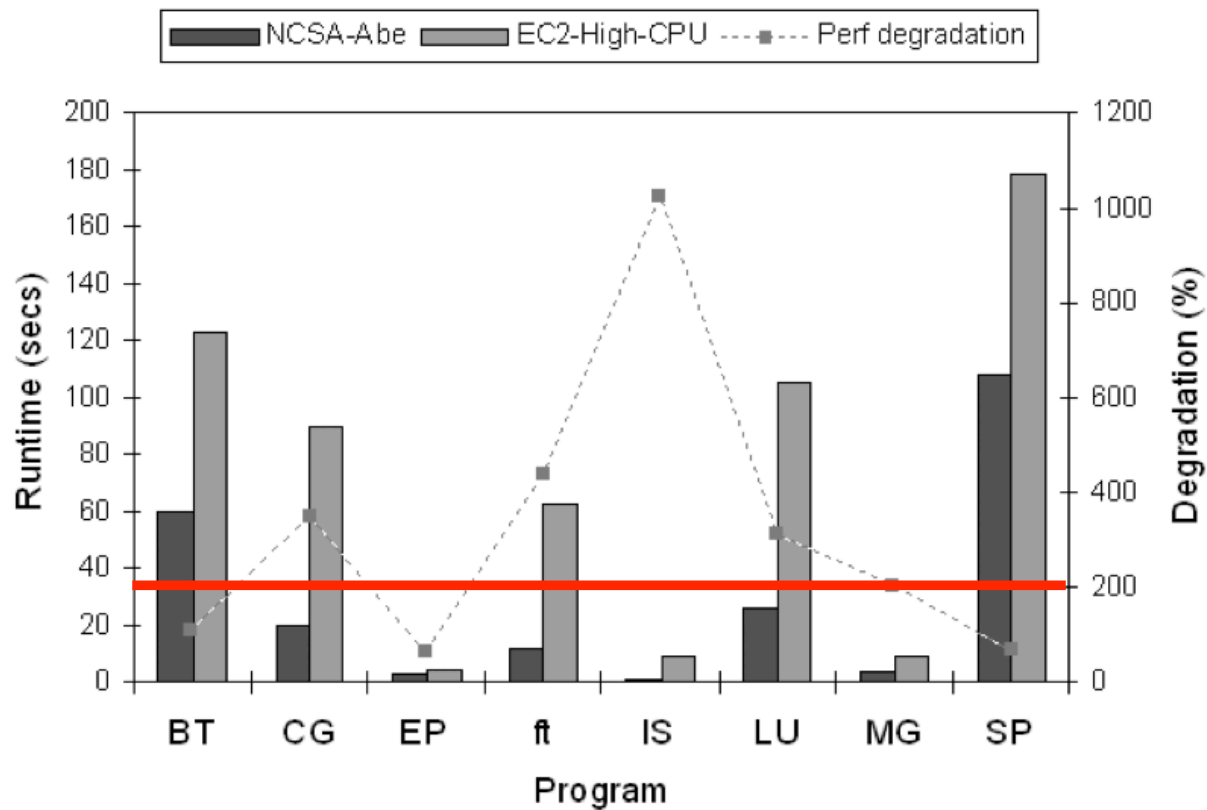
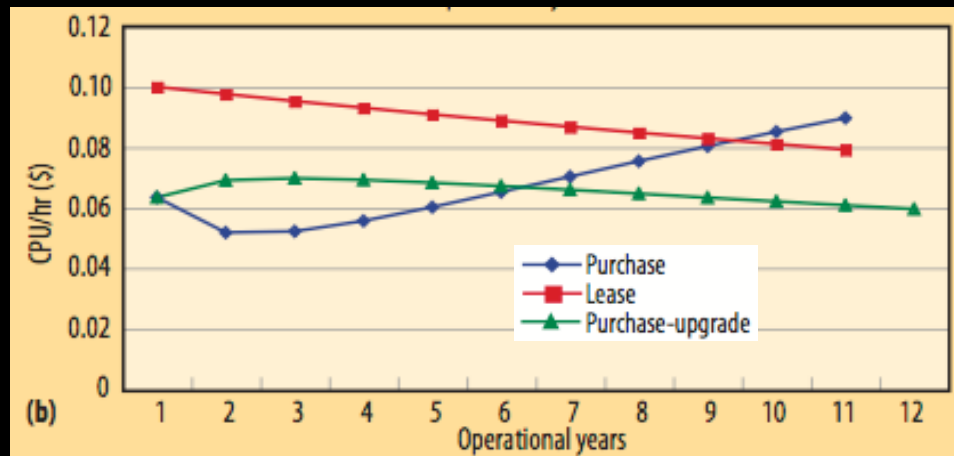
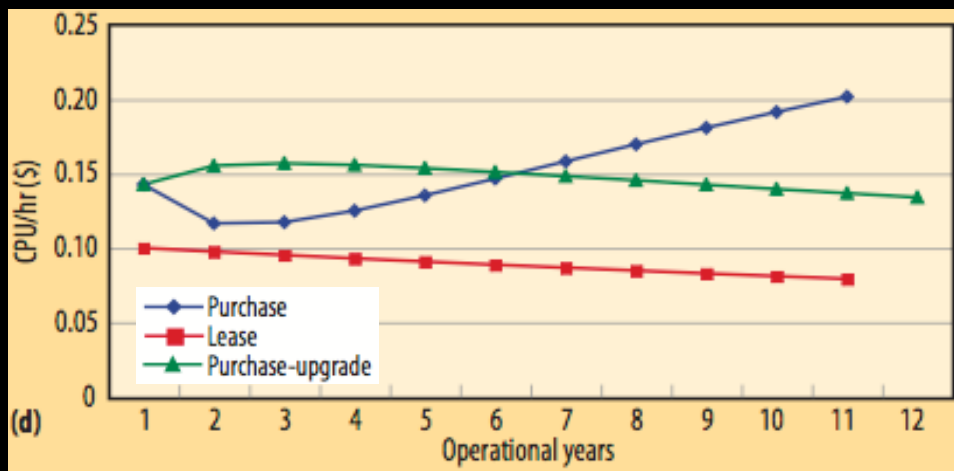


FIGURE 2. NPB-MPI (CLASS B) RUNTIMES ON 32 CPUS ON THE NCSA AND EC2 CLUSTER. BT AND SP WERE RUN WITH 16 CPUS ONLY. OVERLAID IS THE PERCENTAGE DEGRADATION IN THE EC2 RUNS.



90% utilization and low utility cost.

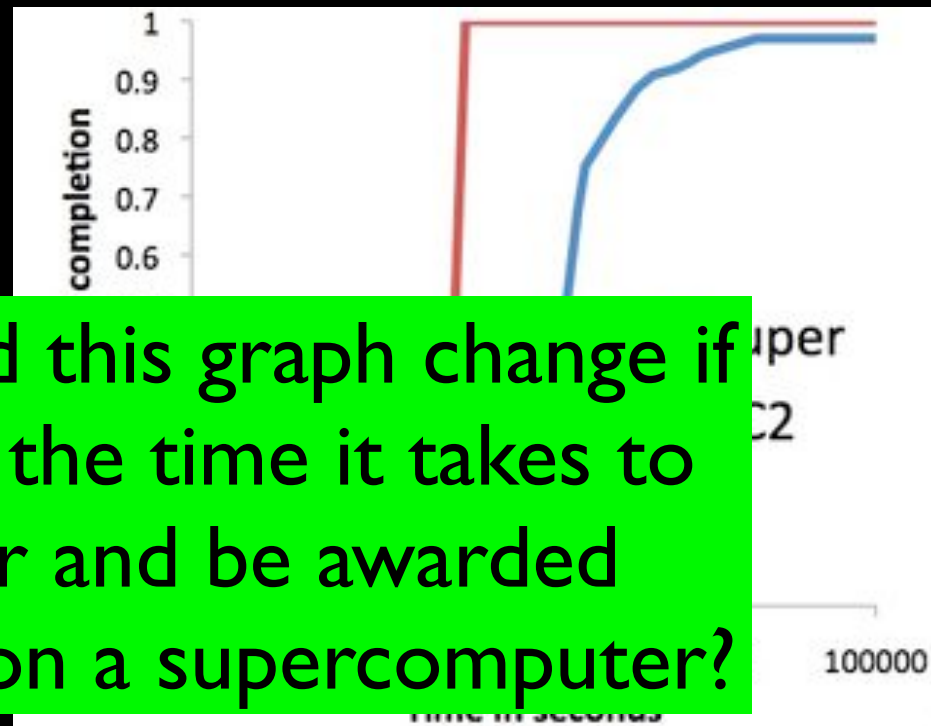


40% utilization and low utility cost.

The Real Cost of a CPU Hour, IEEE Computer Magazine, April 2009.



What's Faster? Amazon EC2 or a Supercomputer?



How would this graph change if we added the time it takes to apply for and be awarded resources on a supercomputer?

<http://ianfoster.typepad.com/blog/2009/08/whats-faster-a-supercomputer-or-ec2.html>

Note: Ian says that this is not precise research - it is to make us think.

The HPC Cycle...


- On campus **scientists complain** that it takes too long to get resources on national resources, or the queues are too long
- We buy “strategic” on campus compute resources...
- We get as many scientists to try to use it as we can
- Success means that local queues are really long and **scientists complain** that we need to get more processors

The one constant in the cycle...


We do have a responsibility to help non-CS scientists make the best use of available HPC resources. Part of that **training** includes smaller local compute resources to be used to tune and adjust software....



www.cnx.org




Is the Singularity at Hand?



http://en.wikipedia.org/wiki/Technological_singularity

It is Getting Slightly Closer...



CAPABILITY



<http://asimo.honda.com/>

iRobot Roomba® 400



25	1	114	117	150	152	120	46	19	126
191	121	104	116	160	105	89	125	40	14
31	139	113	94	97	199	154	140	195	122
112	163	177	48	78	101	130	83	35	197
44	54	106	143	59	38	3	41	93	81
20	164	4	11	131	0	107	71	159	69
181	178	173	148	62	142	170	72	37	145
60	187	169	99	15	85	26	8	192	17
129	73	45	9	24	188	42	151	51	183
179	79	50	76	34	33	185	102	193	184

What is the Largest Number?

110 10 20 10 30 33 182 105 103 180
 150 13 02 0 50 188 05 121 21 183
 00 181 100 00 12 85 50 8 105 11
 181 118 113 108 05 105 110 15 31 102
 10 100 11 131 0 101 11 100
 00 20 100 103 20 38 3 01 03 81
 115 103 111 08 118 100 130 83 32 101
 31 130 113 00 01 100 120 100 102 155
 108 151 100 110 100 102 80 152 00 10
 52 1 110 111 120 125 150 00 10 150

What is the Second Largest
 Number?



What are the last words of
“Where the Wild Things Are”?



what are the last words of where the wild things are

Search

[Advanced Search](#)

Web [+ Show options...](#) Results 1 - 10 of about 31,600,000 for what are the last words of where the wild things are. (0.10 seconds)

[Where the Wild Things Are by Maurice Sendak | LibraryThing](#)

He then travels to where the **wild things are**, where they make him their king. **Last words**, (Click to show. Warning: May contain spoilers. ...

www.librarything.com/work/5537 - [Cached](#) - [Similar](#)

[Where the Wild Things Are Holiday Feature Edition by Maurice ...](#)

Sep 18, 2009 ... Max sails over the ocean to the land of "the **wild things**. ... Where the **Wild Things Are** is a classic book that has been a ... **Last words** ...

www.librarything.com/work/6113813 - [Cached](#) - [Similar](#)

[+ Show more results from www.librarything.com](#)

[Where the Wild Things Are by Maurice Sendak—Children's Book Review ...](#)

Where the **Wild Things Are** is so familiar and well-loved a children's book that ... The **words** and images strike a chord with readers for their direct realism ... tension that carries through to the very **last** line: "And it was still hot. ...

hubpages.com/.../Where-the-Wild-Things-Are-Childrens-Book-Review - [Cached](#) - [Similar](#)



And the point is...

- Researchers will increasingly want to write applications that seem intelligent
- Each of those applications need a few terabytes of spinning, online permanent, network accessible, backed up data that I accumulate over time
- The data will be images, sound, video, or text
- I want to be able to scan all my data, asking a question and getting an answer in less than a second

Your Assignment:
Watch the movie

“Eagle Eye”

and decide what fraction
of the technology in the
movie will be
commonplace within 20
years.



What Is Hadoop?

The Apache Hadoop project develops open-source software for reliable, scalable, distributed computing. Hadoop includes these subprojects:


- **Hadoop Common**: The common utilities that support the other Hadoop subprojects.
- **Avro**: A data serialization system that provides dynamic integration with scripting languages.
- **Chukwa**: A data collection system for managing large distributed systems.
- **HBase**: A scalable, distributed database that supports structured data storage for large tables.
- **HDFS**: A distributed file system that provides high throughput access to application data.
- **Hive**: A data warehouse infrastructure that provides data summarization and ad hoc querying.
- **MapReduce**: A software framework for distributed processing of large data sets on compute clusters.
- **Pig**: A high-level data-flow language and execution framework for parallel computation.
- **ZooKeeper**: A high-performance coordination service for distributed applications.



<http://hadoop.apache.org/>

YAHOO! DEVELOPER NETWORK

<http://developer.yahoo.net/blogs/hadoop/>



Cloud Identity

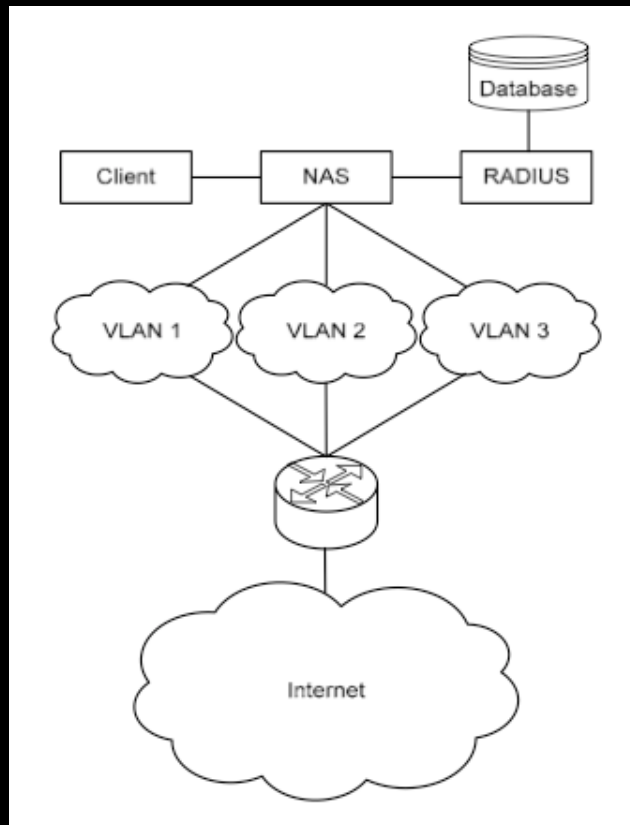
It is actually not that tough...

Unified Identity?

- The commercial cloud space seems to have given up on people having exactly one overarching identity - used everywhere
- People do not **want** one identity but they also do not want to be forced to get one more identity
- Use whatever you have and want to use - Google, Yahoo!, OpenId, Twitter, FaceBook, Blogger...
- Identity is more than just at the moment of login.. OAuth.net

How to you handle
WiFi guests on your
campus?

Eduroam Architecture



- **RADIUS** - Remote Authentication Dial-in User Service
- Created in 1991 by the Merit Network!
- RFC2866

<http://en.wikipedia.org/wiki/RADIUS>

Cloud Learning

EDU@2020



- Ricard N. Katz and Ronald Yanosky
- Google + Sakai = Google Phoenix
- Microsoft buys Blackboard
- Microsoft creates a lifelong personal portfolio
- Virtual teachers and virtual classmates

"College age young people are better socialized to virtual worlds than to real ones."



<rtsp://educouse.rmod.llnwd.net/a680/o1/edu2020.rm>





TsugiProject Goals




- Written in Python, Free, Open Source
- Hostable on Google for **Free**
- Easy to use learning tool building and hosting environment
- Supports **IMS Learning Tools Interoperability**
- **Teachers writing tools - students writing tools (Thousands)**

www.tsugiproject.org

wiscrowd.appspot.com

- A set of tools which allows multi-player educational in-class games
- Support for the teachers using the book “Wisdom of Crowds”
- Works either stand-alone or inside of an LMS using IMS Learning Tools Interoperability

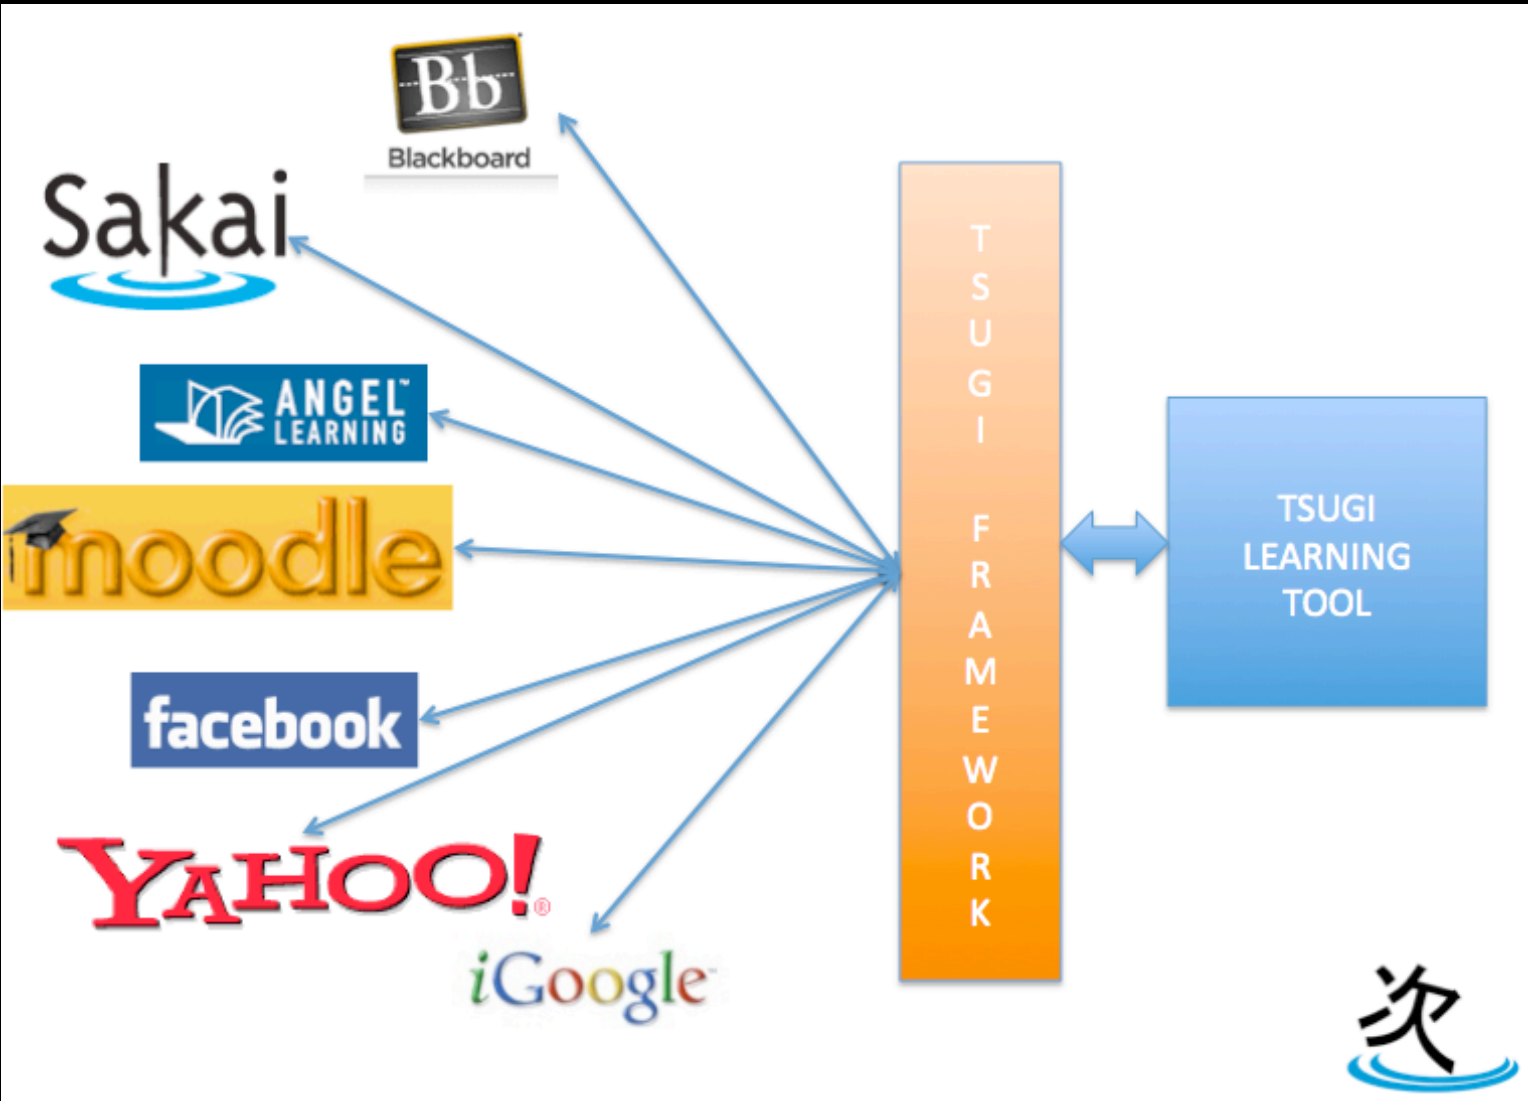
The image is a screenshot of a web application interface. At the top, there is a navigation bar with links for 'My Places', 'Home', 'Help', and 'Logout'. Below this is a header for 'Your Institution UNIVERSITY' with a small image of a building. Underneath the header is a menu with buttons for 'My Institution', 'Courses', 'Community', 'Content Collection', and 'Services'. The main content area shows a greeting: 'Hello: csev from _3_1'. Below the greeting are two input fields: 'Enter Guess: [text input]' and 'Enter Name (opt): [text input]'. At the bottom of the form is a 'Submit Query' button.



High School
College
Freshman
Level



Programming
HTML
CSS
Database
JavaScript



Available Applications



Address Book

Create a dynamic, online address book so your group can stay in touch. You control what fields are part of the address book.

[add](#)



Pearson MyMathLab

Provides practice problems sets and supporting instructional materials

[add](#)



Favorites

Enable your group to build a list of their favorite things. You can choose one of our default list types (favorite posts, movies, music, books, or TV) or create your own custom favorites type.

[add](#)



For Sale

Create an easy way for people to buy and sell from each other. You can choose to customize the application to make it about a specific type of item (e.g. bicycles or tickets or) or leave it open to anything.

[add](#)



Give Away

Whether you're a Freecycling or Freesharing group that is dedicated to giving stuff away or you're a group that is about something else entirely but just wants to make exchanges possible among group members, this application makes it easier for folks to give stuff away.

[add](#)



If we assume tools can be created and hosted anywhere and Plugged into each other... What happens next?



Installing Your Next LMS

- Insert this code in your web page and press “Refresh”

```
<script type="text/javascript"  
  src="http://api.cloudsocial.org/js/ile-main.js">  
</script>  
<script type="text/javascript">  
  ile_init("49c480149a008");  
</script>
```



PythonLearn - Self-paced learning Python


http://www.pythonlearn.com/ RSS Google

PythonLearn [Book](#) [Software](#) [Informatics](#) [Instructor](#) [Python](#) [About](#)

New: Learning Google Application Engine www.appenginelearn.com

For the best effect to learn Python on your own, you should purchase the [textbook](#) and go through the materials in order. If you want to attempt the programming assignments make sure to install the appropriate [software](#) on your system. Installation instructions are provided under the "Software" tab.

This site should not be a substitute for a course you are taking - even if the course you are taking is using the same textbook. Each course and each instructor will take their own approach.



Basic Python

- Writing Simple Programs (and [Assignment Data](#))
- Computing with Strings (C[through a file](#))
- Decision Structures (Chap[Through a File Again](#))
- Computers and Programs [who has the most commits](#)

Wiscrowd grab | close S.

Hello: Joe User from IN101

Enter Guess:

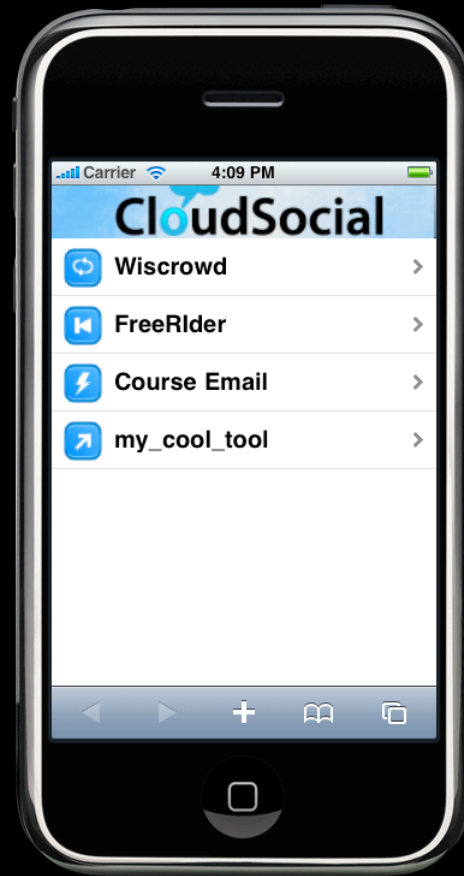
Enter Name (opt):

[Assignment 1 - Getting Started](#), [Assignment 2 - Reading](#), [Assignment 3 - Reading](#), and [Assignment 4 - Figuring out](#)

Powered by [CloudSocial](#)

<http://www.cloudsocial.org/>






Open Educational Resources

- Since CloudSocial allows a user to “bring their LMS along” when visiting content, it is an ideal tool to turn flat OER content into a collaborative learning space – without moving the content.
- Connexions – www.cnx.org
- Open Michigan – open.umich.edu
- MIT Open Courseware



If I worked in Campus IT...



"We don't just think outside the box...
We transform the box into a robot and
then think outside of the robot."

cPanel - Your Own PHP Cloud

- The software that powers ISPs like GoDaddy, RackSpace,
- Give it hardware - it includes and updates Linux, MySQL, Admin, etc
- It deals with run-time performance and load balancing
- Export and import across cPanel-providers
- \$20 per year per server for EDU - give every student an ISP

Outsourcing E-Mail: NO!

- I think everyone should have the choice to read their official campus E-Mail anywhere they like
- People need a low-level e-mail redirect including copy...
- Academic freedom / free speech

Disk Space - More

- Other than a network connection, my 10GB of AFS space at Michigan is my most valued IT benefit
- It is now full :(
- I need about 3GB more per year - for media associated with my teaching output

A Few More Tasks...

- Make EduRoam happen here in the US - I would want to be the first school
- Lets clone Amazon EC2 and S3 protocols, etc and put up our own small clusters that talk these protocols - get involved in figuring out how to help scientists use these servers and protocols
- Start Looking at Hadoop - imagine a multi-institutional Hadoop instance - start with eliminating backup and giving us faculty more, cheap disk space - next move to data-intensive computing


Outline

- Looking at Google for inspiration
- Cloud Applications
- Cloud Computing
- Cloud Intelligence
- Cloud Identity
- Cloud Learning
- A ToDo List of Tasks for IT and CS

www.cloudsocial.org



<< back to portal [CloudCollab](#) [Logout \(drchuck\)](#)

Hello: drchuck (Instructor) from SI101 

Enter Guess:

Enter Name (opt):

```
<script type="text/javascript"
  src="http://api.cloudsocial.org/js/ile-
  main.js">
</script>
<script type="text/javascript">
  ile_init("49c480149a008");
</script>
```

www.tsugiproject.org

www.appenginelearn.com
www.pythonlearn.com



Thank you

www.dr-chuck.com