

HPDC 2008 Cloud Panel

Greg Pfister

Affiliation: Free at Last!



Obligatory cloud icon

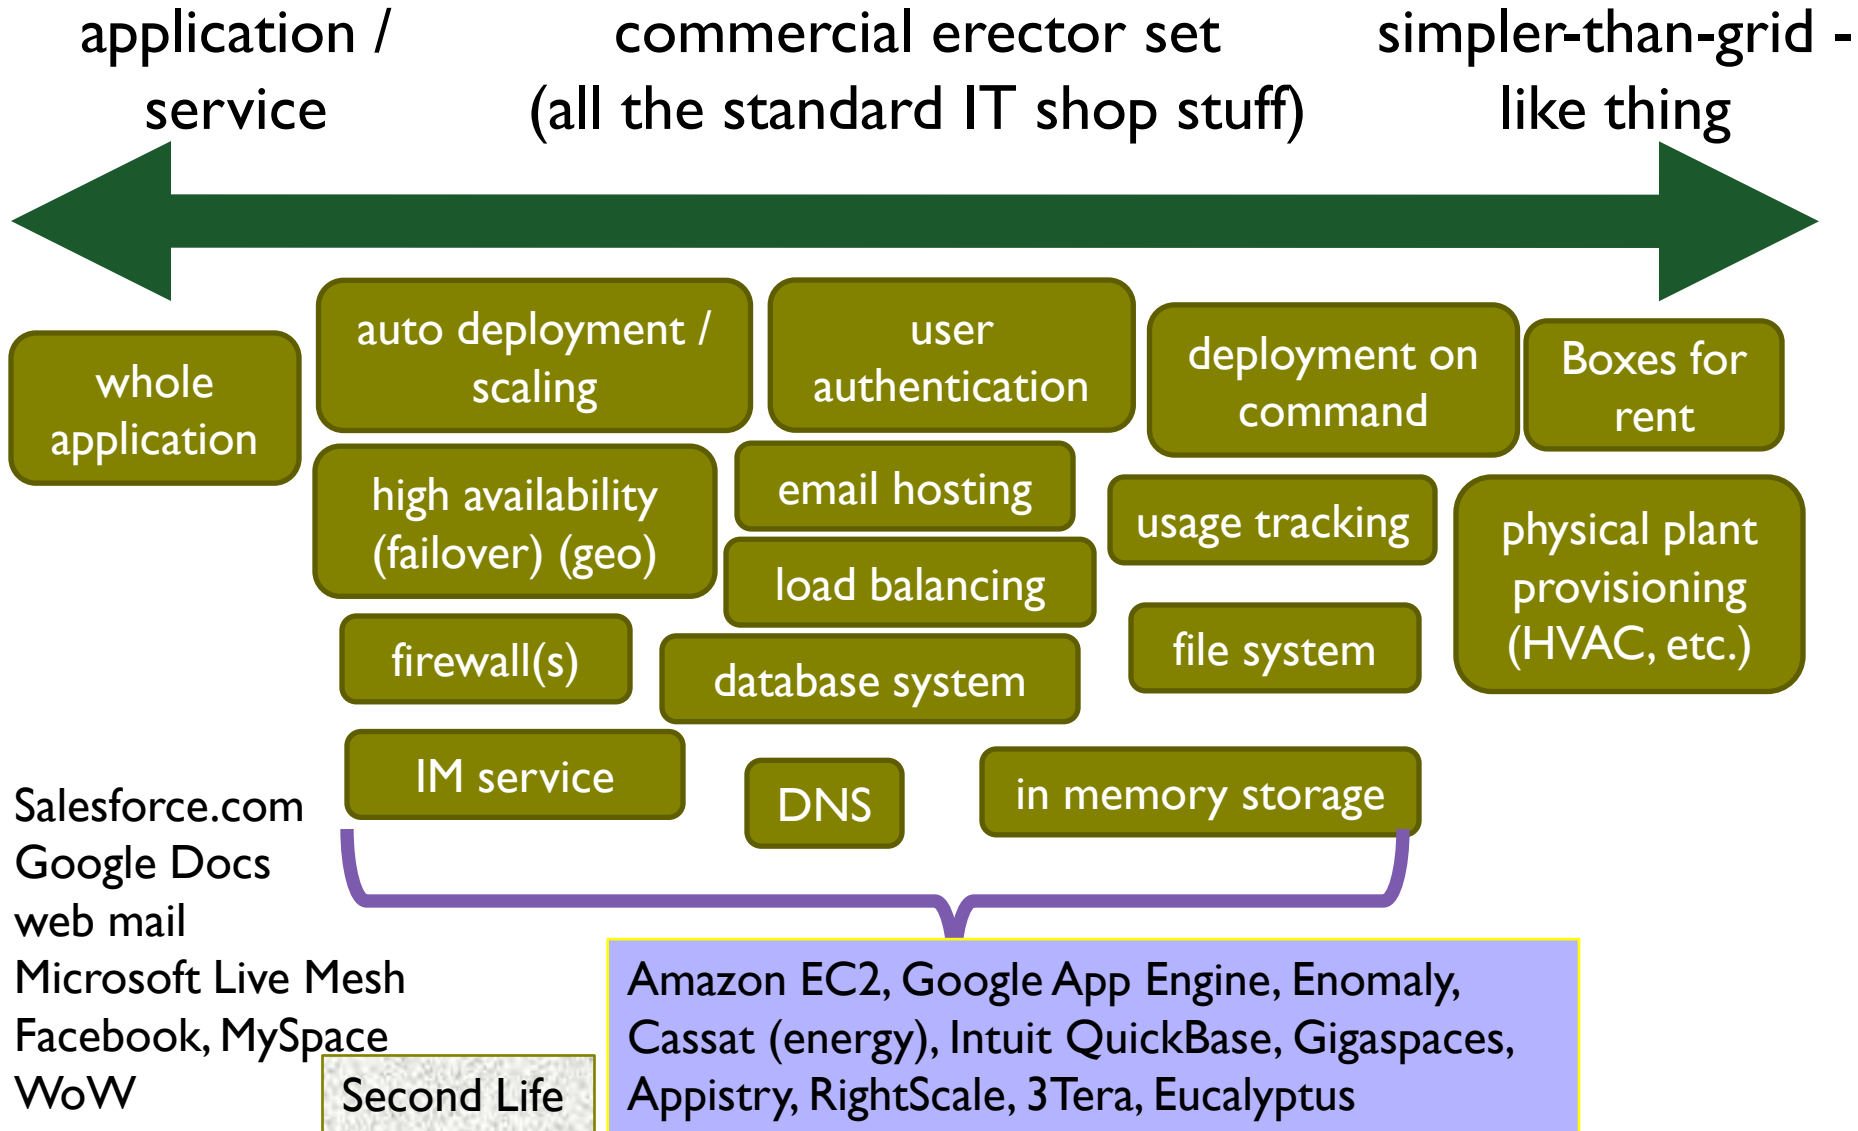
Q1) What is and isn't a cloud?

2

- Close up, a cloud is a fog. (Ray Nugent)
- The term for grid computing that has nothing to do with academia. (Stu Charleton)
- Cloud computing is hiring someone out on the web to host your computing...
- ... where “host” can range across
 - Raw iron
 - Building blocks of varying complexity
 - “standard” commercial infrastructure
 - The whole application



Cloud Hosting Spectrum



Q1b) What isn't a Cloud?

4

- A chipmunk.



- Two things often confused with cloud:
 - Software as a Service (SaaS)
 - A way to provide the infrastructure
 - Virtualization (as in VMWare)
 - A technique to use in deployment
- These can be used to implement cloud hosting, but aren't themselves cloud computing.
- “Pigpen Cloud” is a homonym.
 - Could be implemented using cloud computing.

Q2) What are implications of Clouds for Enterprise Data Centers?

5

- ▶ Cheaper, more responsive, less expensive.
 - ⇒ downsizing
 - ⇒ in some cases, displacement of whole shop
- ▶ A result: Clouds will be resisted.
- ▶ Fortunately (?) many legacy mission-critical applications strongly resist assimilation.



Q3) If TeraGrid evolves to PetaCloud, how would it look?

6

- ▶ As long as grid is tailored to HPC, it's unlikely to become a cloud.
 - ▶ Except, of course, for marketing / funding buzzword purposes.
- ▶ Ways clouds differ from grids:
 - ▶ Significant commercial infrastructure: databases, etc.
 - ▶ Throughput / “transactional” workloads
 - ▶ “automatic” horizontal scaling
 - ▶ Non-optimal client code is OK
 - ▶ Scripting languages, sandboxes,

Yet more random gratuitous clipart



Q4) What are research issues for Clouds?

7

- All issues for commercial computing apply to clouds. E.g.,
 - Improved management
 - Ease / simplicity of use, meaningful performance monitoring, deployment, updating, ...
 - Everything you do with a computer *except* what you bought it for.
 - Scaling of all provided components.
 - Geographical HA / failover.
 - Interoperability and standards (maybe).
 - Specialization to particular workloads
 - Clouds with accelerators and the software to use and manage them?
 - FP, database, crypto, Java, XML, ...



you get the idea

Q5) Can one (or who could) "trust" clouds?

8

There are limits, but we already do.

- Inside the firewall: home-grown clouds, e.g., email
 - (customers invented clouds) (like clusters)
- Outside the firewall: often outsource data (auditing)
 - No intrinsic reason cloud vendors less adept or concerned about privacy/security than clients.
- However:
 - There is certainly a trust hurdle to cross
 - Vendors must earn our trust before we take data out of our mattresses
 - Legal implications for personal privacy may limit use outside the firewall.
 - Geopolitical issues may limit physical cloud locations.
 - Some data will **never** be allowed to reside outside firewalls.



This getting boring now

Q6) Will Cloud Interoperability be Important; if so at what Interface(s) will it be provided?

9

- Interoperability between different cloud implementations – business as usual.
 - SOA, WSS, etc.
- Portability of apps between different cloud implementations; cloud computing standards
 - Would be nice.
 - About as likely & feasible as portability between OSs, databases, ... possible. Takes work.
 - Vendors always like walled gardens and are unlikely to cooperate much.

Q7) What is the Killer App for Clouds?

#1: TWEaaS

10

The Whole Enchilada as a Service

➤ Bill Gates at Microsoft Tech ·Ed 6/08:

- “everything we do at the server level” and “have a service that mirrors that exactly.”
- “hosted in our cloud ... geo-distributed automatically.... many millions of servers.”

➤ Heard from Google CTO several years ago:

- We want to do **all** the programming for **everybody**.



They've got you and me, brother, <clap> in their hands...

Q7) What is the Killer App for Clouds?

#2: Virtual Worlds

11

- Major uptake expected.
 - Training, marketing, meetings, conferences, geographically dispersed teams, recruiting, demonstrations, etc.
- Serious infrastructure requirements
 - Nonstop world simulation, transactional persistent storage, etc.
 - IT shops mostly clueless
- Opening for external vendor (cloud) support
 - ... but then again, commercial cloud vendors are mostly clueless too.

Virtual World Uptake

- 80% Internet user participate by 2011 (Gartner)
- >1B users in 10 years (Strategy Analytics)
- Congressional testimony 4/08:
 - 22-33% CAGR in the number of virtual worlds last year,
 - \$1.02B VC investment 10/07-08.
 - 420,000/week actively *visiting* Second Life
- China CRD: 100 sq km site (former steel mill) converting to house virtual worlds able to support billions of avatars simultaneously.



CRD

Logo of the China Recreation District

Thank you for listening!