CLOUD COMPUTING ARCHITECTURES

Andreas Doerr – Sun Grid Engine

Amir Halfon – CTO, Financial Services

Sun Microsystems, Inc.

Why Cloud Computing Fascinates IT

Developers

- Why won't IT support this?
- Why can't I use the versions I want?
- Why can't I get better availability?
- How can I pay for what I need?
- How quickly can I get more servers?

Deployers

- Why do we have so many versions of everything?
- Where can I cut costs?
- How can I do finer grain provisioning?
- Where do we enforce security, regulation and audit?

Two Perspectives

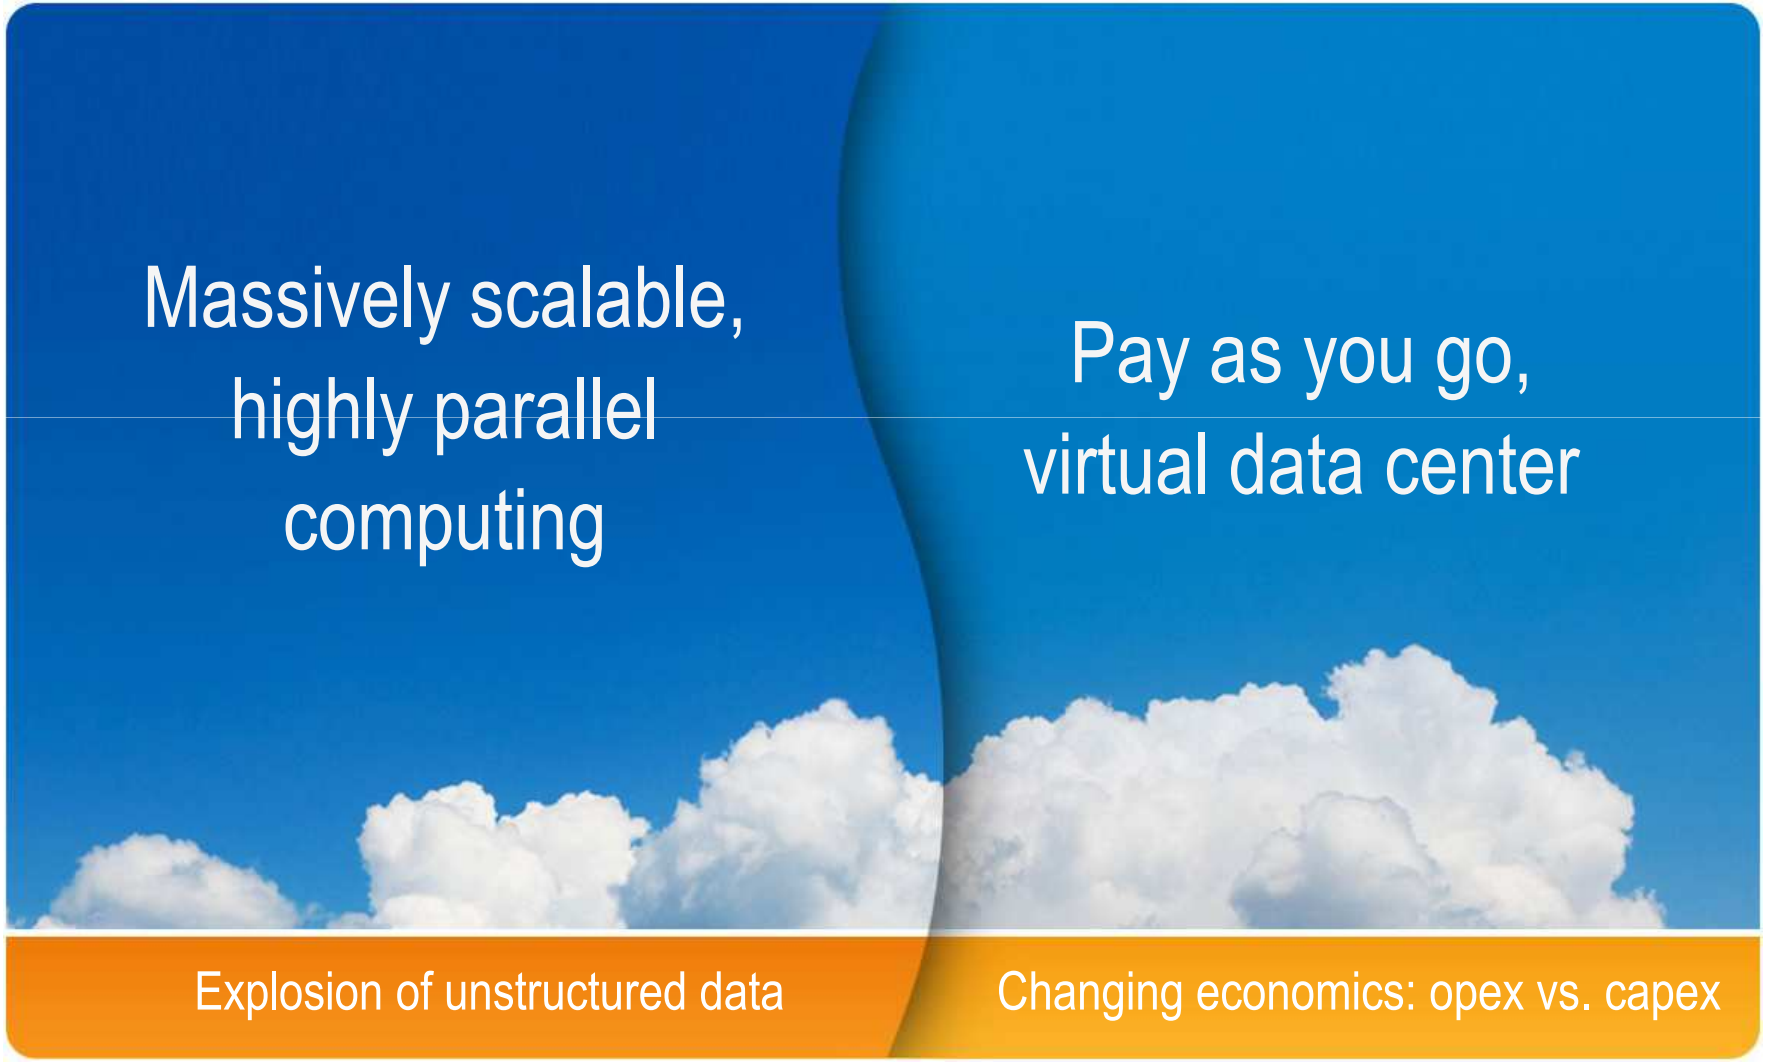

Defining Concepts

From Grids to Clouds

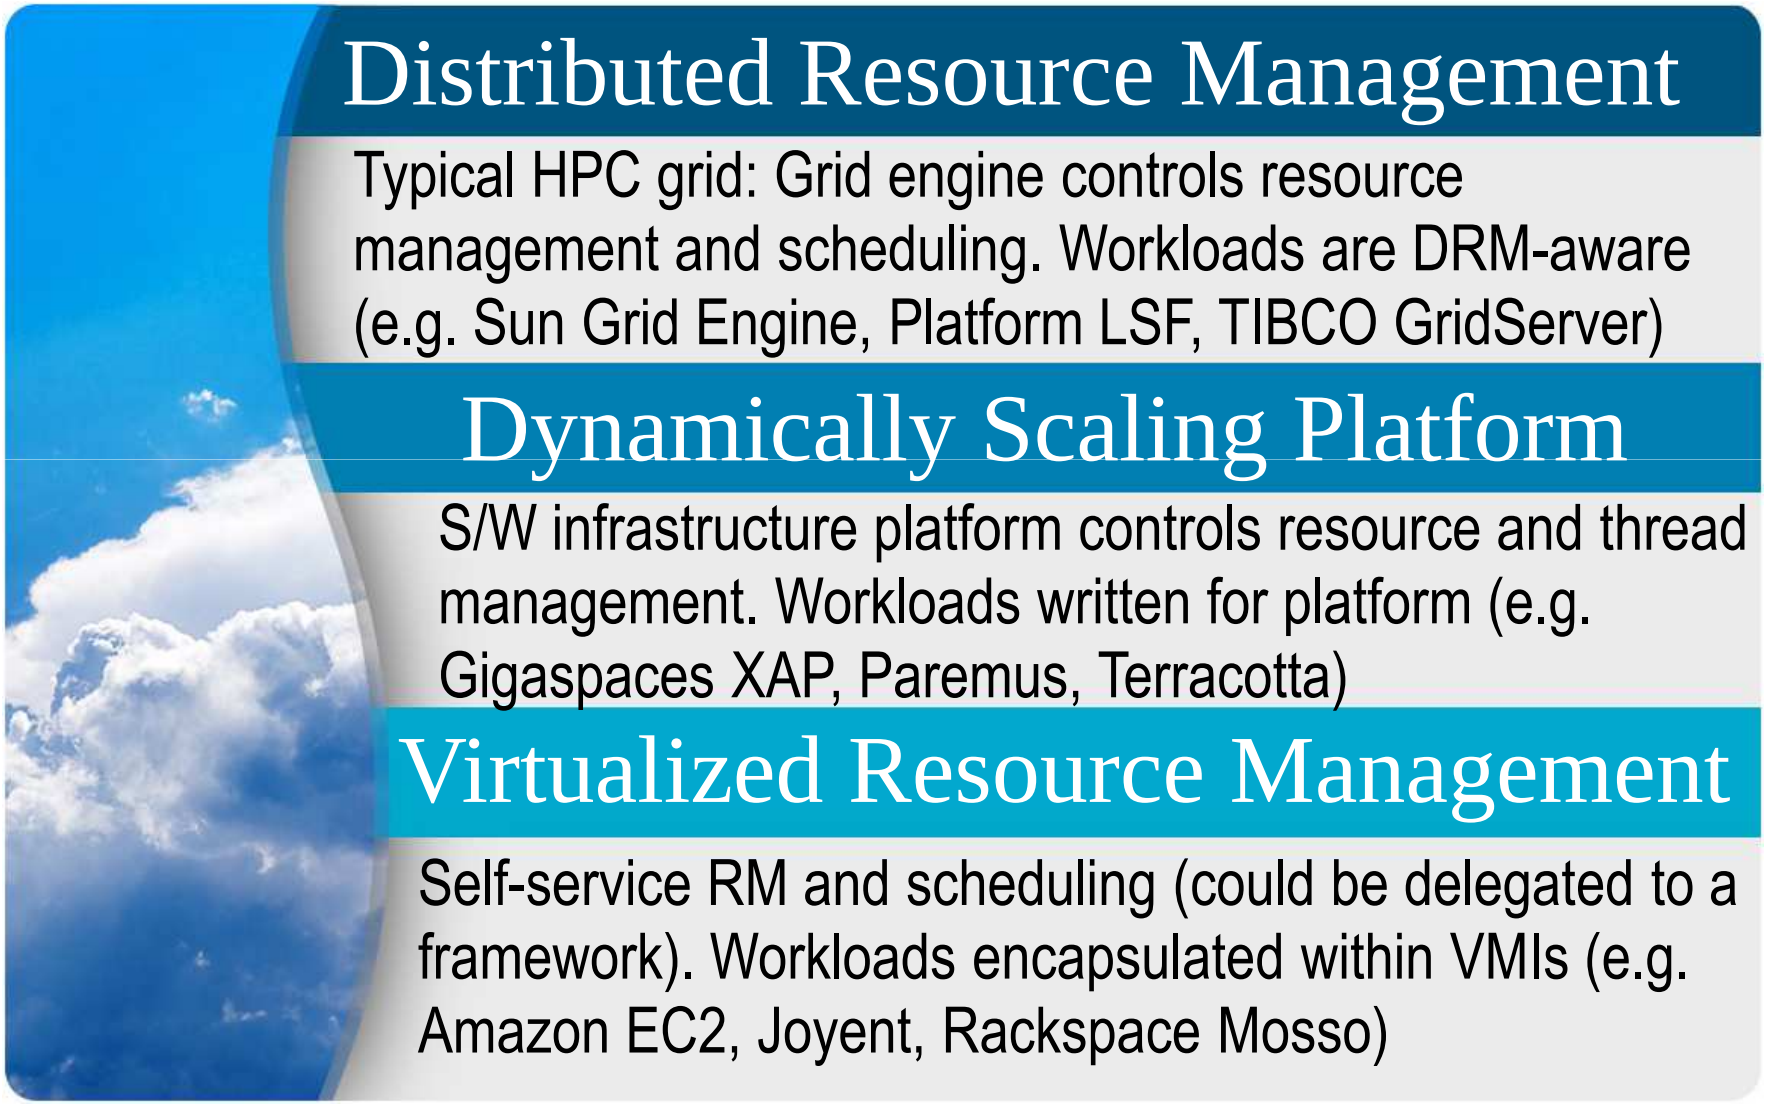

Cloud Service Models

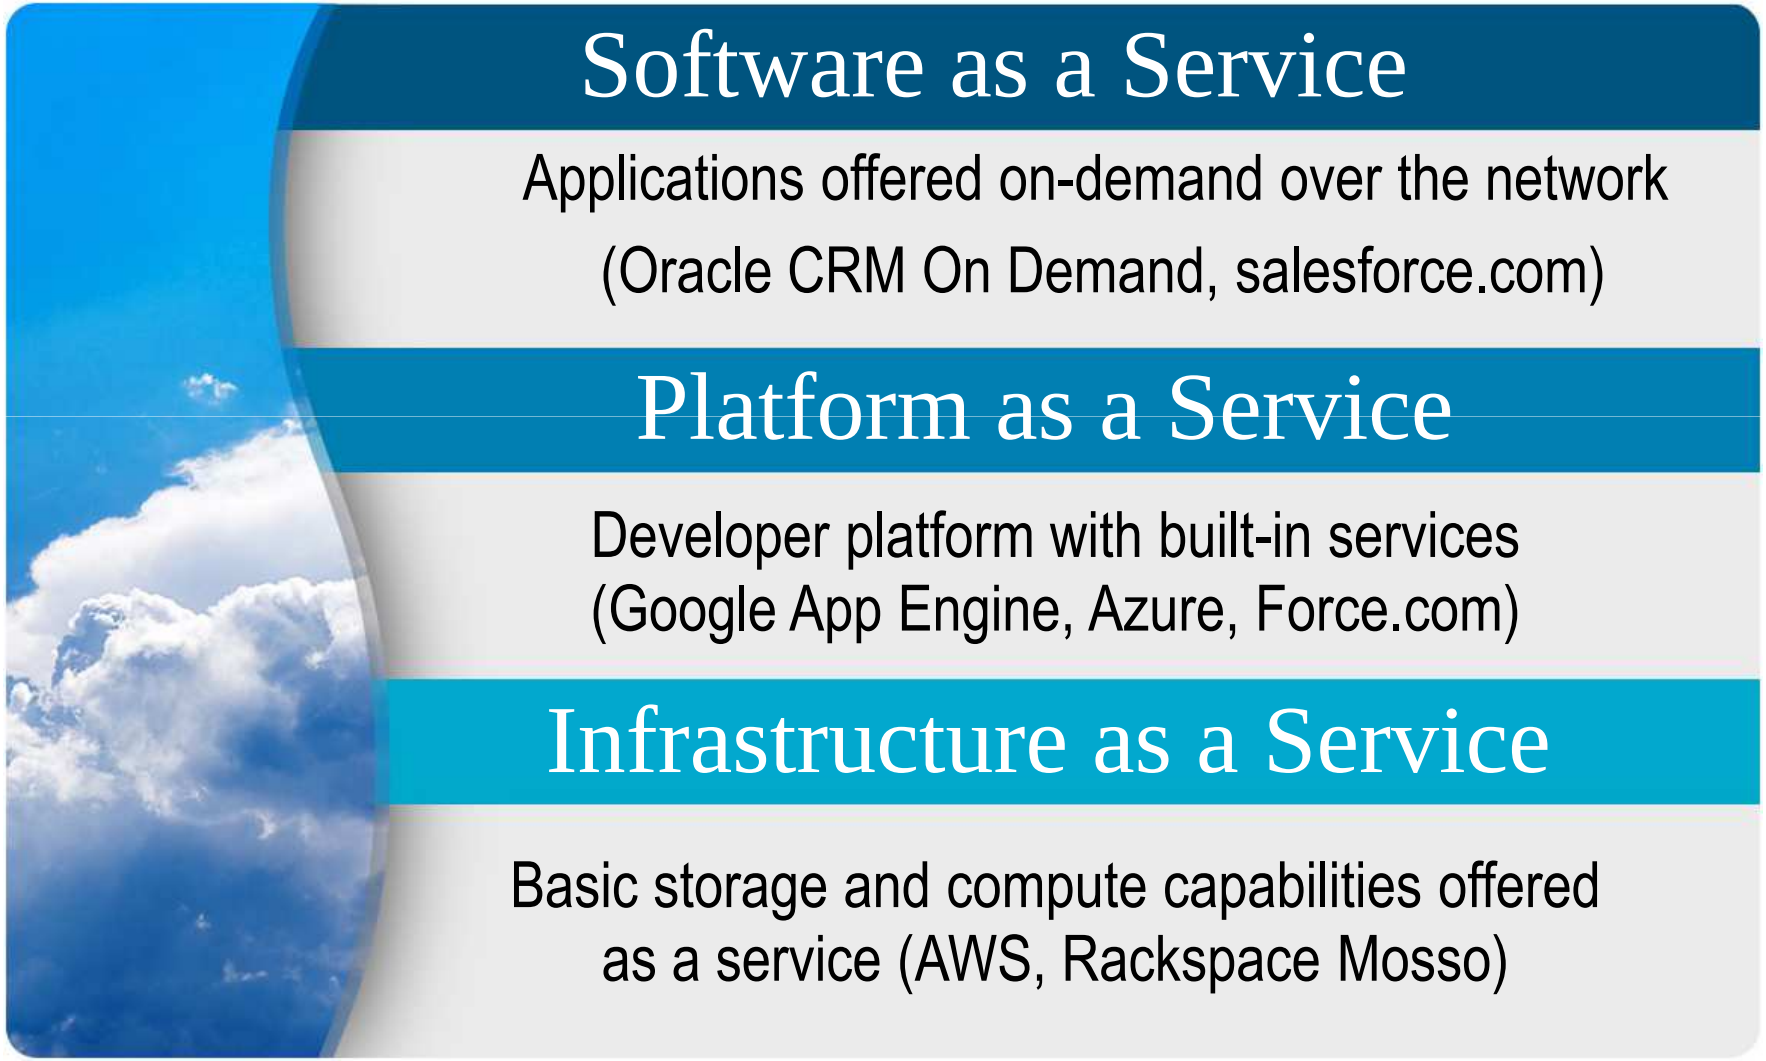

Cloud Ownership Models

Public	Hybrid	Private
		
You don't know who else is on the same servers, networks and disks that you are	e.g.: A private cloud operated out of the resource pool of a public cloud	You own the server, network and disk, and decide who gets to run on it

Sample Use Case: „Cloud Bursting“

- Sun Grid Engine: Maximize resource utilization and horizontal scalability for a wide range of workloads
  - Batch, parallel, parametric, interactive and distributed services
- Bundled components:
  - SGE “Core” – workload and resource management
  - SGE Inspect – Monitoring & management console
  - ARCo – Accounting and Reporting console
  - SDM – Service Domain Manager
    - Power Saving
    - Multi-Clustering and more
    - Cloud connectivity

Sample Use Case: „Cloud Bursting“

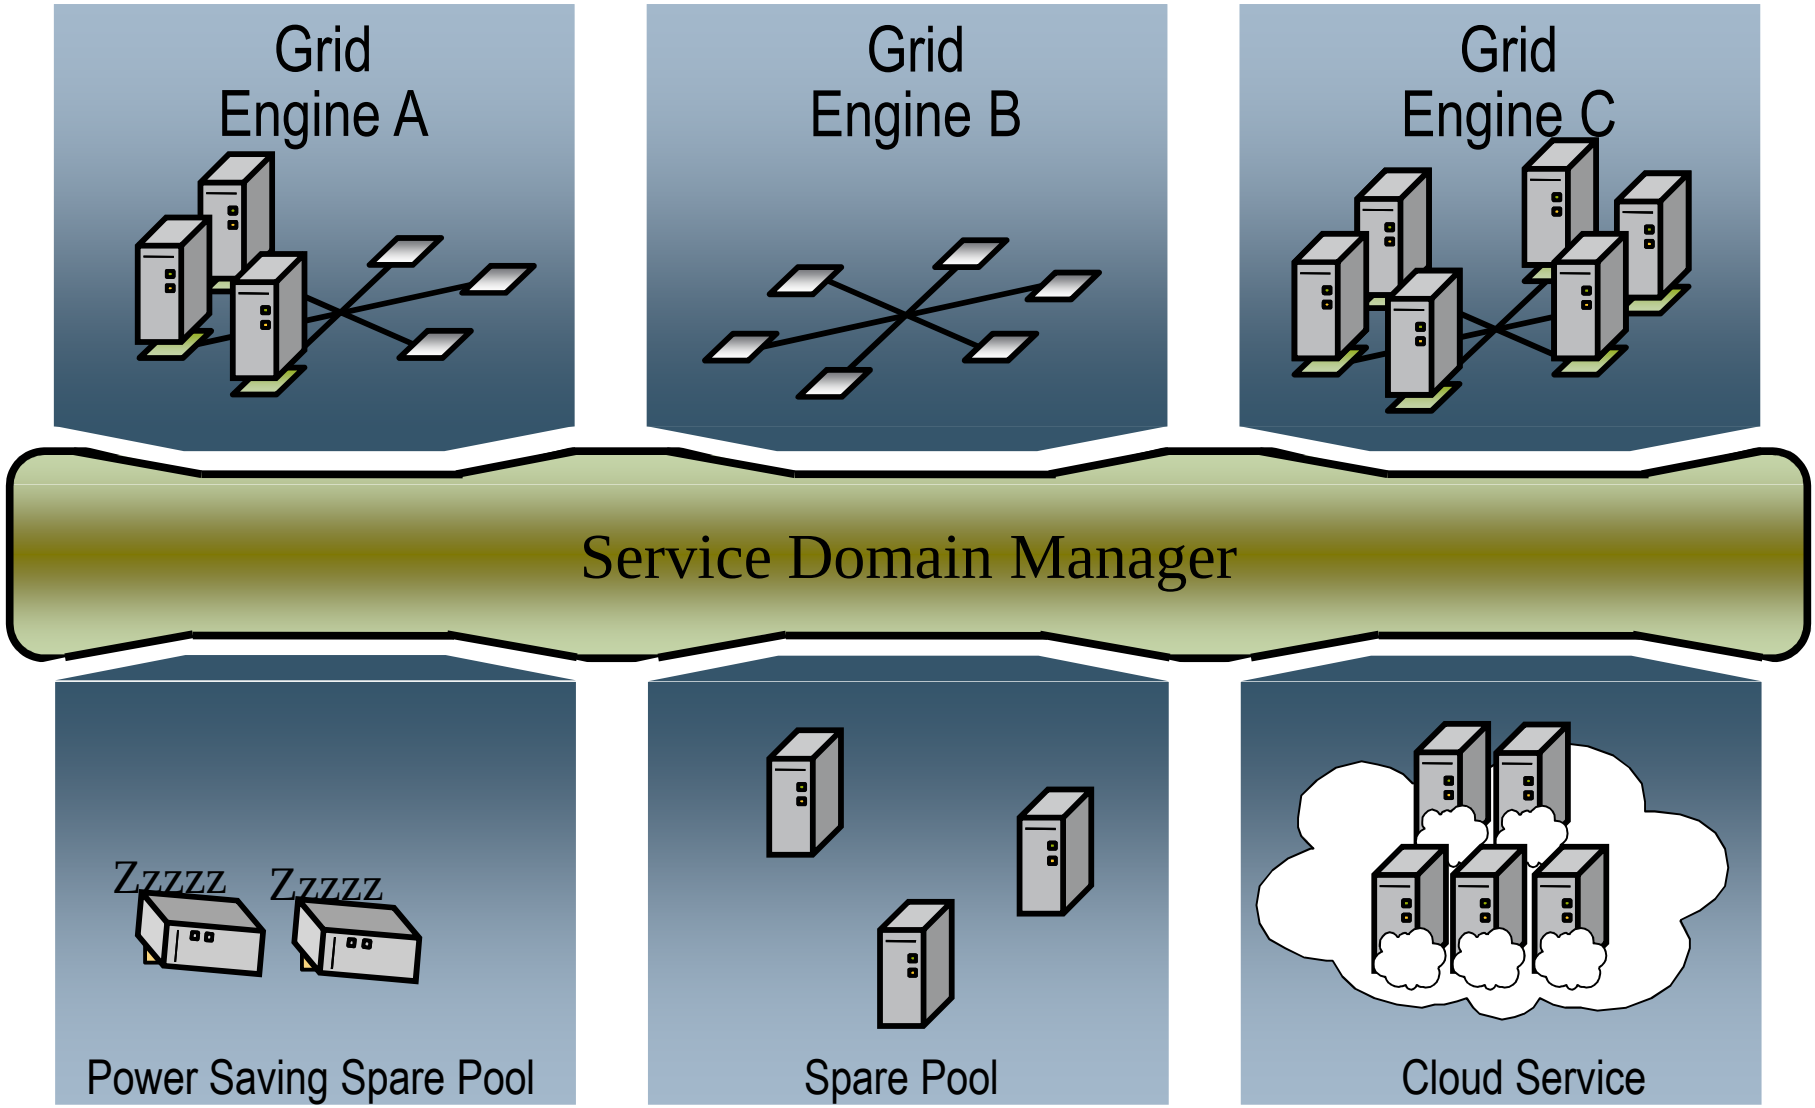

All Clouds Share Key Traits

Sun's View

Building an Enterprise Cloud

Cloud Building Blocks

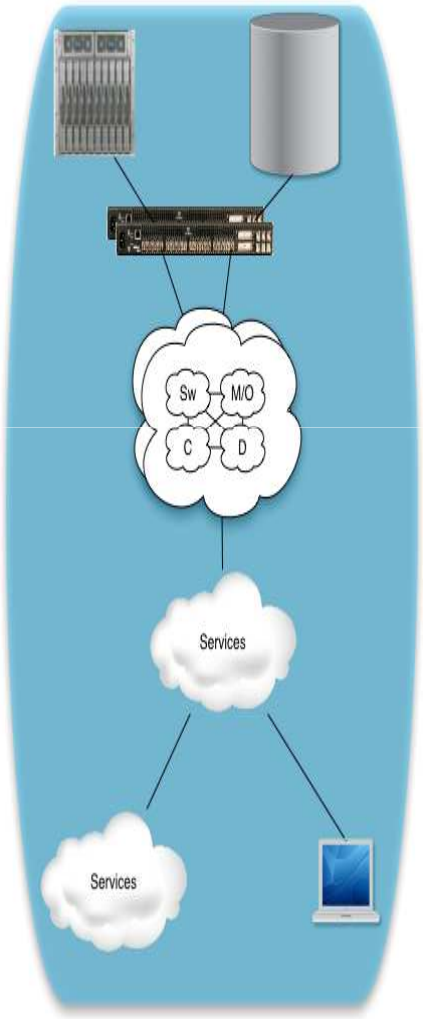

built on NetBeans

MySQL

zembly

hadoop

Joyent

Platforms and Services

OpsCenter

DCRM (aka CMDB)

OpenESB
The Open Enterprise Service Bus

GRID ENGINE

Management & Orchestration

solaris

VirtualBox

vmware

opensolaris

Virtualization

IB

zeus

7000 Series

ULTRASPARC DRIVEN

AMD 64 Opteron

intel

eco friendly

FORCE10

Infrastructure (Networks, Servers, Storage, Appliances)

Summary

- Many clouds: usage will drive redundancy models
- Rich ecosystem evolving
- Architectures are changing
- Developers are impacted
- Forget 9s, think parallel distribution

More Bits

- sun.com/cloud
- Sun Grid Engine

http://www.sun.com/software/sge/

- Security projects:
  - Cloud Safety Box
    http://kenai.com/projects/s3-crypto/pages/Home
  - Hardened VMIs
    http://blogs.sun.com/ec2/entry/hardened_opensolaris_2008_11_on

Questions/Comments?

andreas.doerr@sun.com
amir.halfon@sun.com