



Cloud Computing & Storage

More than a File System in the Sky

Erik Riedel, PhD
Technology & Architecture
Cloud Infrastructure Group
EMC



EMC²
where information lives[®]



HOME

PRODUCTS

SUPPORT

United States

English

Search

About | Headlines | Where to Buy | Partners | Register | Subscribe

Products

- ▶ Desktop Hard Drives
- ▶ Portable Hard Drives
- ▶ Multimedia Drives
- ▶ Home Network
- ▶ Network Storage
- ▶ REV Removable
- ▶ Zip / Floppy
- ▶ Software
- ▶ Online Storage
- ▶ Data Recovery
- ▶ Factory Outlet

» Solutions

» Product Selector



SKIN industries
Skin Hard Drives
BY IOMEGA

Home Users Office Users Creative Professionals



Take your PC
Virtually Anywhere



A Star is Born!
ScreenPlay Director



New eGo USB
with Protection Suite



eGo Gold Edition
Spring Sweepstakes



It's more than a computer. It's your life.



[View the Mozy intro video](#)

Every photo.
Every document.
Every song.
Safe with Mozy.

When you think about it, all the important information in your life is now stored on a computer. Whether it's photos and music or business documents and financial records, everything is digital. With Mozy, you can be sure your digital life will always be there when you need it.



Backing up a home computer or an entire office?

[Mozy has a plan just your size.](#)

[Learn More](#) ▶



[Sign up for Mozy](#)

Cloud Computing

Search

Photos Groups People

Everyone's Uploads

cloud computing

SEARCH

Sort: Relevant | Recent | Interesting




From Ivan Walsh



From Kevin Krejci



From louisvolant



From Black Dust...



From Saxoji




From Gemma...



From Combined...



From ttnk




From michaeljung



From Jonathas...



From naza1715



From Kris Barz



From Ronald Lewis



From jimluce



From web2events

From Si1very



From WELDFoto



From pst418

Sort: [Relevant](#) | **Recent** | [Interesting](#)

View: **Small** | [Medium](#) | [Detail](#) | [Slideshow](#) 



From EAHIV



From termie




From Luke...



From GraemeThicki...



From GraemeThicki...



From GraemeThicki...



From GraemeThicki...



From GraemeThicki...



From GraemeThicki...



From GraemeThicki...



From GraemeThicki...



From GraemeThicki...



From GraemeThicki...



From GraemeThicki...



From GraemeThicki...



From GraemeThicki...



From GraemeThicki...



From GraemeThicki...

[← prev](#) | [1](#) | [2](#) | [3](#) | [4](#) | [5](#) | [6](#) | [7](#) | ... | [next →](#)

(5,245 results)

Sort: [Relevant](#) | [Recent](#) | [Interesting](#)

View: [Small](#) | [Medium](#) | [Detail](#) | [Slideshow](#) 



From PatrickSmith...



From jacii...




From Stuck in...



From Stuck in...



From [Kane]



From enwie



From [Kane]



From mikebaird



From Ivan Walsh



From Kevin Krejci



From IshtiaQ...



From Robert...



From Robert...




From 1f2frfbf



From Amphinomus



From lapsuskalama...



From Kraetzsche...



From Rétrofuturs...

[← prev](#)

[1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) ...

[next →](#)


(5,245 results)

Supporting the Shift to Cloud Inside, Outside, and Across Organizations

Cloud is a model for enabling **convenient, on-demand** network access to a **shared pool** of configurable computing resources (e.g. networks, servers, storage, applications) that can be **rapidly provisioned and released** with **minimal management effort** or service provider interaction



Infrastructure deployed and operated exclusively for an organization or enterprise



Composition of two or more clouds, private and/or public



Infrastructure made available to general public or many industry groups/customers

Source: *National Institute of Standards and Technology, V15 October 2009




Courtesy - wordle.net

Who Is It Really For



VS.



IT Managers

Programmers

Programmers buzz – Ruby/Rails, MapReduce

IT Managers buzz – VM images, vApps, VLANs

Marketing buzz – Virtualization, IaaS, PaaS, SaaS

The previously separate roles of software developer and operations have [become] increasingly intermeshed and intertwined. Things are materially different...

Ray Ozzie, Chief Software Architect, Microsoft



+



IT Managers

Programmers

Programmers buzz – Ruby/Rails, MapReduce

IT Managers buzz – VM images, vApps, VLANs

Marketing buzz – Virtualization, IaaS, PaaS, SaaS

Cloud is often an “excuse” for enterprises to move to “New IT” – away from the old client/server model that has been used for the past ten years [toward Web 2.0 IT]

Werner Vogels, CTO, Amazon

A Few Details

It's not possible to "start over" and re-write all applications using scale-out design patterns in the first few months of a cloud deployment, but it is possible to adapt many legacy applications with the help of virtualization, so cloud infrastructure can support and enable both development models, including mixing the two.

“Developers” Range Widely in Focus/Expertise

- **IT managers/admins** deploying applications encapsulated or pre-packaged into virtual machines
 - Language – configuration scripts, command lines
 - Input – catalog of vApp templates or pre-configured VMs
 - Output – VM images, VM configurations, system configurations
 - Runs on – vSphere/ESX, virtual networks, legacy storage + scale-out storage
- **Programmers** using application frameworks such as Groovy/Grails
 - Language – Grails/Java, Python/MapReduce
 - Input – code (perhaps with help of an IDE)
 - Output – Rails + database configurations, job scripts
 - Runs on – Rails + MySQL, virtual networks, scale-out storage



Apps + Data

- **Development**

- new applications
- explicitly scale-out (e.g. MapReduce, Hadoop)
- built on higher-level frameworks (e.g. Ruby/Rails, Azure)

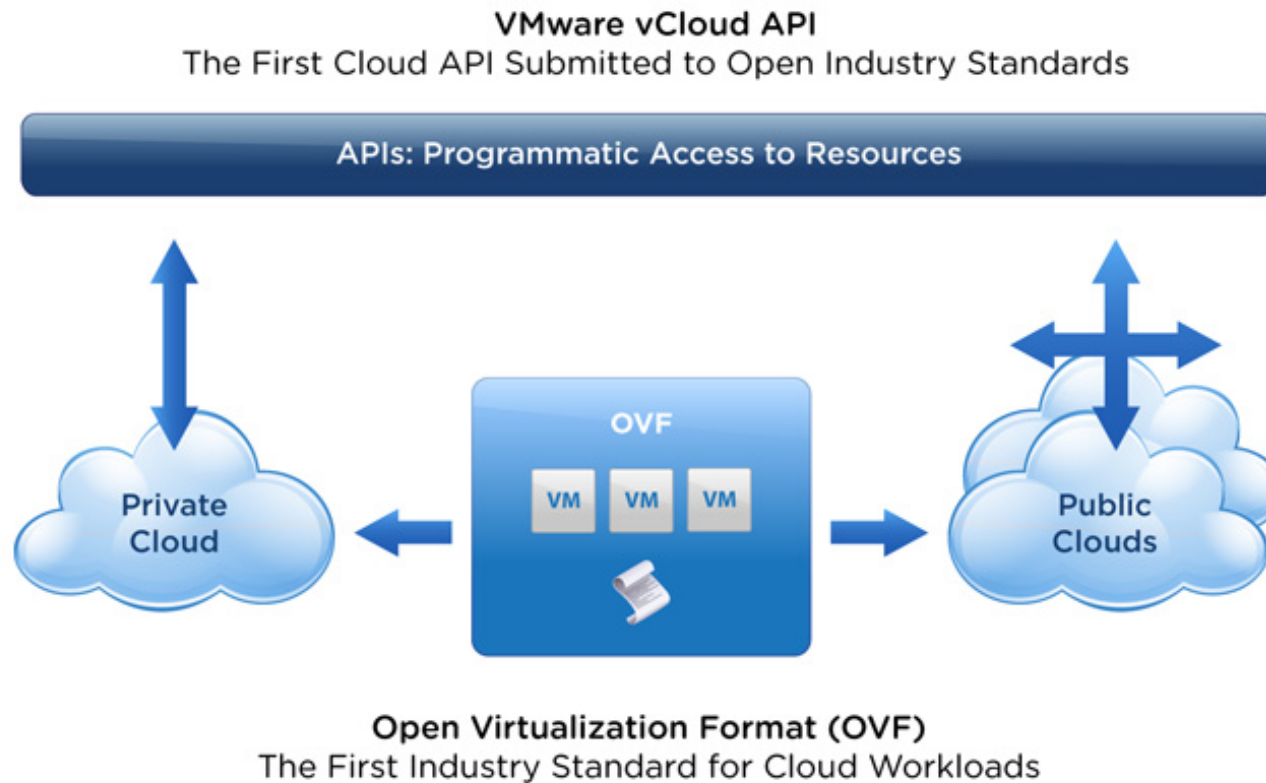
- **Deployment**

- legacy applications
- “packaged” into virtual machine containers
- easy to replicate and migrate across virtual infrastructure

- **Data**

- shared corporate data is the common ground (enterprise apps)
- consumer value centered around their personal data (consumer apps)

Example – Deployment



Marketing buzz – IaaS – Infrastructure as a Service

Example – Development




The screenshot shows the Cloud Foundry website homepage. At the top right is the VMware logo. Below it are links for [Blog](#), [FAQ](#), [Community](#), and [Forums](#). A diagonal banner on the left says "BETA". The main heading is "CLOUD FOUNDRY" with the tagline "No Obstacles: Deploy and Scale Your Applications in Seconds". Below this is a description: "The industry's first open platform as a service. Run your Spring, Rails and Node.js applications. Deploy from your IDE or command line." A large blue button says "Free Sign Up for Cloud Foundry". To the right of the button, text reads: "Request an invite for a CloudFoundry.com account and get notified when your Micro Cloud is available for download".

Marketing buzz – PaaS – Platform as a Service

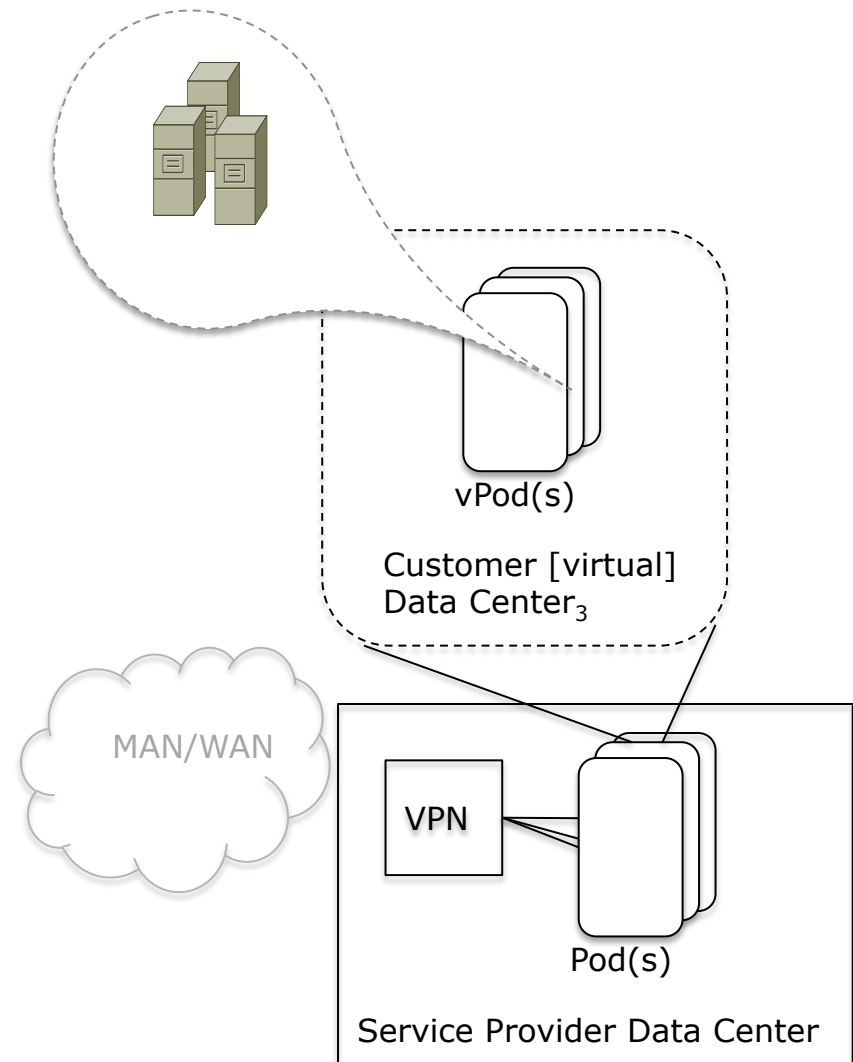
More About Apps + Data

- From the perspective of development & deployment, the key new technology component is a combined data + app (storage + compute) platform where apps are created, deployed, monitored & managed with a common set of tools.
 - *Underlying enablers:*
- Common object space – apps, configs, user data
- Single identity store – public, private, enterprise, consumer
- Federation (public + private) – seamless across infrastructures
- Monitoring – continuous measurement to optimize (and generate bills)



Key Technology Components

- Policy-driven Orchestration
 - Application mgmt – via virtualization
 - Data mgmt – ILM is finally required
 - Continuous measurement & monitoring – to meter/bill; to maintain high efficiency
- vPods (virtualization)
 - Enables easy migration and replication of containerized applications
 - Drives highly efficient resource utilization
 - Eases rapid deployment of new applications & new services
- Pods (packaged racks)
 - Rack-level deployment of infrastructure (compute + network + storage)
 - Drives highly efficient acquisition and deployment vs. traditional full custom or semi-custom design-per-app



Under The Covers

What About The Data?

Atmos Deployment Models



Purpose-built hardware

- Standard components—factory configured, racked, and cabled
- Flexible, cost-effective service model/user-serviceable
- Optional high availability and redundancy; 30 drives per 3U drawer

WS2-120



Small/Medium Compute

WS2-240



Dense Compute

WS2-360



Capacity

Intel Xeon 5500 "Nehalem" architecture

2 TB low-power SATA drives (lower power per BTU at two-times capacity)

- | | | |
|---|--|--|
| <ul style="list-style-type: none"> • 1:15 servers-to-disks • 60 TB cloud entry point • Up to 240 TB per rack | <ul style="list-style-type: none"> • 1:15 servers-to-disks • Up to 480 TB per rack | <ul style="list-style-type: none"> • 1:60 servers-to-disks • Up to 720 TB per rack |
|---|--|--|

Software on VMware

- Certified with EMC Unified Storage (NFS, FC, iSCSI); VMware-supported servers and third-party storage

Atmos Virtual Edition



VMware vSphere

NFS/FC/iSCSI Storage

- Starts at 10 TB
- Up to 960 TB per site



Changing View

- Cloud storage
 - Goal: low(est) cost, tolerate high(er) touch
 - Mask lower hardware reliability with clever software + wetware
 - More limited range of performance/availability (aim for “good enough” for a majority of apps, with some variance)
 - Focus on OpEx (operating costs, human costs)
- Enterprise storage
 - Goal is low(er) touch, tolerate high(er) cost
 - Expect high hardware reliability
 - Wide range of performance/availability trade-offs (optimize for e.g. OLTP all the way to long-term archive)
 - Focus on CapEx (purchase costs, human costs)

Move from higher cost, low touch => lowest cost, higher touch

Reasonable cost investments:


- serviceable components where sensible (e.g. disks/power supplies)
- leverage HA components where sensible (e.g. rack power, connectivity)

Under The Covers

What Is Different?

Today's IT Structure – and Tomorrow's

Traditional Storage





	Today	Tomorrow
Workload	Transactional	Content
Application interaction	Mount point	API
Data management	Admin-driven/ manual	App-driven / policy automated
Data protection / Disaster protection	Additional best of breed products	Built-in data services
Storage scaling and management	File system limits / pre-planning	Unlimited namespace / dynamic
Infrastructure model	Single tenant / storage isolation	Native multi- tenancy throughout the stack

Active Object Storage



Builds on 10+ Years of Storage Research



- APIs vs. mount points – “no slashes required”
 - blocks vs. files vs. objects vs. “APIs”
- App-driven and policy-automated 
 - self-configuring, self-organizing, self-tuning, self-*
- Built in data services 
 - self-healing
- Unlimited namespace, dynamic
 - billions and billions of objects, large and small
- Native multi-tenancy
 - security/auth, monitoring, resource isolation

Summary

Pain Killers, Not Vitamins – Solve Real Problems vs. “Nice To Haves”

Enterprise IT challenges/pain points

- Adapting to the business model changes of cloud
- Answer: private + public clouds with federation
- Adapting to development model changes of cloud
- Answer: leverage new tools, frameworks to develop Web 2.0 and scale-out apps
- Migrating legacy applications to cloud
- Answer: virtualization to encapsulate legacy OS + apps
- Managing data across apps & users – governance
- Answer: a combined + app platform to manage the data flow among apps and virtual machines



Questions?

References

- Geoff Moore *“Partly Cloudy: Business and Innovation in the Internet Era”* September 2010
 - www.snia.org/cloud/Cloudburst/Moore_SNIA_Keynote.pdf
- Peter Mell & Tim Grance *“The NIST Definition of Cloud Computing”* October 2009
 - csrc.nist.gov/groups/SNS/cloud-computing/cloud-def-v15.doc
- Any business or computing magazine published at any point in 2009 or 2010

Why The Cloud Is Here To Stay

Enterprise vs. Consumer Technology

Another angle – cloud computing is really about bringing enterprise computing technology and applications up to the norms and expectations of consumer computing technology.

The way we run our lives has forever changed. The employees we are hiring right out of school are appalled by the technology we use to run our companies. They are more productive at home than they are in the office.

Marc Benioff, CEO, Salesforce.com

The barrier is becoming less and less between enterprises and consumers in application terms [expectations and functionality often are very much the same]”

Eric Schmidt, CEO, Google

The Big Disconnect

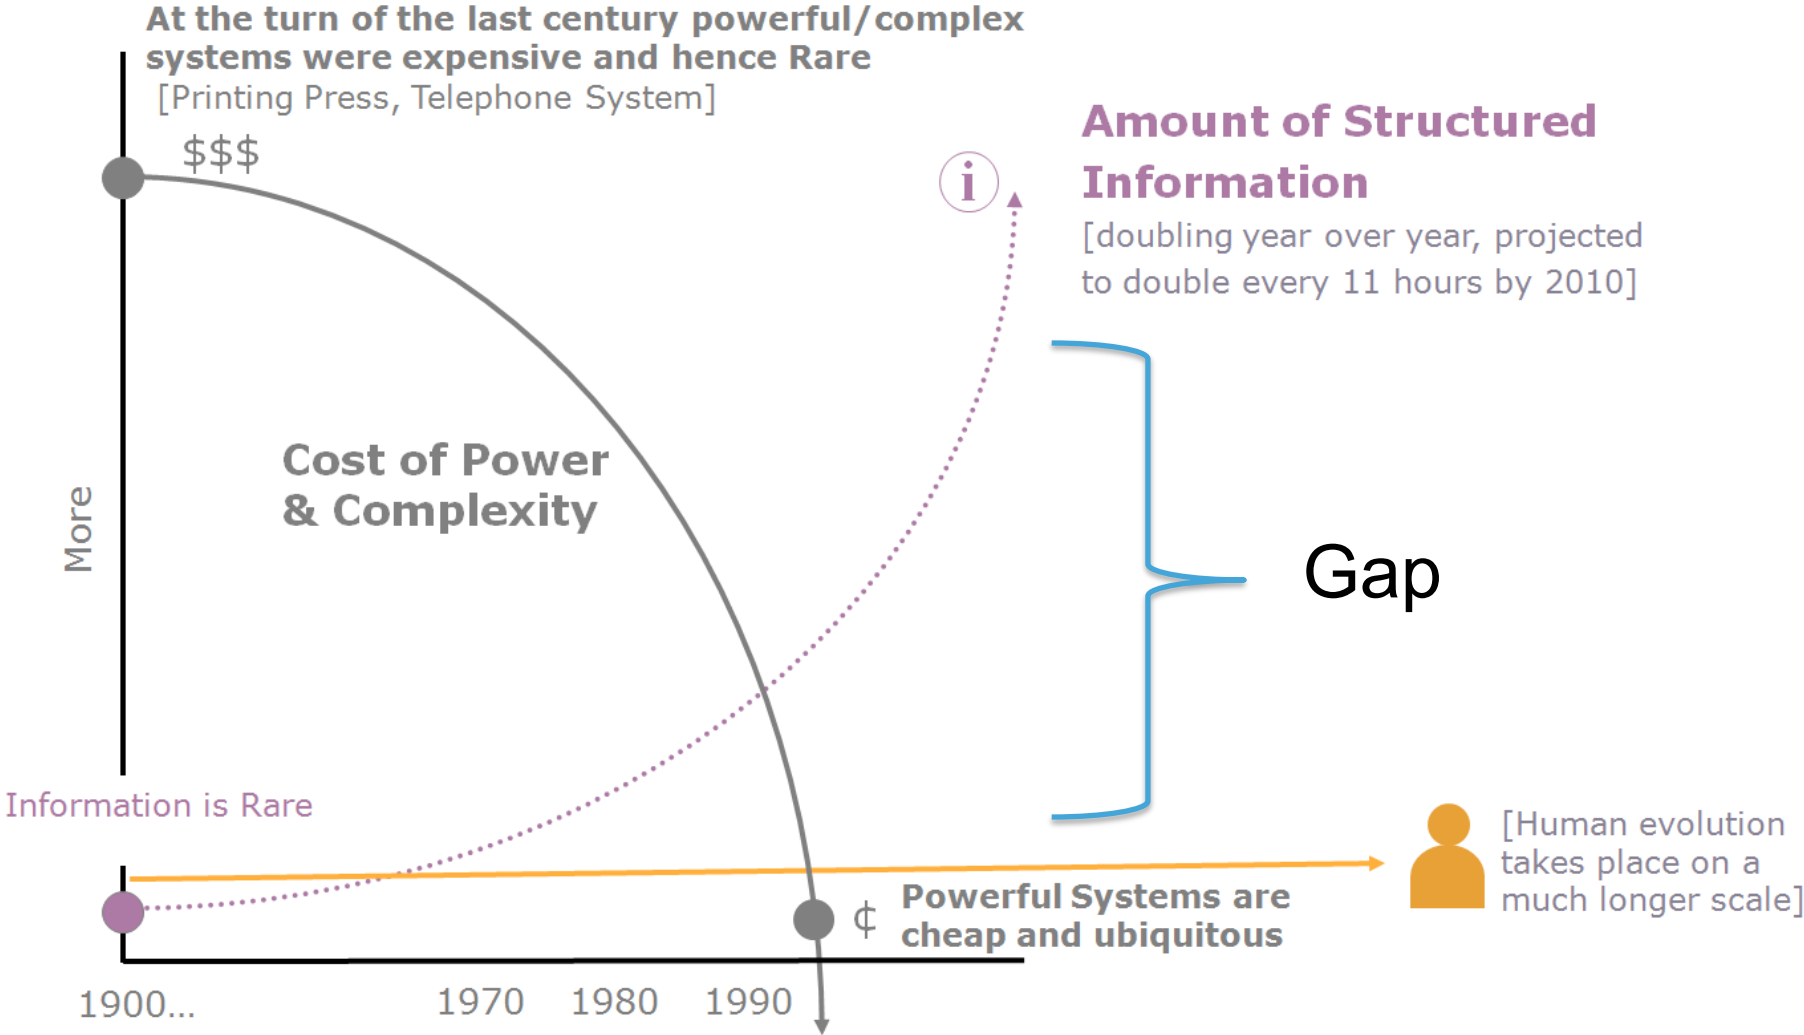
How can it be
I am so **powerful** as a consumer
And so **LAME** as an employee!!??

How disruptive do you think Consumer IT
will be to Enterprise IT?

Why should employees accept a 50% reduction in their productivity when they come to the office on Monday morning? On the weekend, Google can answer any question I have, on Monday, I can't get the answer to "who are my five biggest customers?" On the weekend, someone from my high school can find me and try to be my friend, on Monday, I can't find my VP of Finance.

Geoff Moore, Author, Crossing the Chasm

Consumer Attention


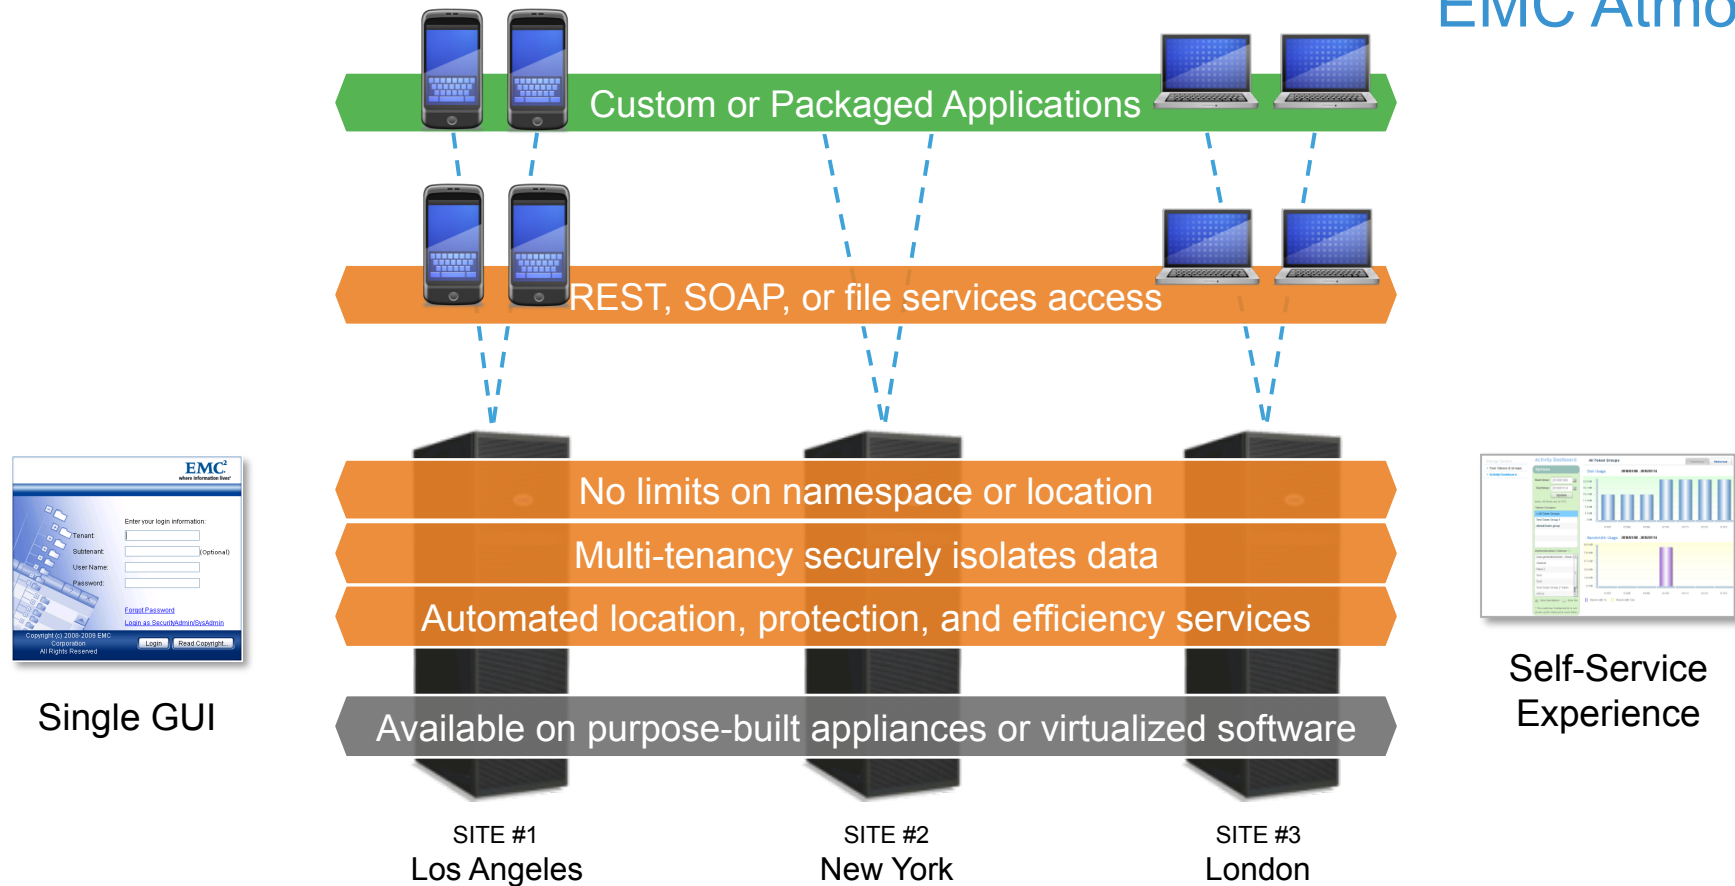



Courtesy Mick McManus, MAYA Design

Atmos

Product Details

What is EMC Atmos?



Single GUI



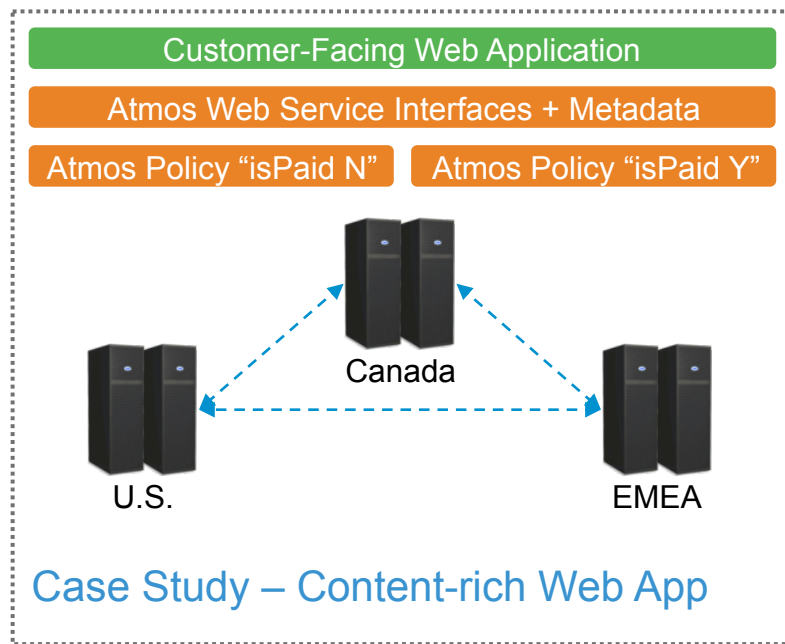
Self-Service Experience

SITE #1
Los Angeles

SITE #2
New York

SITE #3
London

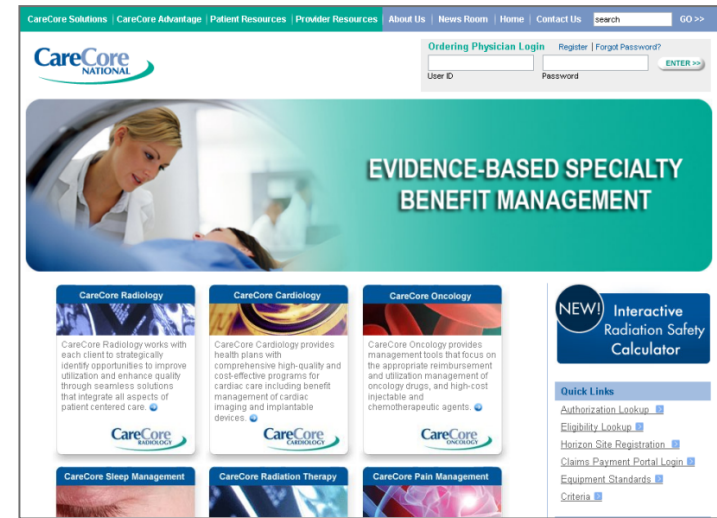
Case Studies



Content-Rich Web App on Atmos

- Global distribution, content mix
- Multi-tenancy, scale to multiple sites
- Policy supports business models

- Petabyte-scale
- Geographic distribution
- Policy-driven storage



CareCore “gets in the cloud” with Atmos

- Wrote to Atmos REST API in one week
- Bought Atmos and deployed in three weeks
- Adding over 2 million objects a day to Atmos
- Started with one app, spreading to many more

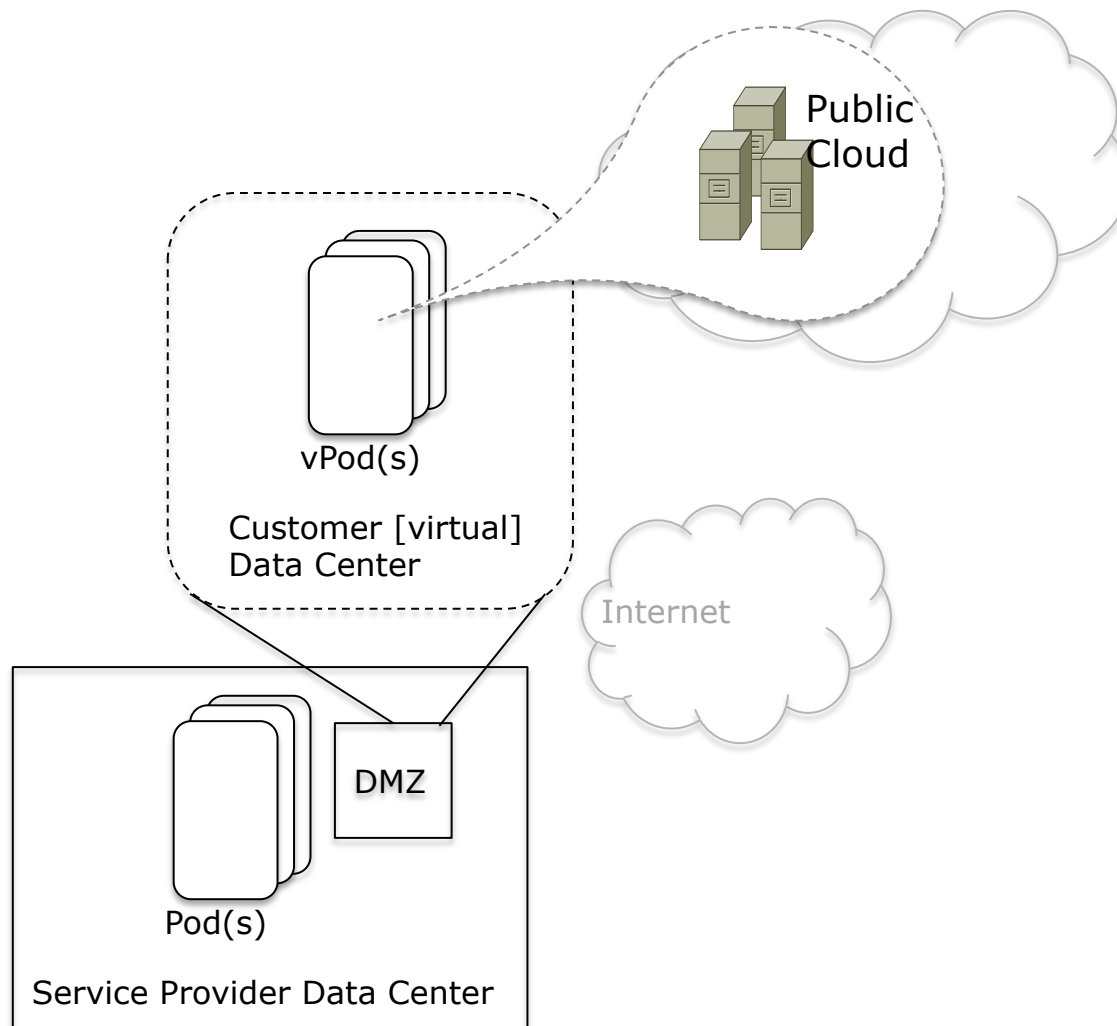
EMC²[®]

where information lives[®]

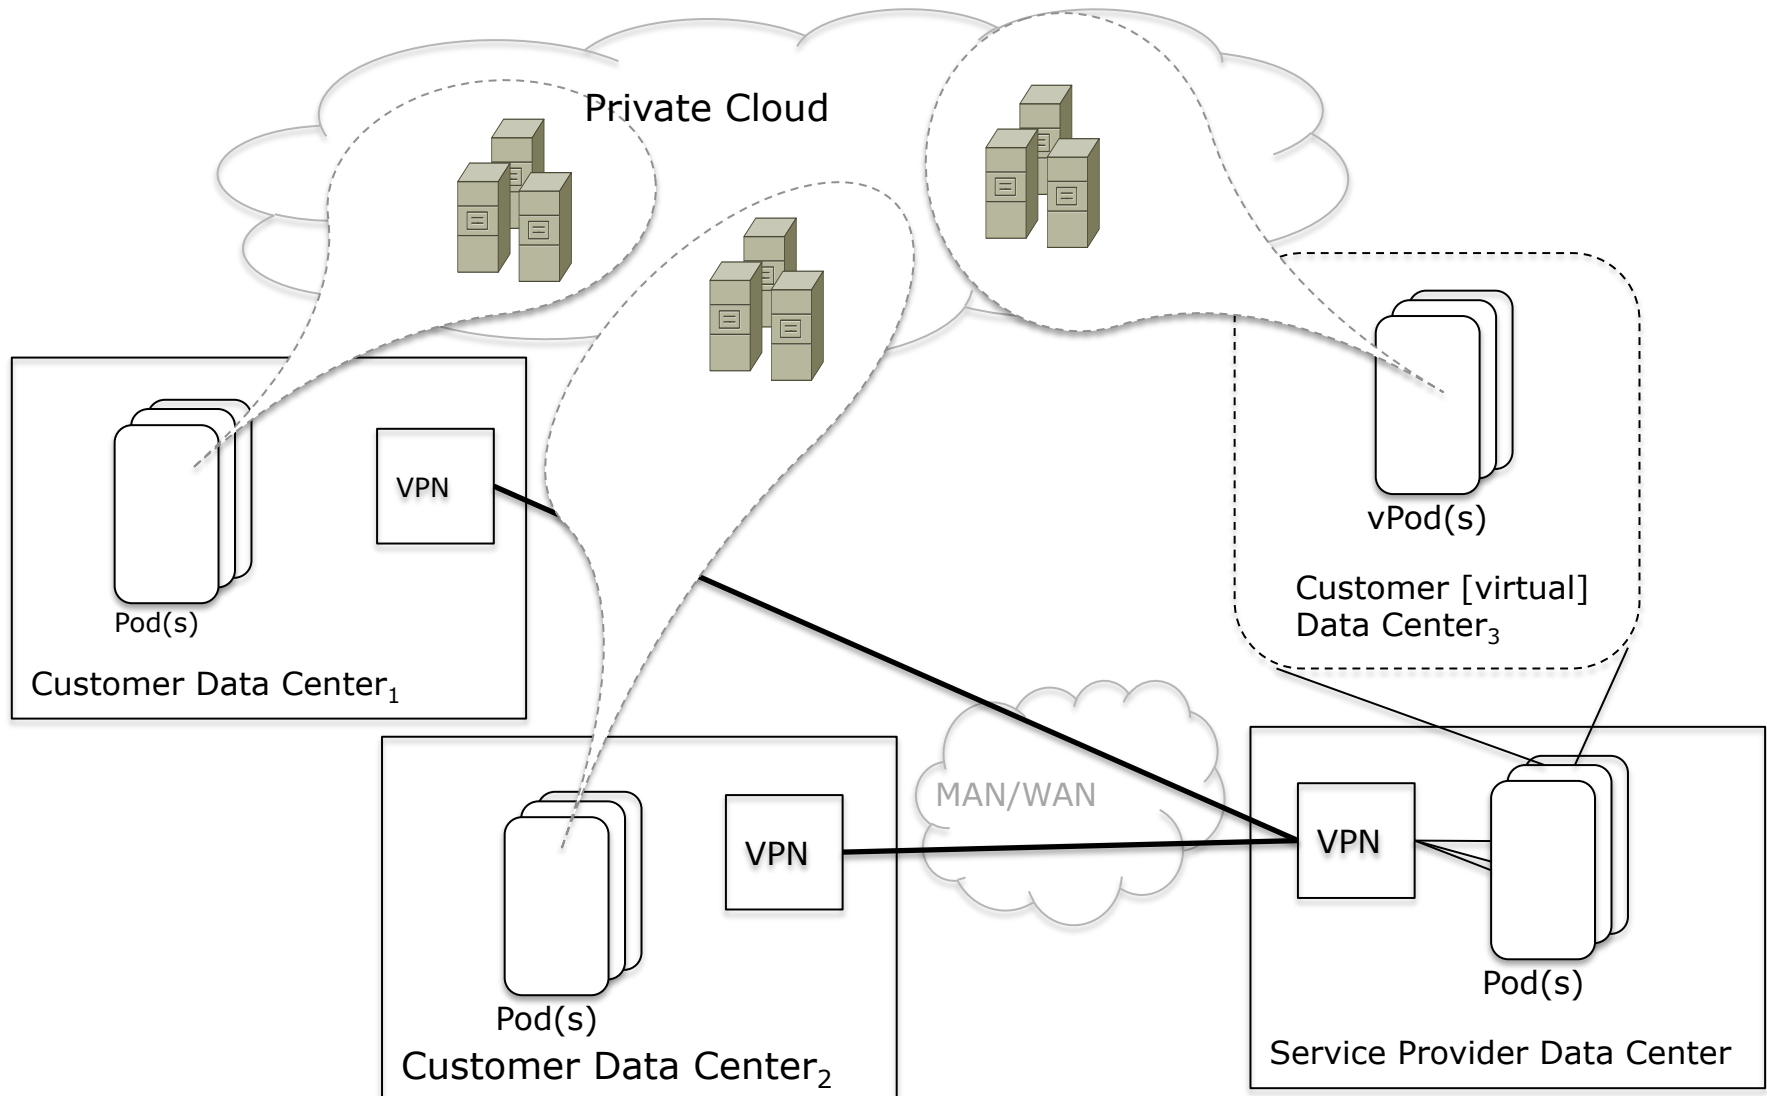
Under The Covers

Deployment Choices

Scenario #1 – Public Cloud



Scenario #2 – Private Cloud



Scenario #3 – Service Provider

