



# Cloud Storage:

## More than a File System in the Sky

Erik Riedel, PhD  
Technology & Architecture  
Cloud Infrastructure Group  
EMC

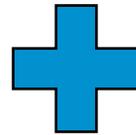
# Development vs. Deployment

“The previously separate roles of software developer and operations have [become] increasingly intermeshed and intertwined. Things are materially different...” Ray Ozzie, Microsoft, October 2009



Programmers

*Ruby/Rails, MapReduce*



IT Managers

*VM images, vApps, VLANs*

# Moving to Web 2.0 IT

“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 10 years [toward Web 2.0 IT]” – Werner Vogels, Amazon, November 2009

It is not possible to “start over” and re-write all applications using scale-out design patterns in the first 12 months of a cloud deployment, but it is possible to adapt many legacy applications with the help of virtualization, so cloud infrastructure must support and enable both development models, including mixing of the two.

# Apps + Data

- **Development**

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

- **Deployment**

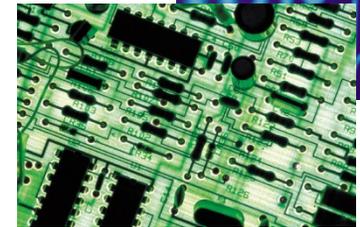
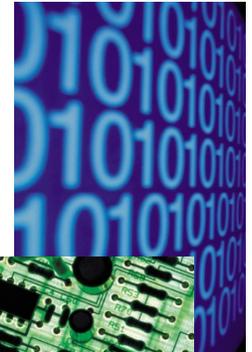
- legacy applications
- “packaged” into virtual machine containers
- easy to replicate and migrate across virtual infrastructure

- **Data**

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

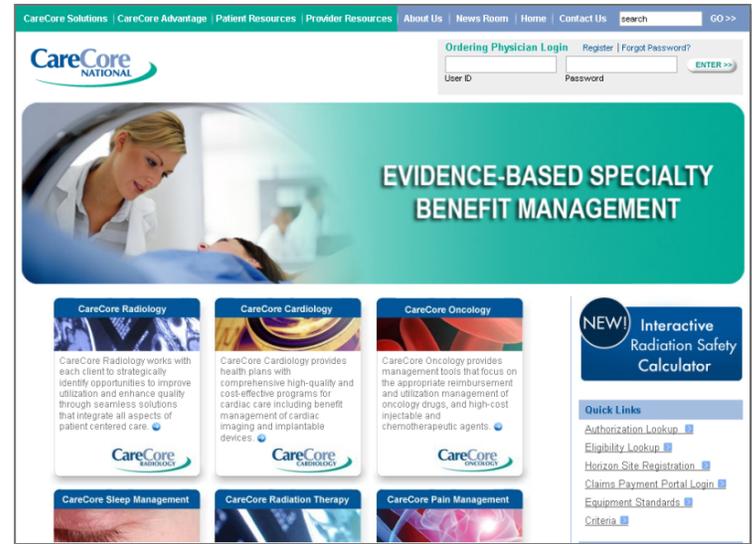
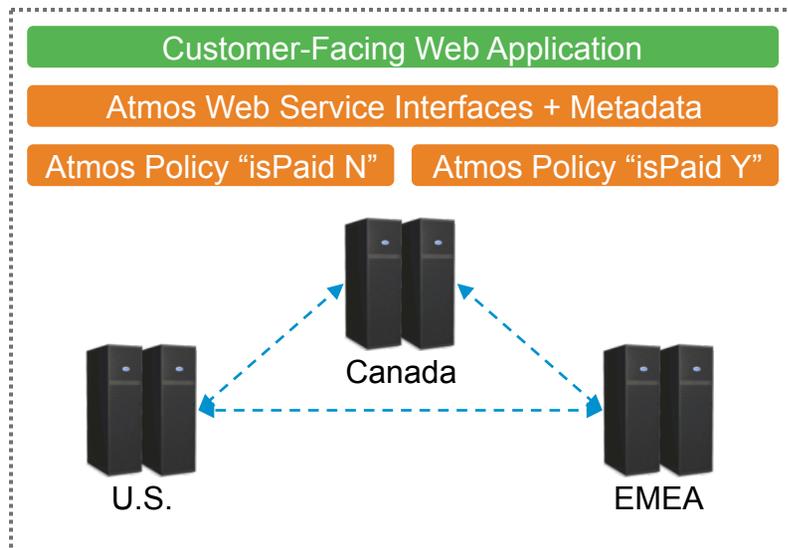
# More About Apps + Data

- 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)



# Case Studies

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



## Content-Rich Web App on Atmos

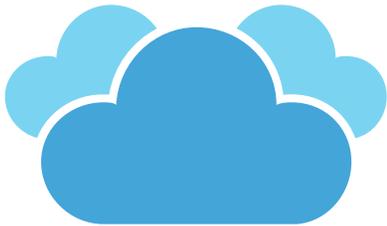
- Global distribution, content mix
- Multi-tenancy, scale to multiple sites
- Policy supports business models

## 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

# Who We Are

EMC Atmos



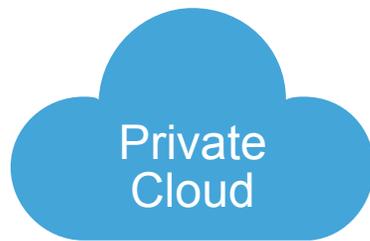
- Atmos brings cloud storage to the enterprise
  - Can be deployed as a product, consumed as a service or a hybrid of both
  - Supports massive content growth in both traditional applications and new applications
  - Enables efficient TCO and cost-to-serve
  - Has a growing ecosystem of solutions and providers

**EMC<sup>2</sup>**<sup>®</sup>

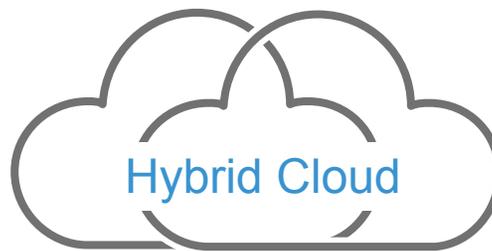
**where information lives<sup>®</sup>**

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

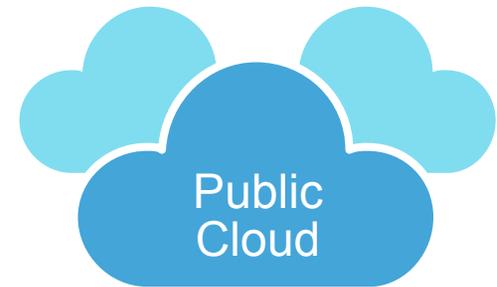
Cloud is a model for enabling **convenient, on-demand** network access to a **shared pool** of configurable computing resources (e.g. networks, servers, storage, applications) that can be **rapidly provisioned and released** with **minimal management effort** or service provider interaction



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 Oct 2009

# 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

