

Reliability in the Cloud – Beyond Infrastructure

Said Syed

IEEE – CQR 2014

May 14, 2014

Today's Cloud Computing Defined

“A standardized, highly automated offering, where compute resources, complemented by storage and networking capabilities, are owned by a service provider and offered to the customer on demand.”

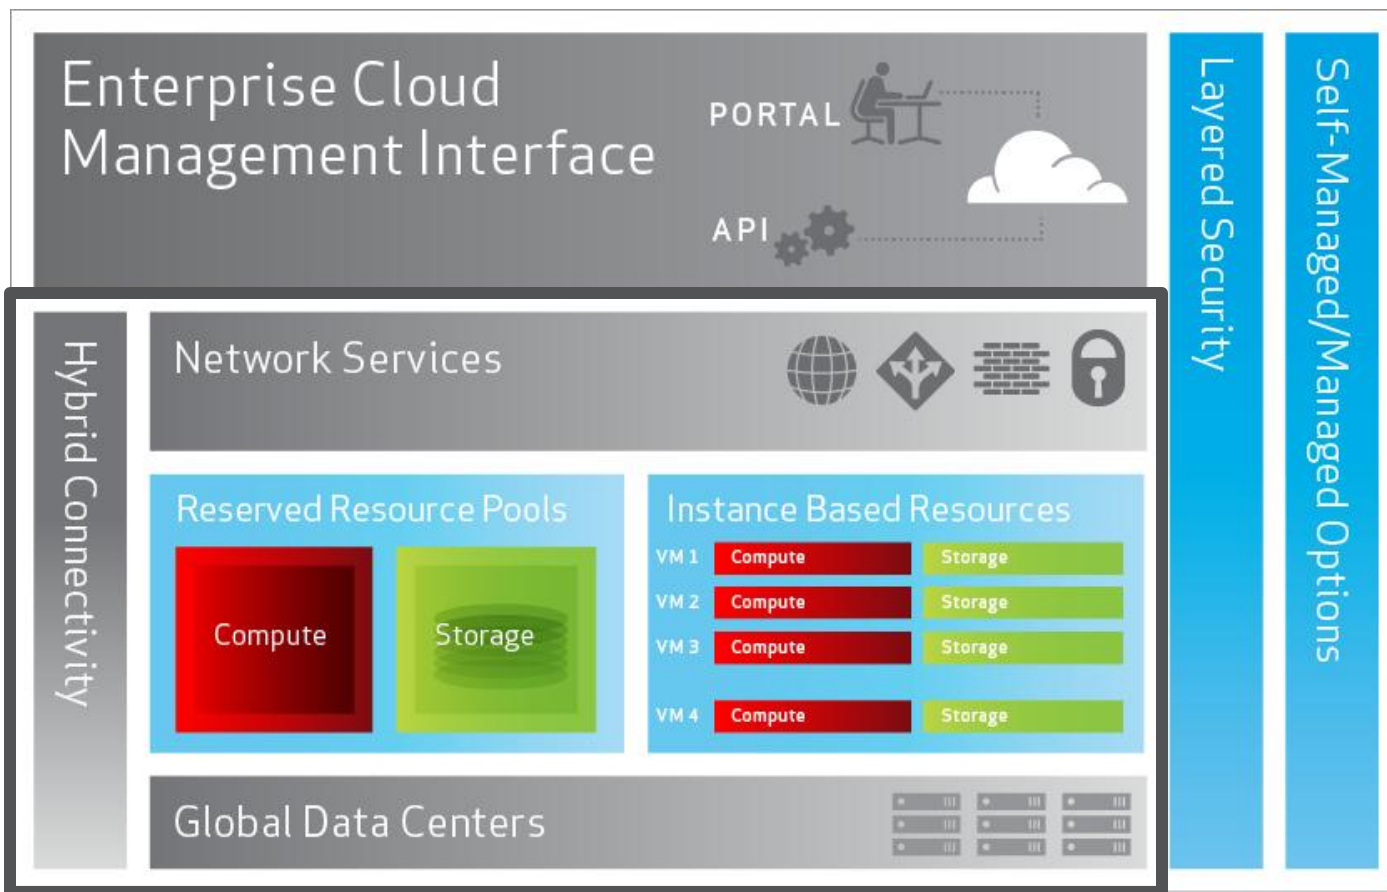
Gartner INC.

October 2012 Magic Quadrant for IaaS Report – ID:G00237002

Is this really Cloud?

IaaS vs Cloud

IaaS



But Is that enough?

What is the Unmet Need

Business Applications Require Infrastructure That's Agile, On-Demand & Elastic

Application

Agile

On-Demand

Dynamic
Infrastructure

Availability



Accessibility



Security



Operational Support



Compute



Storage



Application Monitoring
& Management

Backup & Restore

Business Continuity/
Disaster Recovery

Network Connectivity

Access Control

Application Acceleration

Design, Assessments
& Remediation

Regulatory Compliance

24/7 Monitoring
& Incident Response

Design & Build

Migrate & Test

Operate

Flexible Continuum of Services

In House

Colocation

Managed
Hosting

Private
Cloud

Public
Cloud

Hybrid Configurations

Elasticity

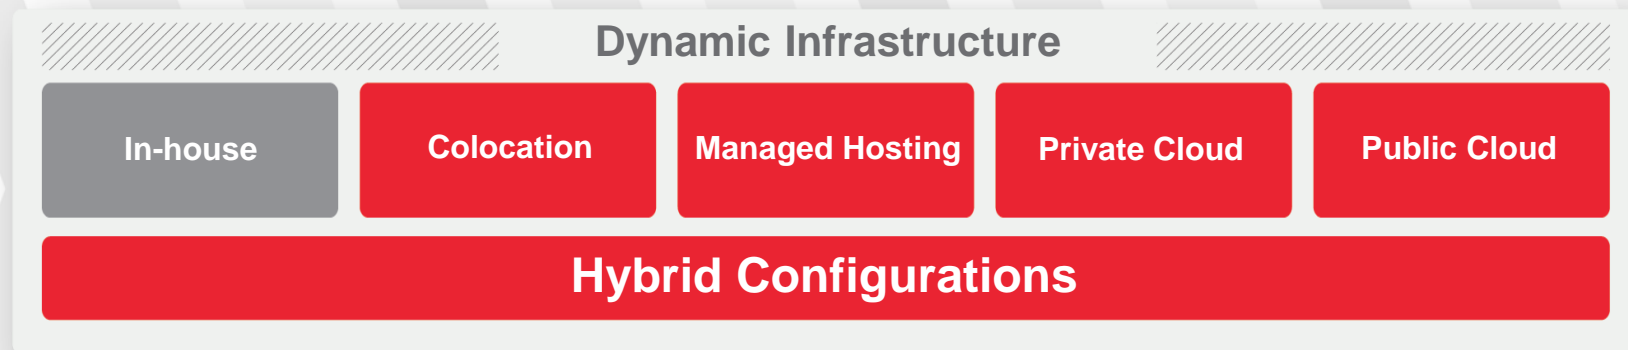
- Resources available when needed with appropriate SLAs for the workload being deployed
 - Automatic expansion and contraction
 - Workload placement algorithms
 - Local, Metro, Geo Clustering

Operational Challenges in Achieving Reliability

- Enterprise Class vs Commodity Platforms
 - Is the platform capable of handling the workload being put in?
- Ability to meet customer SLAs
 - Performance, Availability, Reliability, Serviceability
- Blurred lines of delineation between operational functions
 - Virtualization – Network, Server, Storage
 - Who is ultimately responsible?
- Workloads vs VMs (servers)
 - Cloud Providers don't care or know about the type of workload
 - So how do we design for the unknown and make the platform reliable?
 - Does knowing the workload change the definition of the cloud? To what?
- Orchestration + Automation = Dynamic Infrastructure
 - Its all about the layer above the infrastructure

Future of Dynamic Infrastructures to Achieve Application Reliability

Seamless Workload Mobility Across All Services
Must be as Dynamic as the Workload



Intelligent Infrastructure Layer

- Everything must work together to realize true Reliability
- Local, Metro and Geo Scale
 - Compute, Virtualization, Storage, Data Protection, Workload Mobility (migration)
- Continuous Improvement of SLAs in the Cloud
 - End-to-end Orchestration
 - Automation
 - Dynamic Performance Management to support workload needs
 - Workflow driven Elastic Capacity Management – Scale up/down Automatically
 - Configuration Management to support Compliance and Regulations
 - Life-cycle Management to ensure optimal availability of infrastructure
- Governance
 - Compliance and regulations to achieve required security and availability
 - Distance requirements, DR and business continuity, HIPAA, PCI, Federal

Continuum of Service

One Global Platform vs Distributed Services

