Java on the Cloud - What Java Professionals Need to Know

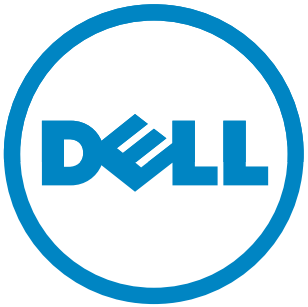

Madhuri Mandava

3rd October 2011

A Few Words About Cloud

- Service-Oriented
- Always Available
- Scale-Out
- Model-Driven
- Failure Resilient
- Federated
- Multi-Tenant
- Elastic

Why the JVM is good for the Cloud

- The Java virtual Machine is already a managed environment
  - Portability across O/S & JVM implementations
  - Remote debugging/profiling/monitoring hooks
  - Security Policies

- Base runtime for different JVM languages
  - Ceylon, Groovy, Java, JPython, Ruby on Rails, Scala etc

What does Java EE offer to Cloud ?

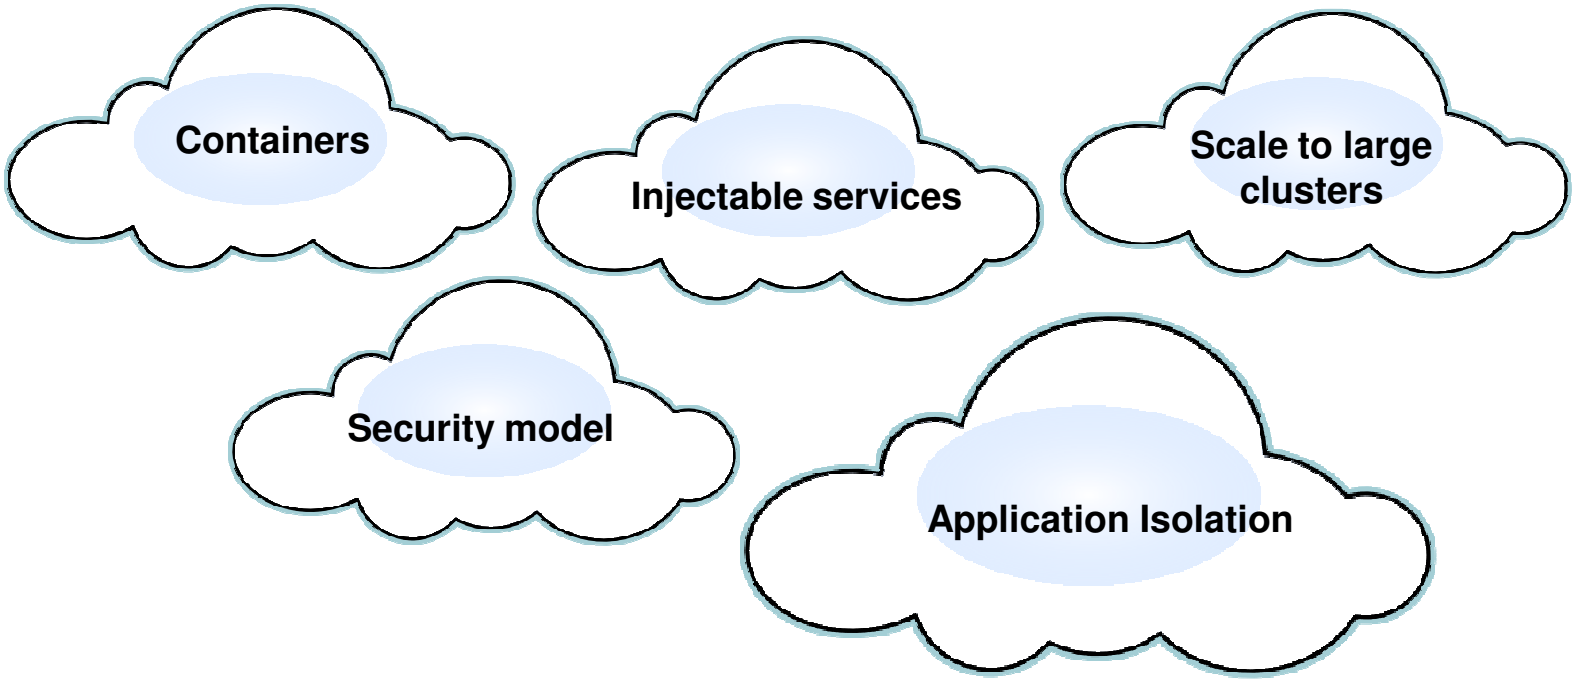

Java EE 7 : The theme - Cloud

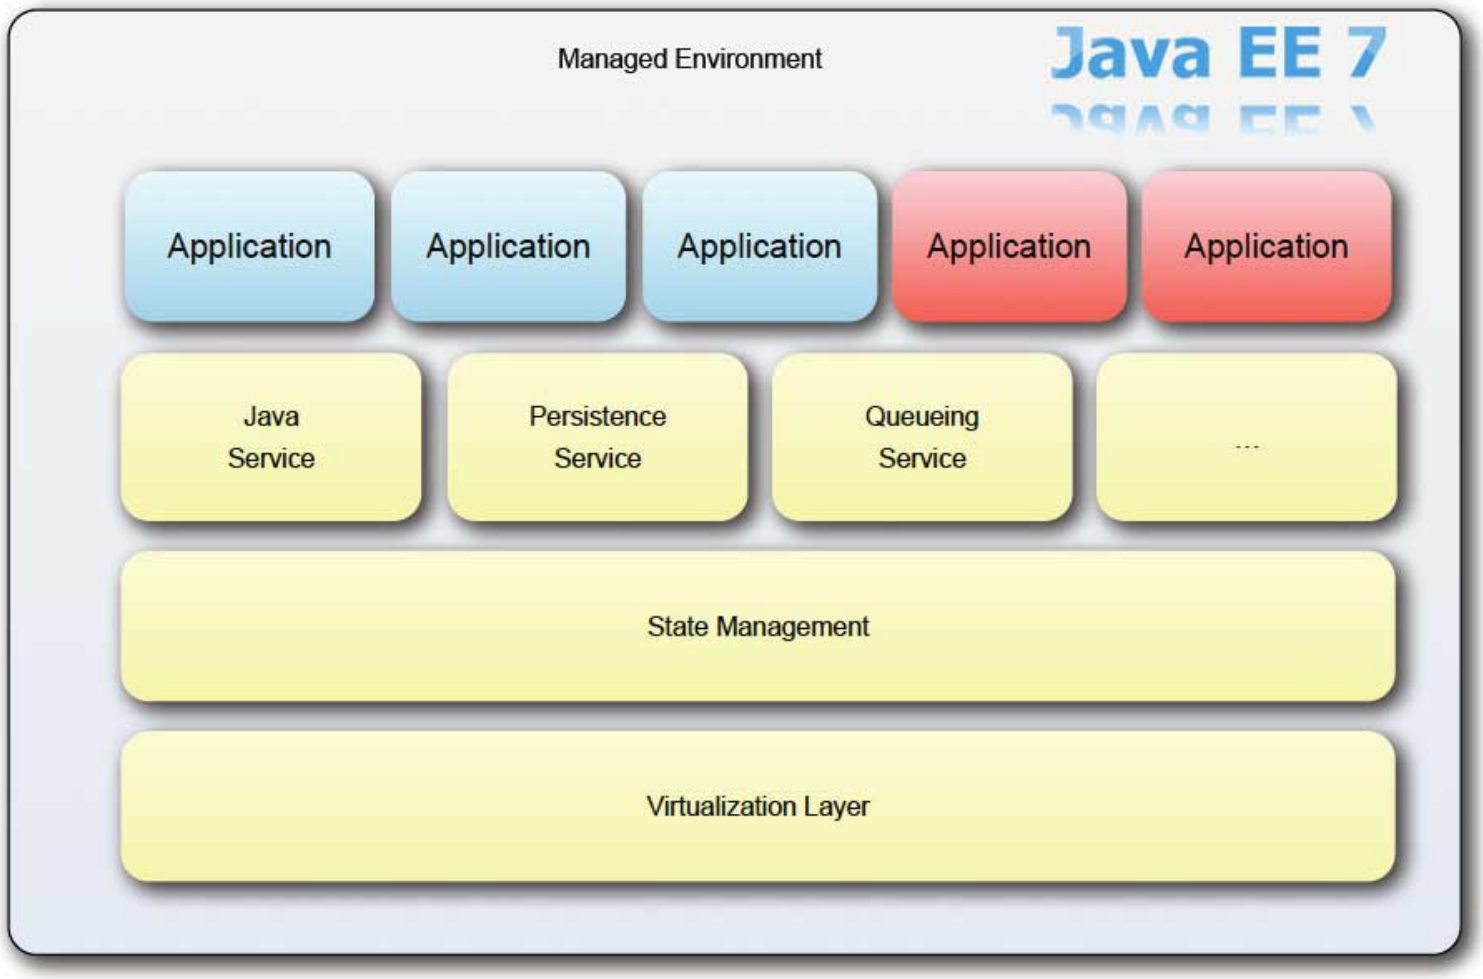

The Java EE 7 Modularity

- Built on Java SE 8 work
- Applications made of modules
- Dependencies are explicit
- Versioning is built-in
- Classloaders are straightened

Java EE 7 as a Cloud Consumer

- JAX-RS, the API for RESTful Web Services

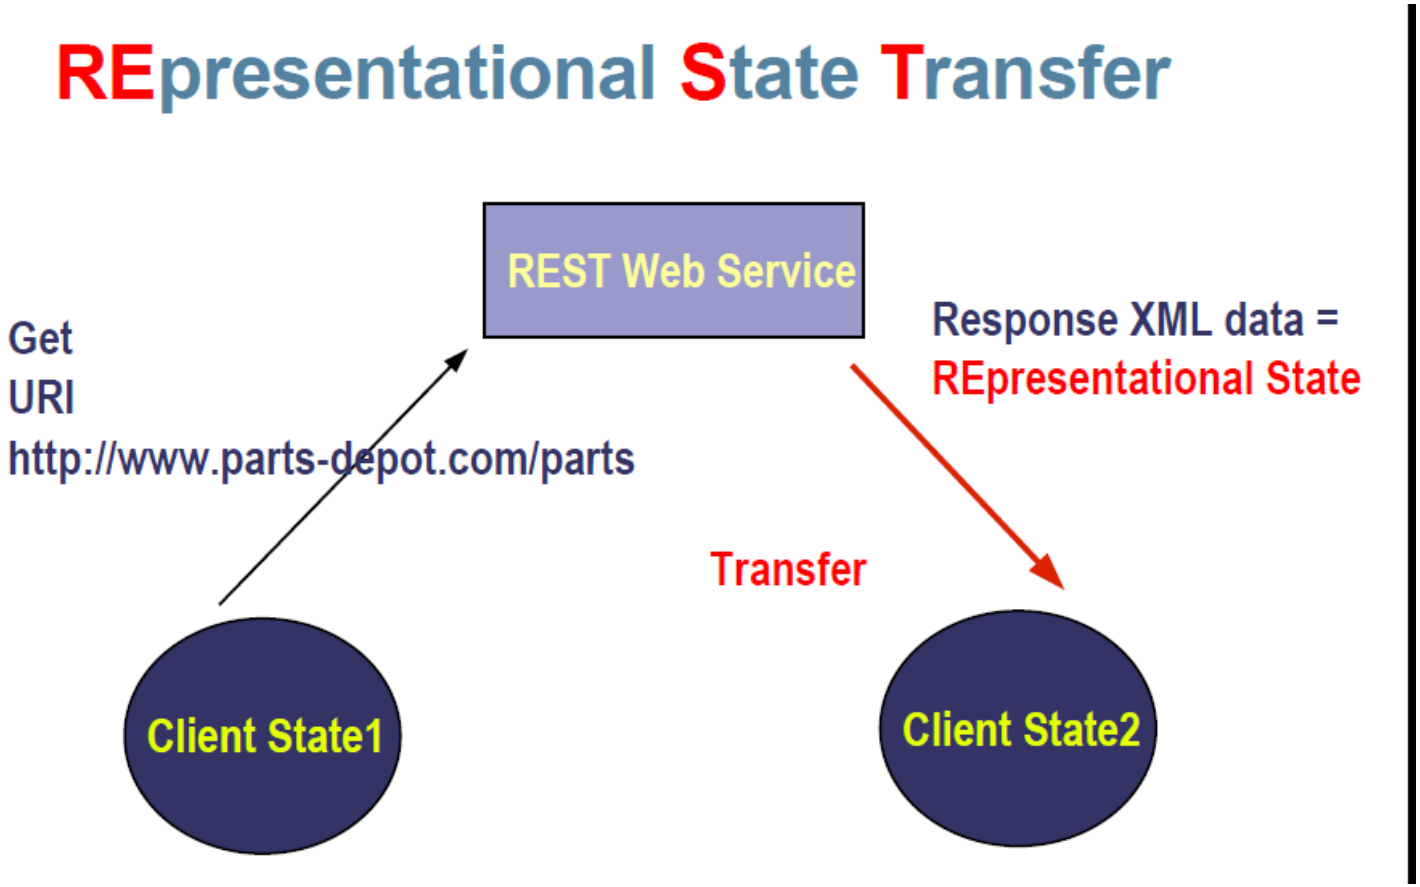

Java EE 7 as a Cloud Integrator

- JBI - Two kinds of components can be plugged into a JBI environment:
  - Service engines
  - Binding components

Java EE 7 as a Cloud Provider

- New roles
- All resource manager-related APIs, such as JPA, JDBC and JMS, will be updated to enable multi-tenancy
- Connectionless versions of the major APIs
- Descriptor for application metadata to enable developers to describe certain characteristics of their applications that are essential for the purpose of running them in a PaaS environment

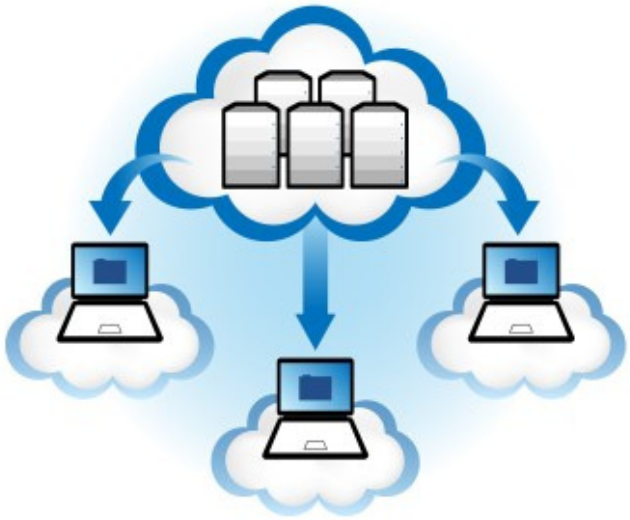

Need for a better Java Infrastructure

- **Fully optimized JVM stack** – which can efficiently leverage higher CPU cores densities, large memory page allocations (>50+ GB/sec) and optimized scheduler that reduces application thrashing
- **Scalable JVM heap sizes** – which can grow under application load through a cooperative memory manager and avoid performance degradation and "out of memory" errors crashes
- **Hit-less production visualization and performance management** – which supports fine gain I/O , lock contention and memory leak analysis that can improve problem resolutions times and identify application bottlenecks
- **SLA enforcement and policy based management** – which provides policy-based management that enables infrastructure-wide utility billing for compute resources, guaranteed transaction quality of service, and the ability to monitor security policies in multi-tenant isolated environment

A Plausible Solution: A Better, More Elastic Java Infrastructure

What JBoss + Azul provide:

- ✓ Fully certified Java "stack"
- ✓ Using only full tested components
- ✓ Integrated and optimized for scale
- ✓ Proven in mission-critical deployments
- ✓ Virtualization & Cloud ready

How Zing Works: AZUL JAVA VIRTUALIZATION

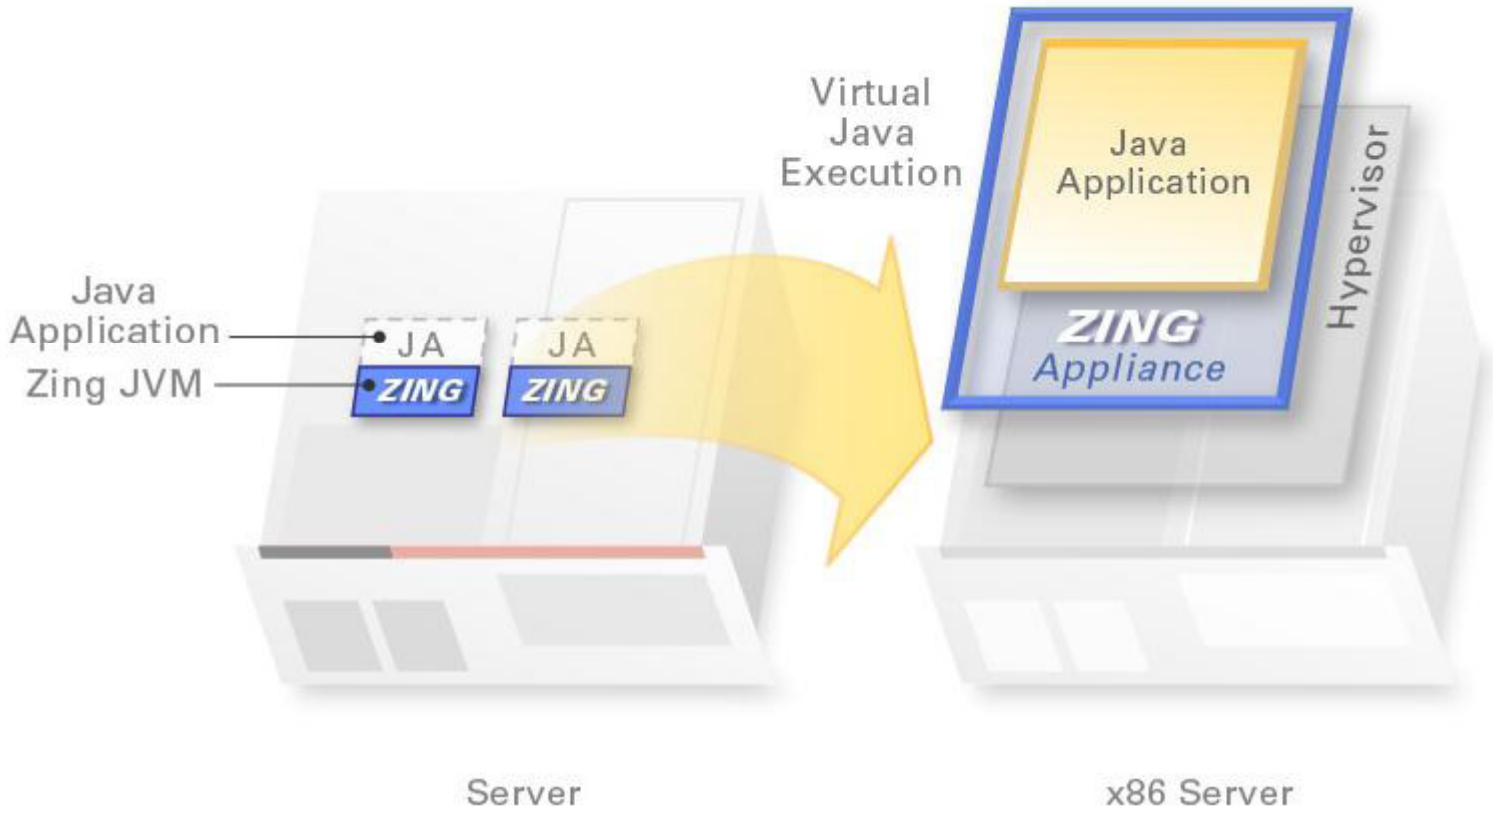

DataNucleus - DataNucleus - Windows Internet Explorer provided by Dell Client Engineering Team

http://www.datanucleus.org/

File Edit View Favorites Tools Help

Favorites | Web Slice Gallery

DataNucleus - DataNucleus

Page Safety Tools

DataNucleus

when you need it!

Search

HOME PRODUCTS DOWNLOAD SERVICES FORUM ISSUES WIKI BLOG

Home | Products | Services | Plugins | Plugin Points | Development

Quick Link

Overview

- **Welcome**
- Why Use DataNucleus?
- Download
- Publicity
- Used by
- News
- Events
- Performance
- Blog
- Release Schedule
- Team
- License
- References
- Donations
- Acknowledgements

Products

- AccessPlatform

Support

- Support
- Problem Reporting
- Contributing

Third-Party Tools

- Third Party Products
- Open Source

DataNucleus

-- DataNucleus AccessPlatform 3.0.2 released ... 27-Aug-2011 -- DataNucleus AccessPlatform 3.0.1 released

The DataNucleus project provides products for the management of application data in a Java environment. DataNucleus provides open source (Apache 2 licensed) products supporting data access using all standardised APIs (JDO, JPA) to a very wide-range of datastores (RDBMS, ODBMS, Map-based, Web-based, documents, etc) supporting querying using a range of query languages. **This flexibility of API and datastore is not available in any other persistence tool** Since the API's are standard and so reasonably well-known this reduces the learner curve, and you can use the same API for all of your data. **It's time to free up your data access layer**

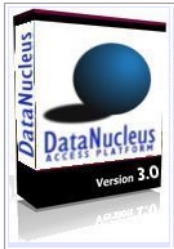

Access Platform is our baseline product. It is the most standards-compliant Open Source Java persistence product in existence. It is fully compliant with the JDO1, JDO2, JDO2.1, JDO2.2, JDO3, JPA1 and JPA2 Java standards. It also complies with the OGC Simple Feature Specification for persistence of geospatial Java types to RDBMS. It utilises an OSGi-based plugin mechanism meaning that it is extremely extensible. It is also free for use. The benefits for your business are significant.

Version 3.0 [Read More] | [Download] | [Get Started]

Version 2.2 [Read More] | [Download] | [Get Started]

http://www.datanucleus.org/services/index.html

Internet | Protected Mode: On

100%

Cloud Providers

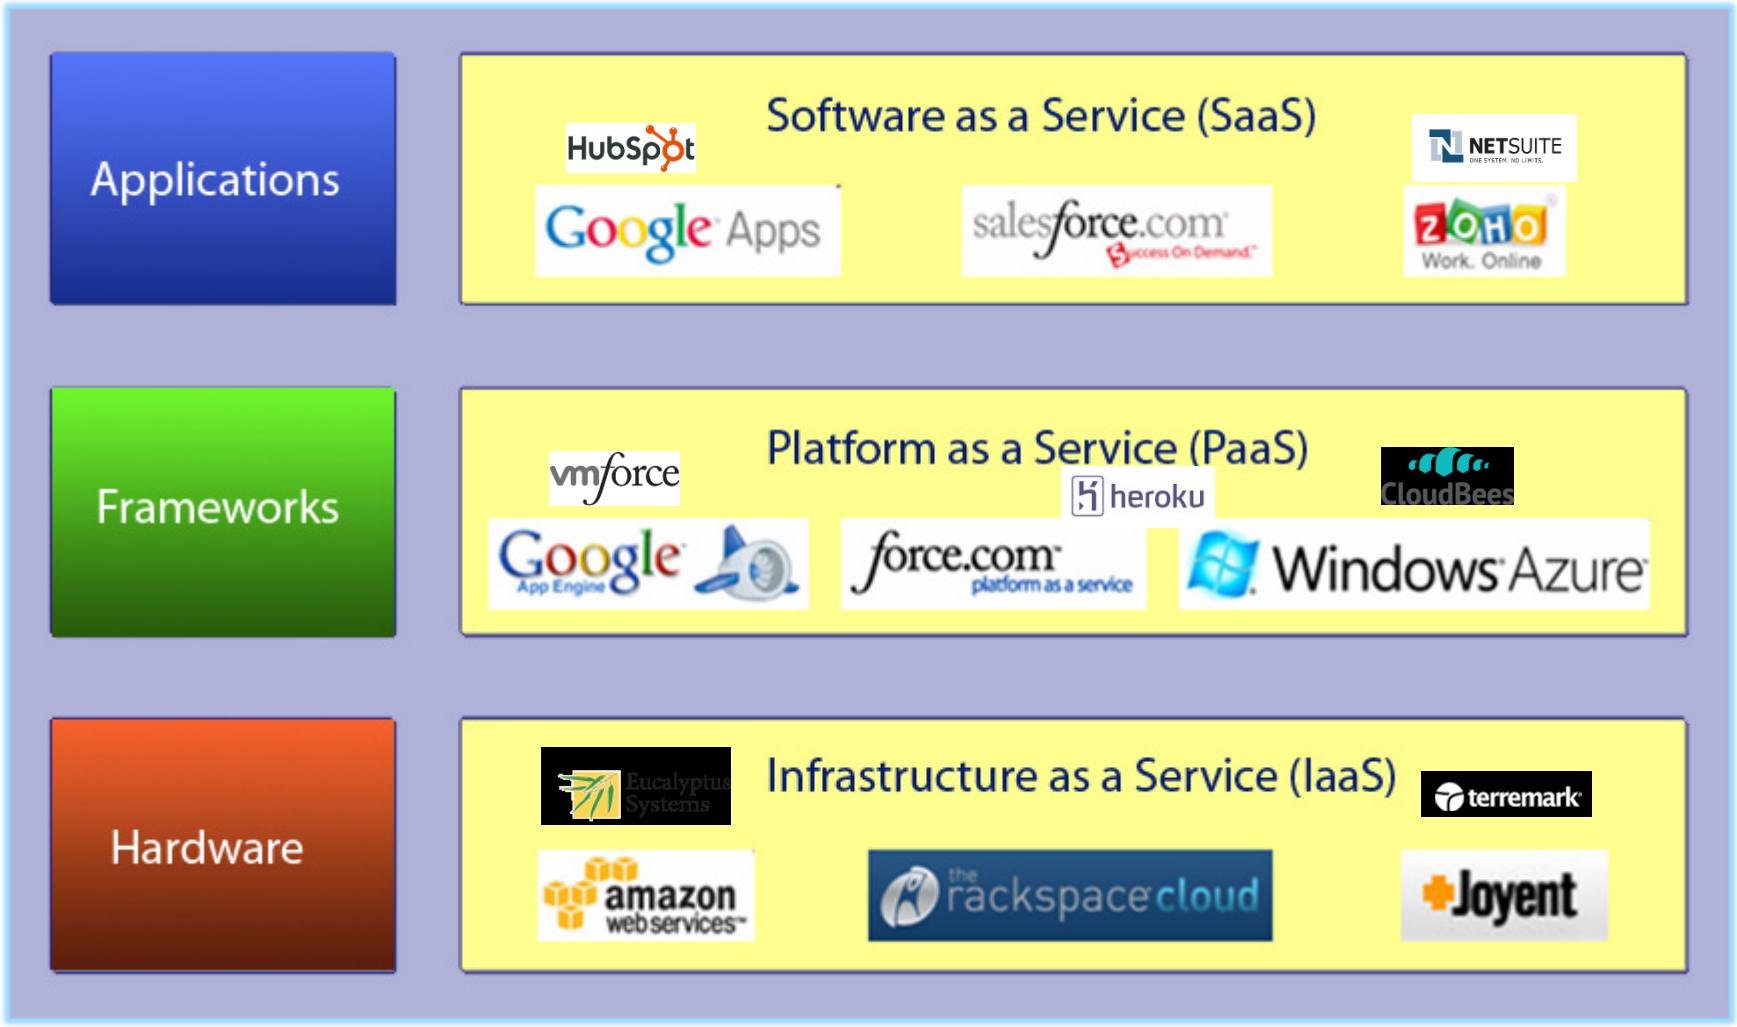

Amazon Beanstalk

- Based on the Amazon EC2 infrastructure
- ...and Auto Scaling and S3
- Add Linux, OpenJDK, Tomcat

- Eclipse Plug In available
- Supports version handling of applications
- Supports elastic scaling depending on load indicators
- Simple Monitoring built in
- Detailed control over the environment (Tomcat parameters, used AMIs, log in to machine etc.)3

Amazon BeanStalk

- Access to Tomcat logs etc.
- Access to the OS
- Fine tuning of Tomcat parameters possible
- Easy, yet powerful

- Videos to get started
- Demo application based on Spring
  - Uses also S3 (storage) and Simple Notification Service (SNS)

- Add Relational Database Service (RDS) for enterprise scale MySQL
- ...and all the other Amazon Web Services (AWS)

Amazon BeanStalk

- Much like your average Enterprise Java environment
- =Tomcat + RDBMS

- Cloud features like elastic scaling available
- Can easily add other AWS elements
- Runs on a proven environment

- But: 1 server = 1 virtual machine
- GAE can run multiple applications on one machine
- More cost efficient (?)

Google App Engine

- Infrastructure offered by Google
- Supports Java and Python
- Infrastructure completely hidden
- GAE sandbox more restrictive than normal JVM sandbox
- Limitations:
  - Focus on NoSQL while typical Java applications use RDBMS
  - Limit on start up time of application etc
  - Limit on response time (30 seconds)
  - no access to File system
  - No control or access to operating system
  - Not even the web server
- Java classes white list
  - i.e. some Java classes must not be used
  - no Threads
  - parts of System class (e.g. gc(), exit()...)

Google App Engine: Storage

- Relational database only in App Engine for Business
- Based on BigTable
  - Google's storage technology
  - High replication or simple master / slave
  - Key/value i.e. objects stored under some key
  - No joins
  - Simple / simplistic (?)
  - Scalable
  - Example of NoSQL

- Max. 1 MB per entity (and other limitations)

CloudBees

CloudBees: DEV@Cloud

- In fact two services: DEV@Cloud and RUN@Cloud
- DEV@Cloud: Developer services
- Continuous Integration (Jenkins)
  - Good application of the Cloud: Peaks and high load only during working hours
  - Standardized and universally applicable service
  - Some Essentials Plug Ins in free version
  - More in Base / Pro / Enterprise pay version
  - Also more parallel build in pay version
  - ...and faster build machines
- Maven repository
  - Snapshot / Release
  - Builds can be automatically deployed
- Potentially other services
- Runs on AWS

CloudBees DEV@cloud

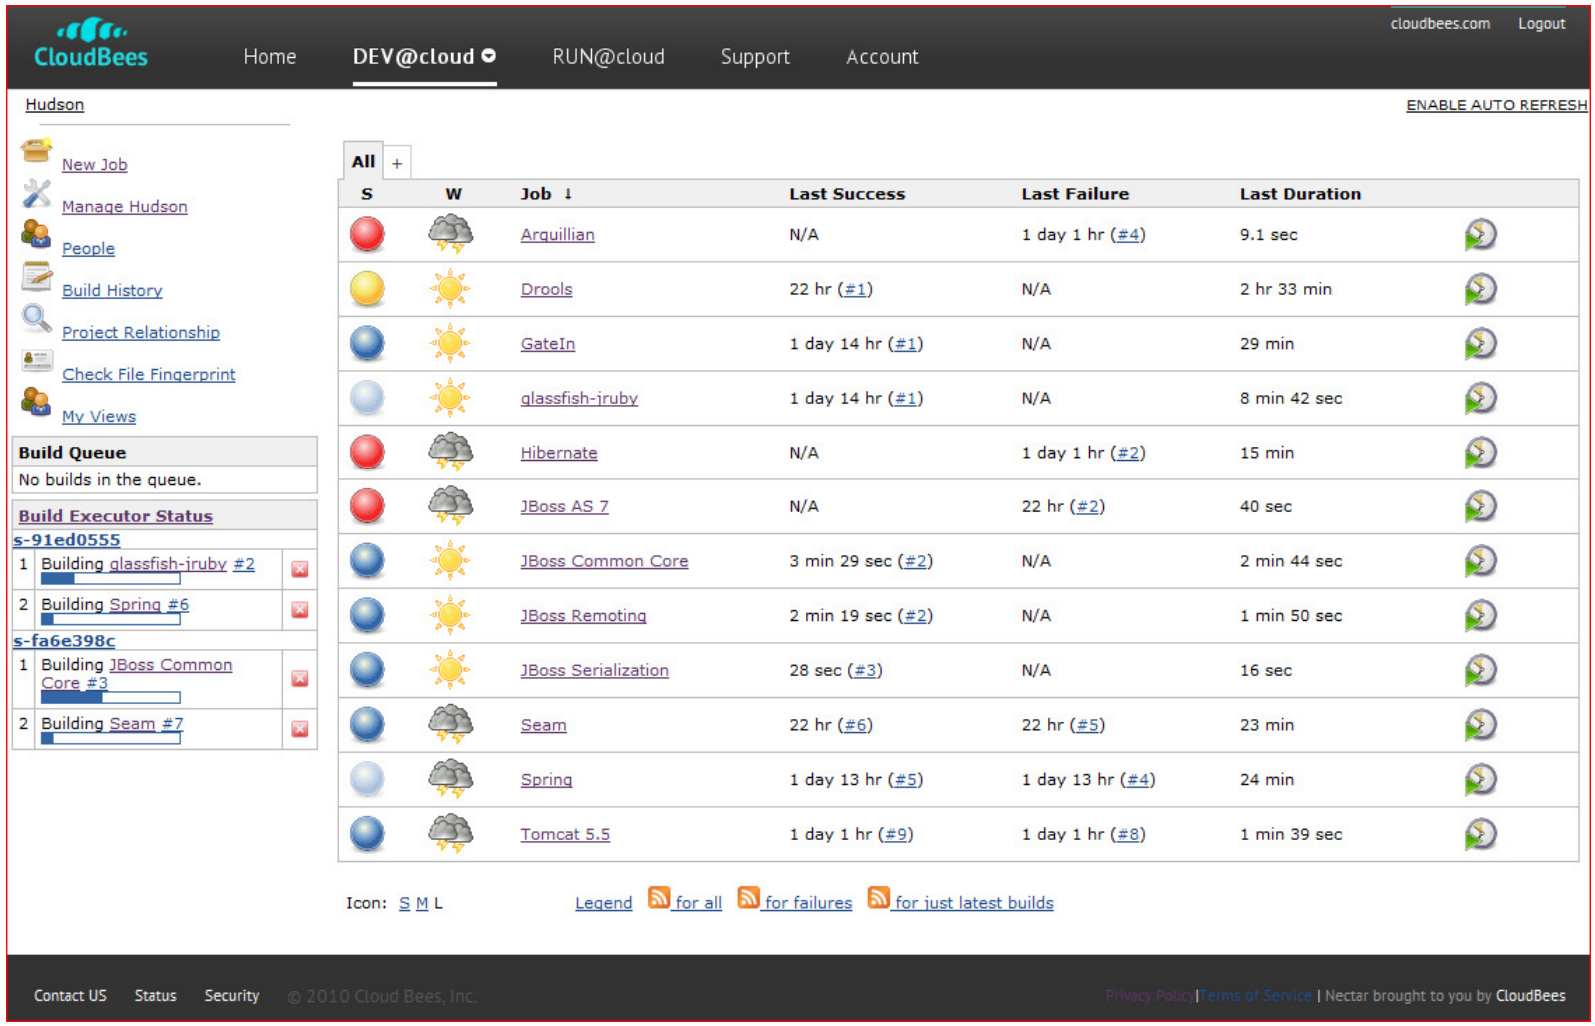

CloudBees: RUN@Cloud

- MySQL database
  - Very simple (i.e. just one server, but backup included)
  - Could use Amazon RDS instead
  - Master/Slave setups
  - Load balanced JDBC driver
  - Automated backups

CloudBees: RUN@Cloud

- Tomcat on EC2
- Easily deploy a WAR
  - either by web interface
  - or command line utility (bees SDK)
- Java as a Service
  - HTTP load-balanced clustering
  - Shared or dedicated server deployments
  - Local development sandbox
- Little control
  - Only 256MB heap
  - Single instance or multiple instance
  - No elastic scaling
  - Potentially cheaper: Multiple applications on one machine
- Simple monitoring (web / command line)

Java in the Cloud: Conclusion

- Feasible to run Java applications in the cloud
- IaaS
  - Maximum control
  - Automatic installation needed
  - Also works on private cloud / virtualized environment
- Google App Engine
  - Very limited sandbox
  - Advantage?
- Amazon Beanstalk
  - Standard Enterprise Java Stack
  - Lots of additional Amazon Web Services (Relational Database Service)
- CloudBees
  - DEV@Cloud: Developer-only features (Jenkins)
  - RUN@Cloud: simple but probably more price efficient

Dell on the Cloud

- Just a month ago, Dell and VMware jointly announced a partnership for creating Dell's first Cloud based on Dell PowerEdge C-Series servers, used in other data centers including Windows Azure, and VMware vCloud Datacenter Services providing the infrastructure needed for public, private and hybrid clouds and the availability of consulting, application and infrastructure services. Dell intends to build private and hybrid clouds in their own or the customer's datacenter using VMware vSphere and vCloud Director.

VMWare – vCloud Director

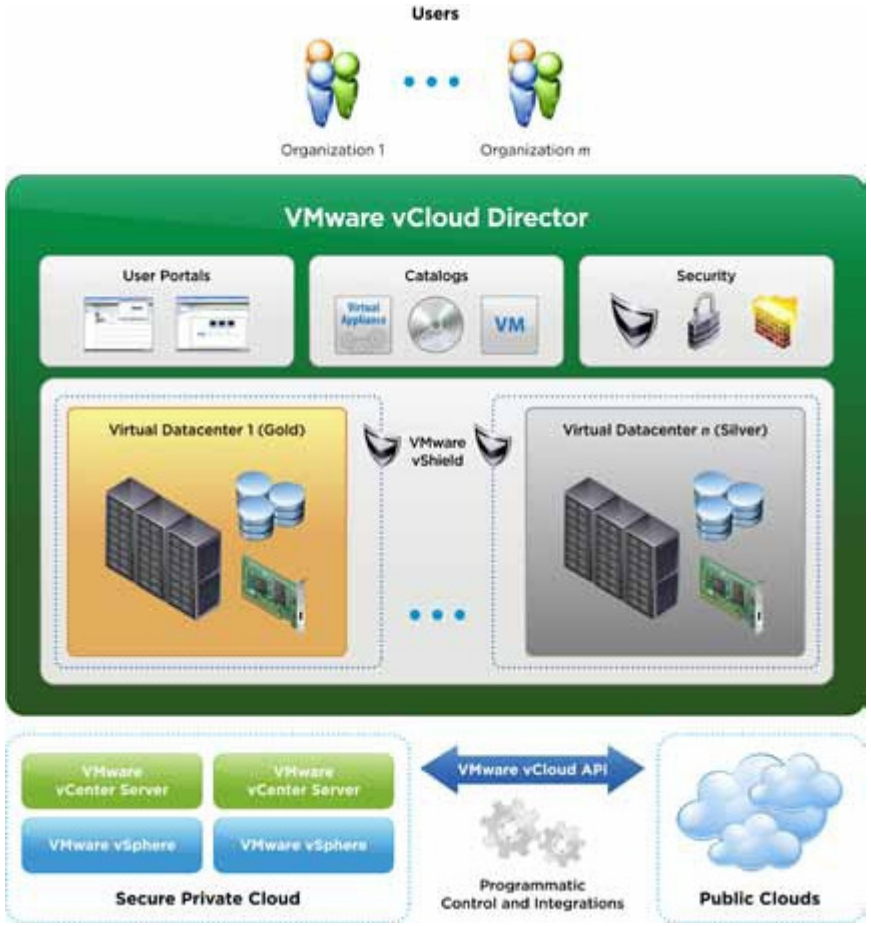

Dell BOOMI

AtomSphere® Integration Cloud™

What is Dell Boomi AtomSphere?

- Make it possible for companies to integrate application, data, and trading partners directly from web by using a visual designer with access to a library components,
- Build, Deploy and Manage Connections Directly from the Web
- No Software Packages or Hardware Appliances to Install
- Connect any Combination of SaaS and On Premise Applications with Unprecedented Ease
- Pay Only for the Connections you Deploy
- Do-It-Yourself Technology - No Coding required
- Self-Provisioning - Sign Up and Begin Building Integrations Immediately

References

- http://www.oracle.com/us/technologies/cloud/036500.pdf
- http://blog.cloudbees.com/2011/03/what-is-PaaS-after-all.html
- http://www.ibm.com/developerworks/java/library/j-PaaSshootout/?ca=drs-
- http://www.cloudcomputingdevelopment.net/running-java-in-the-cloud-with-cloudbees/
- http://broadcast.oreilly.com/2009/04/java-for-google-appengine-fina.html
- http://www.azulsystems.com/
- http://www.vmware.com/