




Why Private Clouds Are Likely to be Open Source

Rich Wolski
CTO
Eucalyptus Systems Inc.
www.eucalyptus.com



Building Your Big Data Future with Open Source



SLAs



Web Services



Virtualization





eCommerce and Infrastructure

- Self-service and “zero touch.”
 - Scalable automatic rental of resource intensive goods
- Transactional and asynchronous
 - Interaction with the site is transactional
 - Delivery is asynchronous
- Site integrity and site availability are critical
 - Individual transactions can fail but the site cannot
- Customer requests must be isolated
 - Service venue must manage competing needs
- Scale out for request volume, scale up for request weight
 - Customer requests are independent
 - Individual requests carry different requirements





eCommerce and Open Source

- eCommerce technologies are mature
 - Half-way through its second decade
- Some of the most prevalent technologies are open source
 - Apache, Tomcat
 - Java
 - J2EE
 - Apache Geronimo, JBOSS, OpenEJB, GlassFish, JOnAS, Spring
 - MySQL, memcached
 - Linux





Linux and Enterprise IT

- Linux is the operating system platform of choice
 - Hardware portable
 - Separates software lifecycle from hardware lifecycle
 - Prevents lock-in
 - Vast ecosystem of software
 - Linux distros provide QA (free or paid)
 - Transparent
 - Possible to own the source code for everything
 - Fast to remediate
 - Open source web community is often faster than paid support
 - Cost effective
 - Possible to mix free and paid offerings fluidly





Enterprise Cloud = Open Source

- If...
 - the most mature eCommerce technologies are open source
- And...
 - Enterprise IT prefers open source platforms for deployment at scale
- And...
 - Private Clouds are the next platform for data centers
- Then...
 - The On-premise Private Cloud will be built from Open Source
- *Will Enterprise IT run a private cloud?*





Private Cloud Advantages

- Separate the application software lifecycle from the infrastructure software lifecycle
 - Application “stacks” can include operating environment
 - Legacy applications can be transitioned gradually
- Separate application software from infrastructure configuration
 - Clouds support a unified configuration model
 - IT can reconfigure in response to technological innovation without breaking application execution
- Customized policy implementation
 - Private clouds must respect local policy definitions
 - One size does not need to fit all





More Private Cloud Advantages

- Flexible resource apportionment
 - The application “mix” can vary based on user needs or load and not infrastructure constraints
- Self-service, scalable, and automatic
 - Resource allocation and decommissioning does not require an infrastructure team and/or a personnel driven process
 - Movie rental analogy: Netflix instead of Blockbuster





Private Cloud Disadvantages

- Cloud application porting is not automatic
 - Can be incremental but requires developer time
 - Elasticity driven by unified configuration is the issue
- Cloud Heisenberg Uncertainty
 - It is not possible to know simultaneously where something happens in a cloud and when it happens.
- Enterprise IT process change
 - Infrastructure administration and management roles change due to the separation of concerns
 - “Application infrastructure” versus “cloud infrastructure”





Eucalyptus

- Elastic Utility Computing Architecture Linking Your Programs To Useful Systems
- Open source web services based implementation of elastic/utility/cloud computing infrastructure
 - Born of an NSF-funded research project in 2008 at UC Santa Barbara
- Software overlay
 - integrates multiple hardware and software platforms
 - Configuration agnostic
 - Policy agnostic
- Built from “best of breed” open source eCommerce technologies





How do we know it is a Cloud?

- The name is power
 - In 2008 the academic and industry communities hotly contested any definition of cloud
- Turing Test:
 - Pick something that is undisputably a cloud and mimic it to the point where an adversary can't tell the difference
 - Amazon AWS
- *Can we emulate Amazon AWS using open source components and commodity hardware so that a user can't detect the difference?*
 - Amazon "the home version" generated a great deal of interest as freely available open source





Eucalyptus as a Business

- Today's Typical Customer:
 - “I need a private cloud, I want it to be open source, and I can't change my infrastructure.”
- I need a private cloud
 - Hype cycle is driving many IT organizations to consider private cloud
 - Web 2.0 intensive companies are early adopters
 - Experienced in scalable eCommerce tech.
 - Amazon AWS is part of the adoption cycle
- I want it to be open source
 - Open source == cheap, transparent, avoids lock-in, etc.
- I can't change my infrastructure
 - Current hardware and software infrastructure must be amortized






Monetizing Eucalyptus

- Open source subscription model
 - Annual fee provides
 - Support
 - Customization expertise
 - Management expertise
 - Use of a small set of proprietary “plug ins”
 - Portability, performance, reliability but not functionality
 - Automatic upgrades
- Aggressive release cycle
 - Goal: new code and features every quarter
- Professional services
 - Training, custom development, etc.





Building Your Big Data Future with Open Source





What Customers Buy

- Configuration and management expertise
 - There are an infinite number of ways to match Eucalyptus to existing infrastructure
- Access to supported upgrades
 - Cloud upgrade is a scary operation
 - New features and release cycle are attractive
- Hybrid usage
 - The ability to move applications and data between private and public clouds makes Eucalyptus support useful for AWS
- Ecosystem
 - Eucalyptus has more than 100 supported partners
 - Anti-lock in, pro choice





Ghosts of Platforms Past

- Then (circa 2000)...
 - AIX, Coherent, Interactive Unix, HP-UX, Irix, NeXTSTEP, SCO, SINIX, SunOS/Solaris, Ultrix
- And now...
 - Linux
- The difference is ecosystem
 - Proprietary Unix still exists but it is highly specialized
- The Platform is Open Source but the ecosystem is mixed
 - Most vendors port their proprietary Unix applications to Linux
 - A few of them just went open source with their platform






Eucalyptus Ecosystem



Four Instance Start-up Time (CDF⁻¹)





Challenges

- Early market
 - Most monetizable combination of features, configuration, and ecosystem isn't yet visible
- Open source QA
 - The space of open source options is bigger than the space of options we sell typically
 - Gets bigger every day
- Big competition
 - VMWare is arguably the most successful software company of the last decade





Winning with Open Source

- Industry standard APIs
 - Technologist driven standards efforts are always more tactical than practical
- Enterprise Open Source governance
 - Open source software successful in the enterprise almost always has an enterprise managing governance
 - Mitigates the risk caused by a lack of focus (software can do anything)
- Packaged and ubiquitous deployment
 - Customized deployments are useful, but static
 - The usage experience must be repeatable and zero touch
- Reciprocal licensing
 - GPL-style licensing prevents “bait and switch”





Which One is Yours?





Building Your Big Data Future with Open Source

Thanks!

Questions?

