

Scalable SOA in Your Own Open Source Cloud



Sanjiva Weerawarana, Ph.D.

Founder, Chairman & CEO, WSO2
Founder, Director & Chief Scientist, Lanka Software Foundation
Member, Apache Software Foundation
Emeritus Board Member, Open Source Initiative
Visiting Lecturer, Univ. of Moratuwa, Sri Lanka

April 2-3, 2009, Arlington, VA.

Overview

- “Cloud computing” - my version!
- Clouds & SOA
- Open {source, standards, services} for clouds
- Avoiding the risks

Cloud computing

- Essentially, the technology to make computing a utility is now widely available
 - “Plug in” and use what you need, when you need it, where you need it
- Brings together
 - Virtualization, Grid computing, Utility computing, Network computing
- Ranges from low level to high level
 - Low level: platform level services for acquiring computation, storage & communication
 - Amazon Web Services (EC2, S3, etc.)
 - High level: fully managed, fixed programming model
 - Google AppEngine
- Public vs. private clouds

Key characteristics & benefits

(Not necessarily what vendors provide today, but this is what makes cloud computing bring real value for SOA!)

- Auto-scaling (elasticity)
- Multi-tenancy
- Scalable parallel computing
- Other benefits include
 - Cost savings by not peak provisioning
 - Faster deployment
 - Outsource non-core skills
 - Utility billing

Auto-scaling (elasticity)

- Ability for the resources to be scaled up, *and down*, based on demand
- Enables efficient use of compute resources
 - No need to provision for maximum load
- Requires application or middleware to do work to manage load
 - Queue up incoming requests
 - Distribute amongst available worker nodes
 - Monitor queue length and adjust worker pool accordingly

Multi-tenancy

- Ability for a single application to be used by multiple customers simultaneously
- Enables scaling of customers without having to assign a server cluster per customer
- Critical is data isolation
 - Separate databases to tables to multi-tenant tables
- Application authors must be able to write multi-tenant applications
 - Middleware can provide tools to help

Scalable parallel computing

- Ability for the elastic group of nodes to work together as a parallel computer
- Enables application to engage as many nodes as needed for the problem at hand
- Requires programming model
 - MPI-style low level to Map-Reduce/Dyad style
- Requires large scale data management solution
 - Distributed file systems to DHTs to group communication

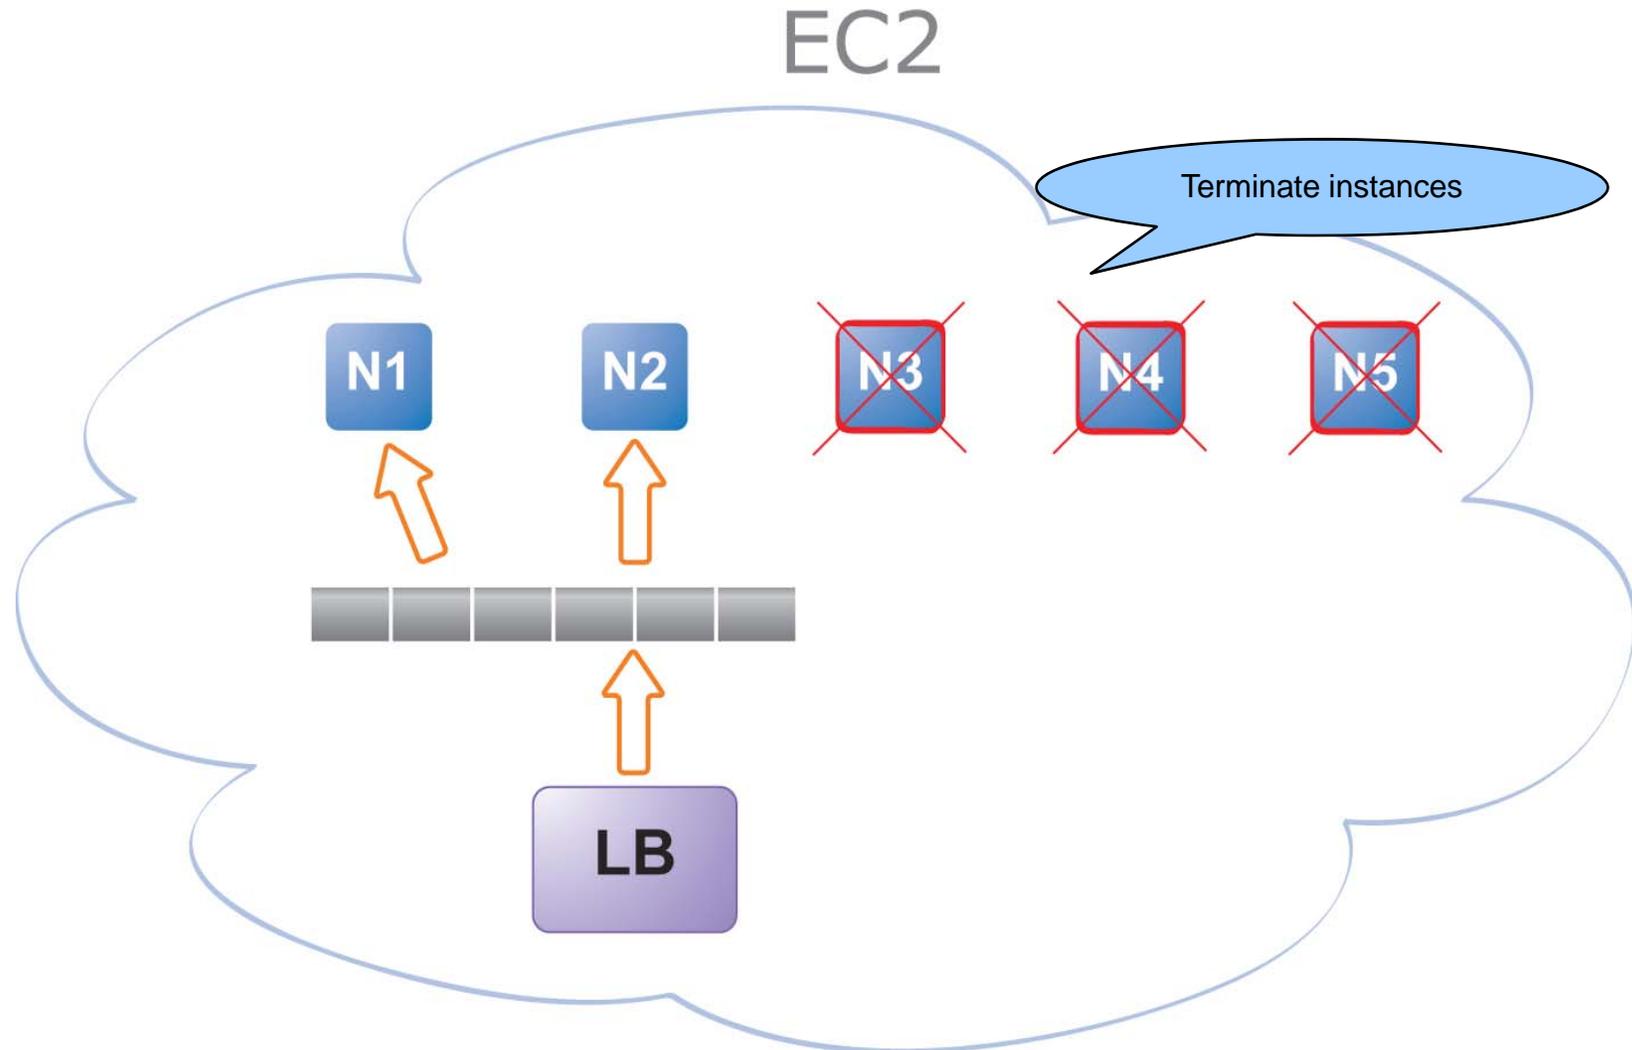
Role of clouds in SOA

- Summary: a cloud is a cost-effective deployment platform for your SOA
- Basic level
 - Use virtualized platform as a convenient provisioning tool
- Advanced level
 - Service and application authors write services which auto-scale, are multi-tenant and possibly use parallel computing

Virtualized deployment in a cloud

- Low level is easy
 - Make VM image with your software and boot it up
 - Most “cloud” products do just that!
- Not any different from just a regular deployment
 - Still advantageous – no need to go thru deployment/configuration pain!
- Any product can be deployed this way
 - No need for product to be modified (as long as OS is supported on the virtualization platform)

Auto-scaling services: Apache Axis2 / WSO2 WSAS approach



Writing multi-tenant services

- Objective: allow different customers / users to use the same service as if it were their own
 - Having the service be multi-tenant vs. running a different server for each gives all the sharing advantages
- Issues
 - Authentication & authorization
 - Data
 - From database per tenant to multi-tenant tables
 - Security / protection
 - Server architecture for scalability
 - Programming model for service authors

Scalable parallel computing

- Service implementation may need to use variable amount of compute resources
- Can bring in parallel computing infrastructure
 - E.g., map-reduce

Open source / standards / services for clouds

- Risk of *hard* vendor lock-in with clouds!
- Services in the cloud are a new lock-in risk
- Open cloud manifesto
 - (IMO) too much marketing, too little real value
 - Really, too early to attempt to standardize and define what “open” means

Open source cloud platforms

- Open source alternatives for all aspects of cloud computing exist and are evolving
- Join the communities or support them commercially
- Set up your own cloud, deploy scalable middleware and governance
 - “Private cloud” is a good thing, especially for large organizations

Summary

- “Cloud computing” provides a cost-efficient, convenient deployment architecture for SOA
- Vendor lock-in risks are higher with cloud computing
- Open source is driving innovation in cloud computing and provides great way to avoid lock in
- Still very early days .. will take time to mature & stabilize!



sanjiva@wso2.com

<http://sanjiva.weerawarana.org/>