

HEY YOU, HAVE YOU MOVED TO THE CLOUD?

An introduction to the Cloud: why and how

JARED STILL – PRINCIPAL CONSULTANT

About Jared Still

- Principal Consultant at Pythian
- Oak Table Member
- Oracle Ace
- Perl aficionado
- Twitter: @PerlDBA
- jkstill@gmail.com

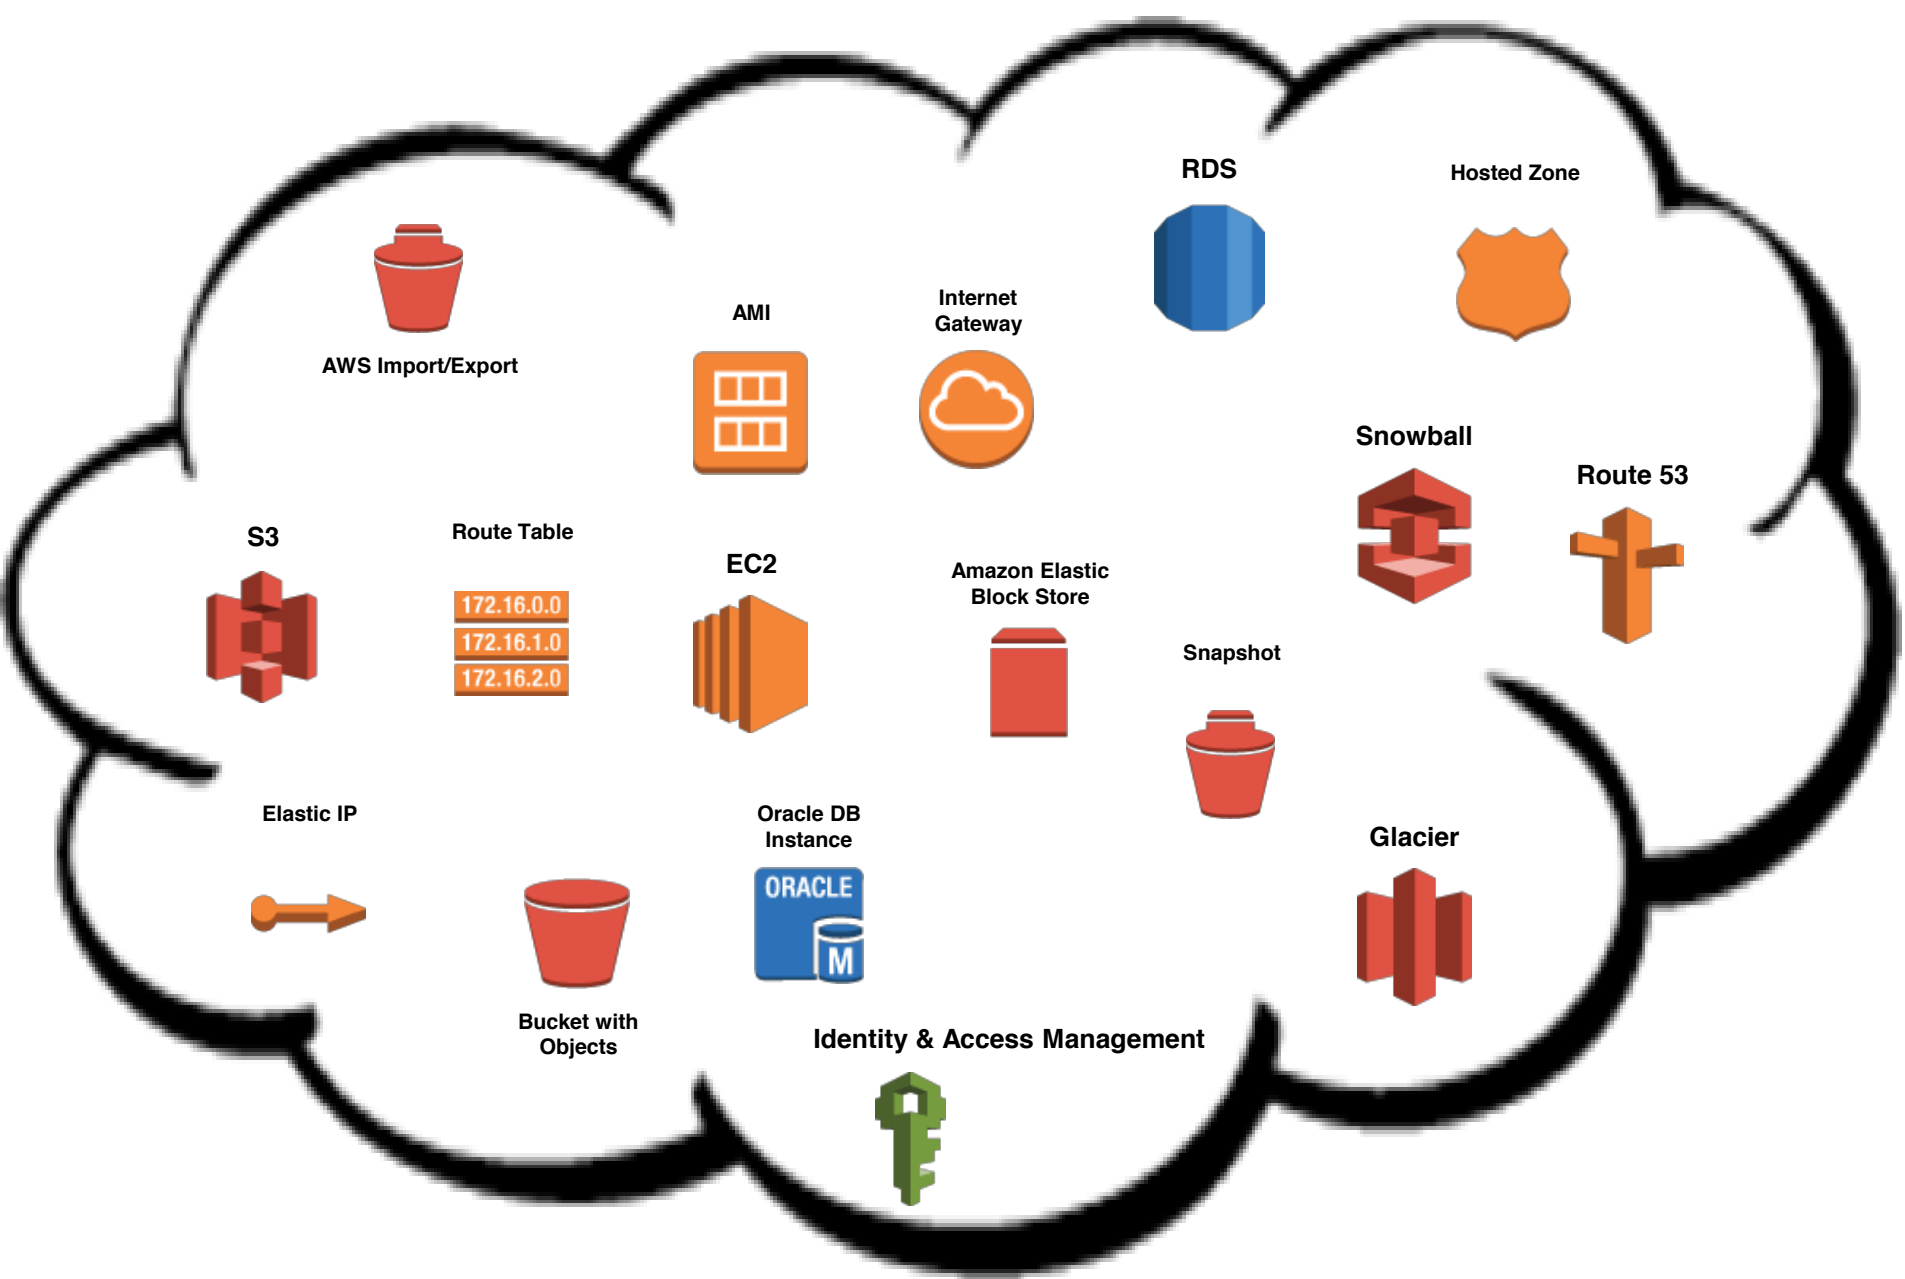


WHAT IS AWS?

A Brief Primer

AWS = AMAZON WEB SERVICES

- Compute Resources
 - EC2
- Storage Resources
 - EBS
 - S3
- Network Resources
 - Routing
 - DNS
 - Gateways





EC2: ELASTIC COMPUTE CLOUD

- Computing ‘instances’
 - Linux
 - Windows
- Scalable
- Scriptable
 - Instantiate a computer via CLI
 - `aws ec2 run-instances --image-id ami-f0091d91 --count 1 --instance-type t2.micro --key-name my-key-pair --security-group-ids sg-945663ae --subnet-id subnet-17e6fb92`



EC2 INSTANCE TYPES

- Hardware Attributes
 - Number of vCPUs
 - Memory
- Storage Type
 - Local (ephemeral or non-persistent)
 - EBS (persistent)
- Optimizations
 - Compute
 - Network
 - IO



AMI: AMAZON MACHINE IMAGE

- Images used to launch EC2 Instance
- Linux
 - Several Flavors
- Windows

EBS: ELASTIC BLOCK STORE

Amazon Elastic
Block Store



- SSD backed storage
- Persistent
- Snapshots

S3: SIMPLE STORAGE SERVICE



- Organized in Buckets
- Highly Available
 - 99.9999999999% Durable
 - Lose one object every 10^{10} years.
 - 100 bucket limit
 - Unlimited objects per bucket
- Use for:
 - Backups
 - Static Web pages
 - Personal file backups

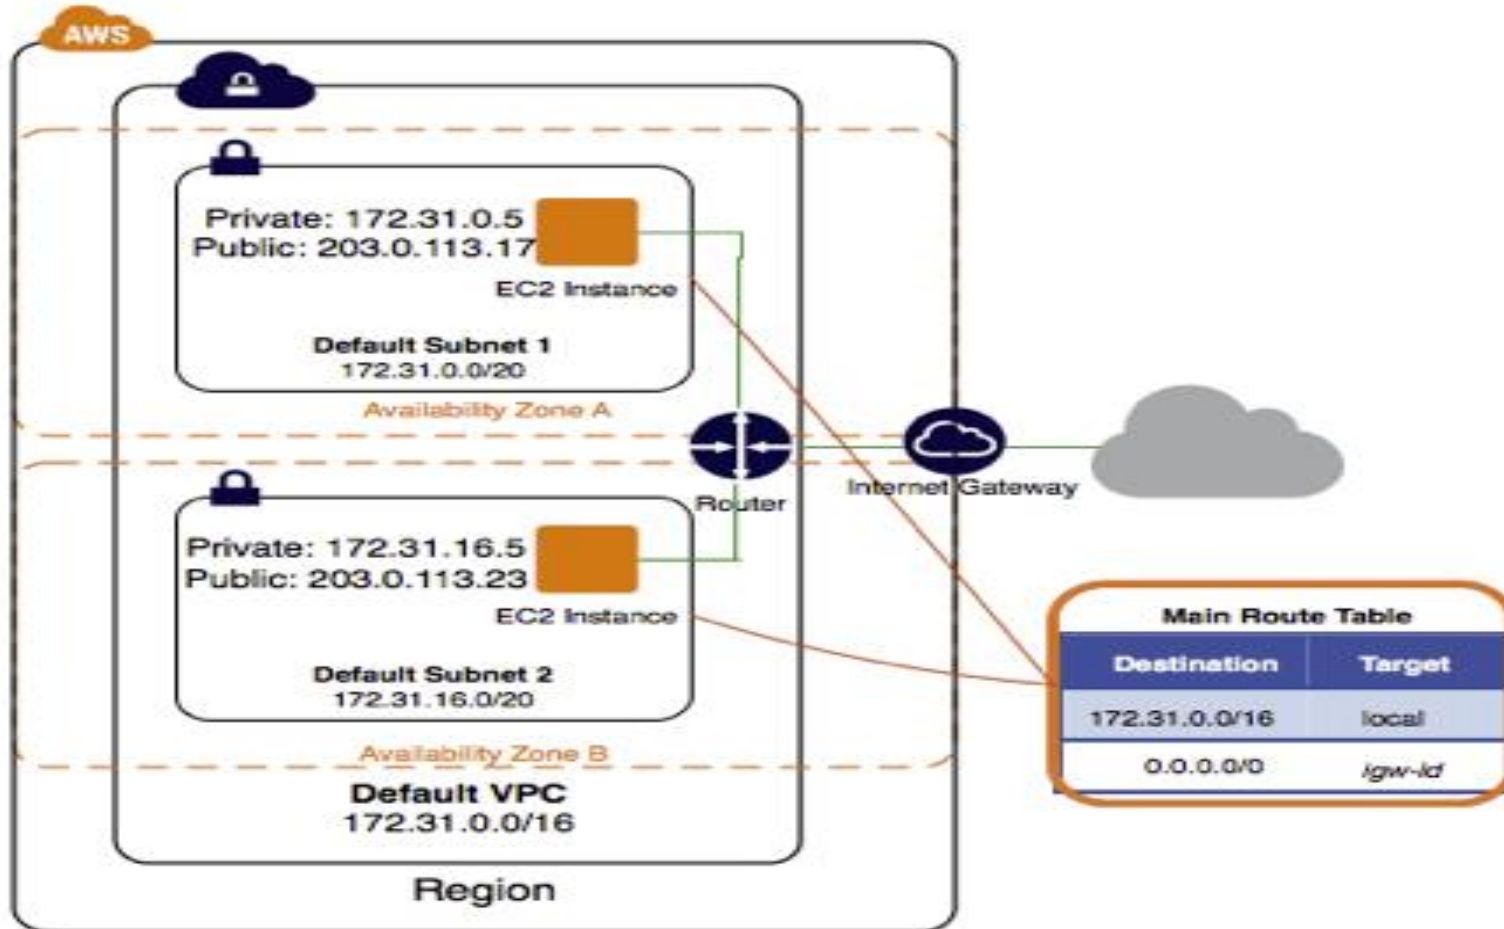
Bucket



VPC: VIRTUAL PRIVATE CLOUD



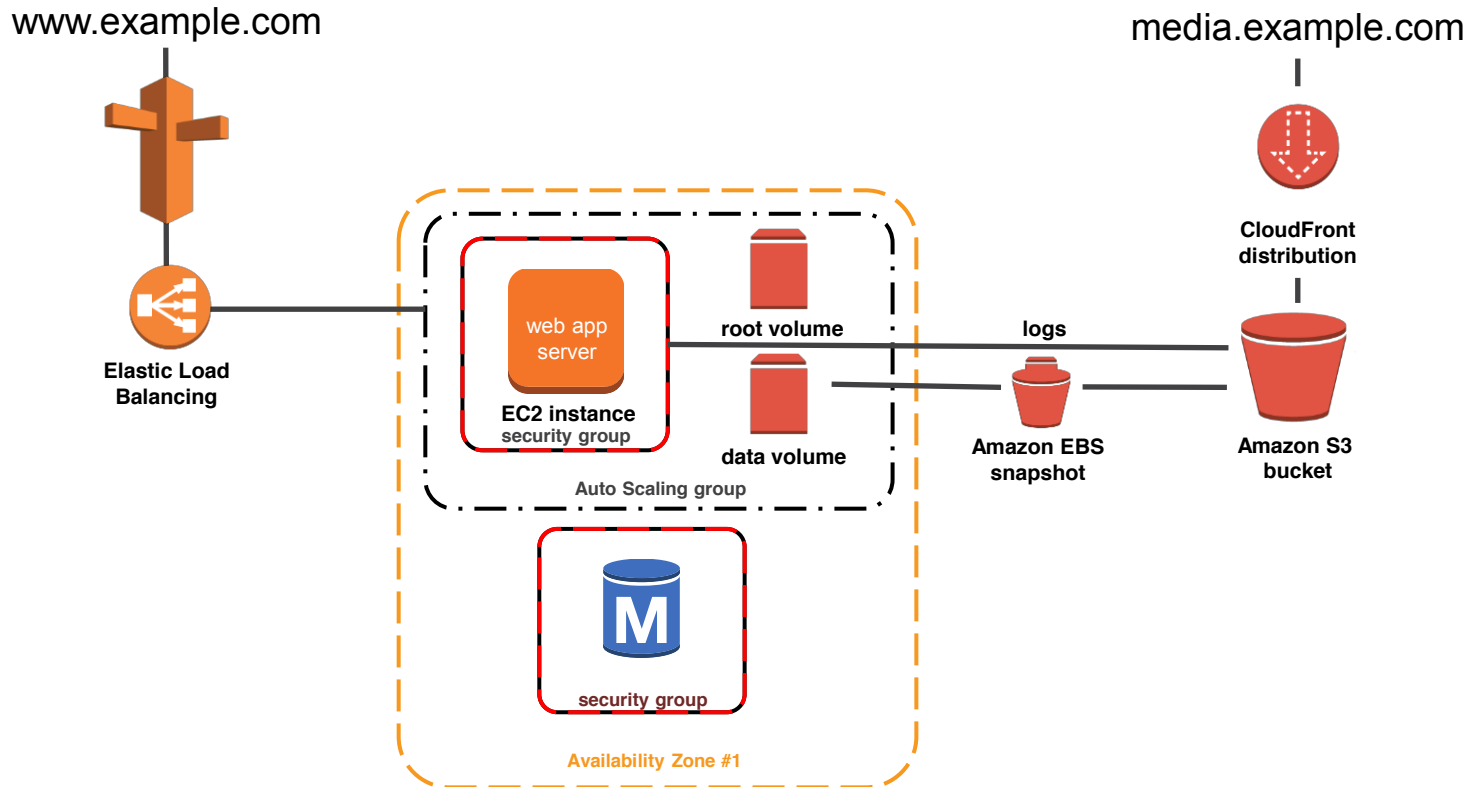
- Network Infrastructure





VPC: VIRTUAL PRIVATE CLOUD

- Network Infrastructure





RDS: RELATIONAL DB SERVICE

- Amazon Controlled RDBMS
 - Oracle
 - SQL Server
 - MySQL
 - MariaDB
 - Aurora (Amazon's optimized MySQL)
 - PostgreSQL
- Many DBA Tasks Automated
 - Backups
 - Patching
 - More on this later...

COMMON TRAITS

- Scalable
- Scriptable
 - “Infrastructure as code”
 - Chef, Puppet, Vagrant, OpsWorks, Cloud Formation
 - DBA’s may not create these, but...
 - DBA’s should be aware of them
- Highly Available
 - HA does require some design work

WHY MOVE TO THE CLOUD?

FOLLOW THE MONEY

- CAPEX vs OPEX
- Warning!
 - I am not an accountant
 - Nor do I play one on TV
- This presentation will contain only some simple comparisons.

REDUCE CAPEX AND OPEX

- CAPEX – Capital Expenditure
 - ie. Hardware in this case
- Overprovisioning for seasonal business
 - 40 cores
 - 1T RAM
 - Nearly idle at 10% load 10.5 months/year
- Dell R930
 - 4 x Xeon® E7-4830 (10 core each)
 - 1T ECC RAM
 - \$32,397.22
 - ~ \$39k with support

AWS = EXPENSIVE?

- AWS Equivalent (approximate)
 - 2 x m4.10xlarge (40 vCPU each)
 - Pay 3 yrs up front (includes support)
 - Total RAM 320G – 32% of dedicated server
 - ~~\$64,000~~ \$46,936.00 total
- Not much savings there!
 - But...
- Paying for more than server
 - Infrastructure

SCALABLE SAVINGS

- Think about scaling
 - $80 * 10\%$ load = 8 threads required 87.5% of the year
 - Single m4.2xlarge (32G RAM)
 - \$4694.00 3 yrs
 - Scale out 2 months per year
 - 2 x m4.10xlarge \$3456.00 monthly
 - Over 3 years \$20,736.00
- Total Server cost \$26k

HAVING A GOOD YEAR?

- Sales exceeding all expectations
- Legacy method
 - Try to keep things working
 - Systems will be slow...
 - Plan for next year – too late for this year.
- Scalable method
 - Add more servers.
 - Point, click, done.

AUTOSCALING

- Architect for auto-scaling
 - Monitor load
 - Add EC2 instances at 70% load
 - Remove instance when not needed
- Use for
 - Application servers
 - Report servers
 - Other?
- RAC?
 - Sadly, no.
 - It can be done, but quite complicated.
 - Ephemeral (non-persistent) storage
 - ASM mirroring for data integrity

OUT OF SPACE? – LEGACY METHOD

- Leaving out the paperwork...
- Buy new drives
 - Possibly new SAN disk drawer
 - Additional SAN licensing
- Install the Hardware
- How much time required?

OUT OF SPACE? – CLOUD METHOD

- Provision new drives
 - Add to EC2
 - Add disks to ASM
- How much time required?
 - Create volume
 - Select the EC2 instance
 - Add volume

CAPACITY PLANNING

- Seasonal capacity requirements?
- Add/Remove as needed
 - Space
 - CPU
- Need extra space required for data load?
 - Spin up volumes
 - Remove when no longer needed.

SERVERS ON DEMAND:

- Have you ever?
 - Asked all IT buddies for a server to test restore?
 - Begged for server to do POC?
- With AWS (or cloud of choice)
 - Test backup, restore, proof of concepts
 - All without begging, borrowing or purchasing hardware

AUTOMATION (OPEX)

- automation allows DBA's to do more with less
 - Yes, that is a trite cliché
- Scripting tasks that were previously unscriptable reduces the time needed for routine tasks.

How to Deploy Oracle

EC2 OR RDS?

- EC2
 - Traditional Oracle install on dedicated server
- RDS – Relational Database Service
 - Amazon controlled instance running DB
 - Oracle
 - SQL Server
 - MySQL
 - MariaDB
 - Aurora (Amazon's optimized MySQL)
 - PostgreSQL

RDS

- No root access
 - No ssh access at all
- Automated by Amazon
 - Backups
 - Patching 😊
 - Monitoring
 - Scaling
- Not for all DBs
 - But, pretty cool for those that fit the model.

Moving Data to the Cloud

DIY

- SSH for small stuff
- VPN
 - Bond VPN channels
(Ask Network Engineer about this for some detail)

AMAZON IMPORT/EXPORT

- provided by Amazon

- Import/Export Disk

AWS Import/Export



- Import/Export Snowball

Snowball



- Direct Connect

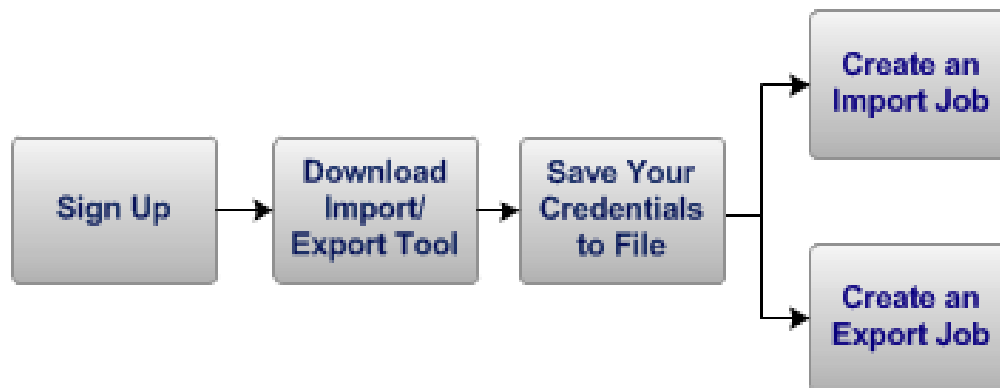
Direct Connect



IMPORT/EXPORT DISK



- Purchase Compatible NAS Device
 - eSATA
 - Compatible with Red Hat Linux
 - (not all commodity NAS devices are compatible)
 - \leq 16TB






IMPORT/EXPORT SNOWBALL

- Petabyte Scale Device
 - 50TB per device
- Supplied by Amazon
 - Order from your AWS account
 - \$200.00 US for 10 days
- Secure
 - Encrypted
 - Keys are NOT stored on device





DIRECT CONNECT

- Dedicated Network
 - Your data center  AWS
 - 1G or 10G speed
- Hybrid Data Center Uses
 - Backups offsite to S3/Glacier
 - Data Transfer to AWS Analytics
 - POC

AWS LIMITATIONS

- Limits per service
 - Can request increase
- Hard Limits
 - vCPU per instance
 - Memory per instance
 - Network IO
- Largest
 - 36 vCPU
 - 244G RAM
- If you need more than that...

APP PARTITIONING / DATA SHARDING

- Multiple instances
- Multiple databases
- App server directs connections
 - Server 1: Clients A-D
 - Server 2: Clients E-H
 - ...

Operating in the Cloud

DAY TO DAY

- Not
- That
- Much
- Different

BACKUPS

- Common Method - EBS Snapshots
 - Database in backup mode
 - Snapshot volumes
 - Archive all current redo
 - Save all to S3
- Kind of like old hot backups
 - Seems like a step backwards
- Use RMAN
 - Backup to F/S
 - Move to S3
 - Then archive to Glacier
- License OSB Cloud Module
 - Backup directly to S3
 - Archive to Glacier

POC, NEW FEATURES, RESTORE TESTS

- How often have you
 - Begged ...
 - Borrowed ...
 - Stolen
- A server so you can
 - Explore new features
 - Run a restore test
 - Perform a POC

DEMO'S AND TESTING

- 20 minutes
 - Launch Instance from Oracle AMI – c4.large (2 vCPU)
 - Create 11.2 db from Oracle AMI
- Needs more CPU!
- Traditional
 - New server (\$\$\$)
 - Restore and Recover new server
- AWS
 - Stop db and EC2 instance
 - Change instance type to c4.xlarge (4 vCPU)
- Total Cost of test:
 - \$0.48 US

DEMO'S AND TESTING CONT.

Billing for increase capacity test

EC2 Instance running Red Hat Enterprise Linux

\$0.17 per On Demand RHEL c4.large Instance Hour 1 Hrs \$0.17

\$0.28 per On Demand RHEL c4.xlarge Instance Hour 1 Hrs \$0.28

Total: \$0.45

EBS

\$0.065 per IOPS-month provisioned - US West (Oregon)

200G provisioned IOP at \$0.065 per month for 2 hrs

$.065 * 200 * (2/24/31) = \0.03

Total cost of test: \$0.48

Licensing

AWS CPU

- vCPU != Core
- AWS Core = 2 vCPU

- Don't trust me!
 - Talk to your Oracle Sales Rep

ORACLE ON AWS

- Simple
 - Or at least as simple as Oracle licensing always is...
- Example instance type: c4.8xlarge
 - vCPU: 36
 - Virtual Cores: 18 Intel x86
- Oracle Core Count
 - Don't forget core factor of .5
 - $18 \times .5 = 9$ cores
- 9 Cores must be licensed on AWS

THANK YOU!

- Download
 - <http://tinyurl.com/have-you-moved-to-the-cloud>



RESOURCES

- AWS Import/Export Disk
 - <https://aws.amazon.com/importexport/disk/>
- AWS Import/Export Snowball
 - <https://aws.amazon.com/importexport/>
- AWS Direct Connect
 - <https://aws.amazon.com/directconnect/>
- Virtual Cores chart
 - <http://aws.amazon.com/ec2/virtualcores/>
- AWS Instance Types
 - <https://aws.amazon.com/ec2/instance-types/>

RESOURCES

- AWS Icons
 - <https://aws.amazon.com/architecture/icons/>
- AWS Limitations
 - http://docs.aws.amazon.com/general/latest/gr/aws_service_limits.html
- AWS Total Cost of Ownership Calculator
 - <https://aws.amazon.com/tco-calculator/>
- Optimizing AWS Costs
 - <http://www.slideshare.net/AmazonWebServices/optimizing-total-cost-of-ownership-for-the-aws-cloud-36852296>

RESOURCES

- Dell R930 Pricing
 - <http://tinyurl.com/dell-r930>
- Deploying Scalable Oracle RAC on Amazon EC2
 - <https://aws.amazon.com/articles/7455908317389540>
- Linux Academy – AWS Training
 - <https://linuxacademy.com/>

RESOURCES – ORACLE LICENSING

- Pythian Blog on Licensing on AWS
 - <http://www.pythian.com/blog/licensing-oracle-in-a-public-cloud-the-cpu-calculation-impact/>
- Licensing Oracle Software in the Cloud
 - <http://www.oracle.com/us/corporate/pricing/cloud-licensing-070579.pdf>
- Oracle Processor Core Factor Table
 - <http://www.oracle.com/us/corporate/contracts/processor-core-factor-table-070634.pdf>
- Oracle Software Investment Guide
 - <http://www.oracle.com/us/corporate/pricing/sig-070616.pdf>

RESOURCES – AWS CORES/vCPU

- Virtual Cores
 - <http://aws.amazon.com/ec2/virtualcores/>
- Virtual CPUs
 - <http://aws.amazon.com/ec2/instance-types/>
- 1 Core = 2 vCPU

