

ORACLE®

Virtual Technology Summit

Hands-On Learning With Oracle and Community Experts

Where Technology and Community Meet



Oracle NoSQL Database 3.0

Installation, Cluster Topology Deployment, HA and more

Seth Miller, Oracle ACE

Robert Greene, Product Management / Strategy

Oracle Server Technologies

July 09, 2014

Safe Harbor Statement

The following is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described for Oracle's products remains at the sole discretion of Oracle.

Download Oracle NoSQL Database

- Go here and download ONDB and unzip into a directory (will be your KVHOME)
 - http://download.oracle.com/otn-pub/otn_software/nosql-database/kv-ce-3.0.9.zip

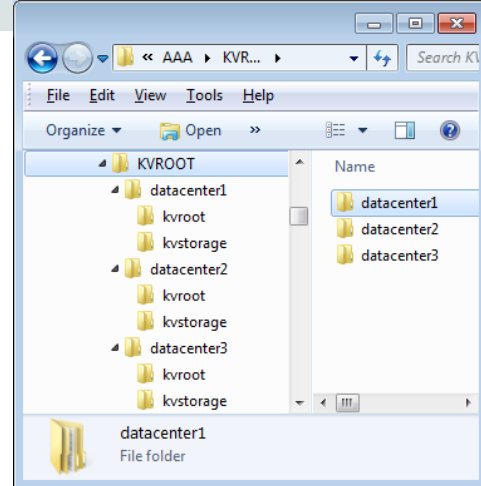
Steps to Deploy ONDB

- Create a config file with host, port, capacity, memory etc.
- Start Storage Node Agent
- On this storage node agent
 - Name your KV store
 - Deploy a data center
 - Add storage node to the data center
 - Deploy admin service for web GUI
 - Create a storage node pool
- Add other Storage node agents to the data center
- Create and deploy topology

NoSQL HOL - Getting to Work

- Setup for work

- Create an installation directory somewhere called “kvhome”
- Unzip the Lab content into that kvhome directory
- Create a data directory called “KVROOT”, subdirs “datacenter1-3” subdirs “kvroot”, “kvstorage”
 - » See Picture for details
- Setup your environment for the scripts
 - Look at setenv.bat
 - set KVHOME=%USERPROFILE%\nosql\kvhome
 - set KVROOT=%USERPROFILE%\nosql\KVROOT
 - set CLASSPATH=%KVHOME%\lib
 - set PATH=C:\Program Files (x86)\Java\jre7\bin;%PATH%



NoSQL HOL - Getting to Work

- Make sure your environment is finding the install
 - User> java -jar %KVHOME%\lib\kvstore.jar

Usage: java -jar KVHOME/lib/kvstore.jar
<kvlite | makebootconfig | securityconfig |
start | stop | restart | runadmin | load | ping | version |
generateconfig | help> [-verbose] [args]

NoSQL HOL - Getting to Work

- Step 1 - Create a Storage Node configuration file
 - This is done 1 per physical machine or 1 per VM
 - Examine the script `makeboot.bat`
 - Notice: capacity, CPUs, memory, etc
 - Notice: we play with ports because VM's on 1 machine
 - If separate machines, then all ports can be the same
 - Run it **user>makeboot.bat**
 - Go checkout the `config.xml` file

NoSQL HOL - Getting to Work

- Step 2 – Start the Storage Node agents
 - This is done 1 per physical machine or 1 per VM
 - Examine the script `start.bat` (Notice: a simple command points to KVROOT)
 - Use processExplorer or Task Manager: (Notice no processes from kvstore)
 - Run it **user>start.bat**
 - Use processExplorer or Task Manager : (Notice processes from kvstore)
 - Processes listen for boot strapping commands to setup / extend clusters

NoSQL HOL - Getting to Work

- Step 3 – Start an administration CLI session
 - This is done on one of the storage nodes in the cluster
 - Examine the script `startadmin.bat`
 - Notice: uses the agent port, puts you into a bootstrap mode
 - Run it `user>startadmin.bat`
 - Notice: you are now at an administrative command prompt

NoSQL HOL - Getting to Work

- Step 4 – Create / deploy the bootstrap zone and admin U.I. service
 - This is defining a logical Zone, assigning Storage Nodes, setting replication
 - Examine the script `cstore1_win.src` (file extension is meaningless)
 - Notice: this is a series of commands for the admin CLI
 - Notice: we are running them as a script, but you can do each line in CLI
 - Run it `kv>load -file cstore1_win.src`
 - Notice: you just defined a cluster (store) named “ondb”
 - Open a browser and type in the url: `localhost:5899`see the topology

NoSQL HOL - Getting to Work

- Step 5 – Create / deploy 2 more Zones, with Storage Nodes
 - This is finishing the declaration of resources and defining a Topology
 - The topology is not actually deployed until the very last step
 - Be prepared, go to your browser as soon as you run script
 - Examine the script `cstore2_win.src`
 - Notice - adding resources to a pool, then telling Topology use the pool
 - Run it `kv>load -file cstore2_win.src` ...hurry to your browser
 - Now you have a full Topology created and a cluster Deployed

NoSQL HOL - Getting to Work

- Step 6 – Lets talk about the Topology created, SN, RG and HA topic
 - While examining the Admin Browser, hover over RG links and notice SN shift
 - What you will see is that each replica in RG lands in separate SN
 - Examine the SN's, replicas and masters
 - Run it **kv>ping**
 - Now you can see where the master is For the moment

NoSQL HOL - Getting to Work

- Step 7 – Lets kill some hardware and see what happens
 - Open for editing the stop.bat file, and copy “datacenter3” line
 - In another terminal window, execute the line to stop a SN
 - like pulling power on machine
 - **user>start java -jar %KVHOME%\lib\kvstore.jar stop -root %KVROOT%\datacenter3\kvroot**
 - Go checkout your browser and see it is telling you there are problems
 - Back to the admin shell, run it again **kv>ping**
 - You can see master has changed, one data center down
 - Still have read / write access even after having killed a master

NoSQL HOL - Getting to Work

- Step 8 – Bring the cluster back fully online
 - Open for editing the start.bat file, and copy the “datacenter3” line
 - In another terminal window and execute the line to start the SN
 - like turn power on machine
 - **user>start /b java -jar %KVHOME%\lib\kvstore.jar start -root %KVROOT%\datacenter3\kvroot**
 - Go checkout your browser and see it is telling things are getting better
 - Back to the admin shell, run it again **kv>ping**
 - Now you can see the master has changed again, no admin, but automatic HA optimization

NoSQL HOL – Using the NoSQL Cluster

- We are going to define some Tables, import data, do some Queries
 - Using data CLI to avoid JDK, IDE, OS issues in compiling apps
 - Creating tables is straight forward, nothing to do on this slide, just review it:

Create table

table create -name shardUsers

Add the fields to the table

add-field -name firstName -type STRING

add-field -name lastName -type STRING

add-field -name email -type STRING

Define primary keys, then shard key as subset

primary-key -field lastName -field firstName

Add a shard key.

shard-key -field lastName

exit

Add the table plan to the database

plan add-table -wait -name shardUsers

NoSQL HOL – Using the NoSQL Cluster

- Step 1 - Lets put some Table schema into our cluster
 - will be using the data CLI to avoid having to deal with JDK 1.7, IDE's, etc
 - Start the DataCLI and define some Table definitions
 - Note – Schema is not strictly necessary, except if you want indexing
 - Run CLI **user>datacli.sh**
 - Import Table definitions
 - Run load command **kv>load -file createTable.src**
 - Run load command **kv>load -file createChildTable.src**
 - Run load command **kv>load -file createComplexFields.src**

NoSQL HOL – Using the NoSQL Cluster

- Step 2 - Lets take a look at our Table schema created
 - Assume the DataCLI is still running, if not start it
 - Show the Table definitions
 - Run command **kv>show tables**
 - This will give you the list of all tables names
 - Run command **kv>show table –name <a name from the list>**
 - **kv>show table –name shardUsers**
 - **kv>show table –name shardUsers –child shardUsers.address**
 - Notice: Shard Key, Primary Key, Fields <typed>, Indexes, etc

NoSQL HOL – Using the NoSQL Cluster

- Step 3 - Lets import some data into our cluster
 - Assume the DataCLI is still running, if not start it
 - Data in a standard JSON file format can be imported / exported from database
 - Run commands:
 - **kv-> put table -name complexUsers -file complexUsers.txt**
 - This will take the JSON data in .txt file and load into cluster
 - **kv-> put table -name shardUsers -file shardUsers.txt**
 - **kv-> put table -name shardUsers.address -file addresses.txt**

NoSQL HOL – Using the NoSQL Cluster

- Step 4 - Lets look at the data in our cluster
 - Assume the DataCLI is still running, if not start it
 - Run commands:
 - **kv-> get table -name complexUsers**
 - **kv-> get table -name shardUsers**
 - **kv-> get table -name shardUsers.address**
 - **kv-> get table -name shardUsers -child shardUsers.address**
 - Notice here that it's a hierarchy, and data is pushed to the child

NoSQL HOL – Using the NoSQL Cluster

- Step 5 - Lets Query the data in our cluster
 - Assume the DataCLI is still running, if not start it
 - Use get on a KEY FIELD to query based on KeySpace, no need for index
 - Note – these are ordered by their definition
 - Run commands:
 - **kv-> get table -name complexUsers -field userID -value 2**

NoSQL HOL – Using the NoSQL Cluster

- Step 6 - Lets Query the data in our cluster
 - Assume the DataCLI is still running, if not start it
 - Use get on an Index to query based on non-key fields
 - Note – these results are ordered store wide
 - Note – see index names from previous step kv->show table –name <tname>
 - Run commands:
 - kv-> **get table -name complexUsers -index arrayIndex -field likes -value movies**
 - kv-> **get table -name shardUsers.address -index addressIndex -field addressName -value work**

NoSQL HOL – Using the NoSQL Cluster

- Step 7 - Lets Query the data in our cluster (Assume the DataCLI running)
 - Use “get” to query based on compound key fields
 - Run **kv->show tables -name shardUsers** :: see PK definition ordering
 - Run commands:
 - **kv-> get table -name shardUsers -field lastName -value Robertson** :: Succeeds
 - **kv-> get table -name shardUsers -field firstName -value Beatrix** :: Fails
 - **kv-> get table -name shardUsers -field lastName -value Robertson -field firstName -value Beatrix** :: Succeeds ... note ordering does not matter

NoSQL HOL – Using the NoSQL Cluster

- Step 8 - Lets Query the data in our cluster (Assume the DataCLI running)
 - Use get to query based on a range specification
 - Run **kv->get table -name shardUsers** :: see the data ranges
 - Run commands:
 - **kv-> get table -name shardUsers -field lastName -start Q -end T**
 - **kv-> get table -name shardUsers -field lastName -start I -end K**
 - Try the same query from its child table, notice returns all hierarchy data
 - **kv-> get table -name shardUsers.address -field lastName -start I -end K**

NoSQL HOL - Conclusions

- Getting a cluster setup is super simple
 - Install the software and create config files on each storage node machine
 - Start some agents on those machines
 - Use Admin to create Topology definition of Zones and add your storage nodes
 - Deploy the Topology
 - Use an Admin to create some Table schema, load some data, query

NoSQL HOL - Conclusions

- Easy access and simple query operations cluster wide
 - Use key space definition (primary key of the Table)
 - Use indexes for secondary, value based queries of non-primary keys
 - Access to array fields and ordered compound keys and range filtering

Oracle NoSQL Database

Product editions and licensing

- **Community -or- Enterprise**
 - Apache licensed client drivers
 - Facilitates OEM product integration
 - **Community Edition has all of the basic functionality** and APIs. Gets you started
 - Subscription option (2k/server)
 - **Enterprise Edition** for multi-data center, **Oracle integration-centric** customers and/or non-GPL compliant customers
 - Standard Oracle core based (10k/cpu)

ORACLE®

SUPPORT

Oracle NoSQL Database

Summary

- **Enterprise ready**- for administrators and developers
- **Integrated** - into Oracle software stack of tools and solutions
- **Engineered** - fast deployment and comprehensive support
- **Best NoSQL choice for any company using Oracle products**

Join NoSQL Database Community

Oracle.com/BigData



Twitter

<https://twitter.com/#!/OracleNoSQL>



LinkedIn

<http://www.linkedin.com/groups?gid=4147754>



Oracle's NoSQL DB blog

<https://blogs.oracle.com/nosql>



Oracle Technology Network

<http://bit.ly/1f0d8wU>



Developer Webcast Series

<http://bit.ly/1doV2jl>



Hardware and Software Engineered to Work Together

ORACLE®