

# Assessment Plan - Clustering

## Plan:

Set up a Fedora 4 load balanced cluster, perform tests including load-balance read requests across the cluster, nodes joining/leaving the cluster, and throughput/response time comparison of a single Fedora and a cluster.

## Description:

Currently the Fedora 4 cluster performance testing was conducted in a single machine [1], and load balanced cluster setup was conducted using Fedora Beta 1 version [2]. After Fedora 4 Beta 2 was released, the version of Modeshape and Infinispan were changed and the existing cluster configuration document does not apply to the new Fedora 4 release. Our plan at VT is to set up a Fedora 4 load balanced cluster using n AWS EC2 instances, where  $n \geq 3$ , to do the following tasks:

1. Verify the load balanced cluster setup using the updated Modeshape and Infinispan configuration.
2. Demonstrate the nodes joining and node leaving in the cluster.
3. Measure the cluster's throughput and response time
  - a. Continue increasing the load until server is unable to handle further requests from the client.
  - b. Exam if the n-node ( $n \geq 2$ ) cluster handles n times as much requests as a single instance
  - c. n-JMeter ( $n=2$ ) clients to send requests in different region/availability zone and measure the response time.
  - d. Simulate real world case, 95% read requests and 5% write requests.
4. Write up installation document
5. Write up performance report

## Software:

Fedora 4: master branch version (09/18/14) <https://github.com/fcrepo4/fcrepo4>

JMeter: <http://jmeter.apache.org/>

Loadosophia: <https://loadosophia.org/>

## Test Environment:

AWS: EC2 t2.small or medium instances.

OS: Ubuntu 14.04 LTS.

## Test Procedure:

1. Verify the load balanced cluster setup using the updated Modeshape and Infinispan configuration.
  - a. Following updated configuration files and document to set up Fedora 4 cluster
  - b. Configured the cluster in replication mode
  - c. Load objects into one of the instances
  - d. Configure load balancer to equally balance requests across these three instances.
  - e. Use JMeter to send read/write requests.
2. Demonstrate the nodes joining and node leaving in the cluster.
  - a. Create new Fedora instances
  - b. Add new instances into cluster
  - c. Perform read requests as in Test 1
  - d. Remove instances from cluster
  - e. Perform read/write requests as in Test 1
3. Measure the single and cluster's throughput and response time
  - a. Use JMeter to send read requests to a single Fedora instances
  - b. Use JMeter to send read requests to a Fedora Cluster
4. Continue increasing the load until server is unable to handle further requests from the client.
5. Exam if the n-node ( $n \geq 2$ ) cluster handles n times as much requests as a single instance.
6. n-JMeter ( $n=2$ ) clients to send requests in different region/availability zone and measure the response time.
  - a. We choose 4 different zone, US-East, US-West, Asia and EU.
  - b. Each zone run 2 JMeters and calculate the average response time.
7. Simulate real world case, 95% read requests and 5% write requests.
  - a. One JMeter do write request while another JMeter do read request.

## Reference:

[1] <https://wiki.duraspace.org/display/FF/Response+Time+Comparison+of+Single+Fedora+VS+Cluster>

[2] <https://wiki.duraspace.org/display/FF/2014-08-10+Acceptance+Test+-+Load+balanced+cluster+setup>