

Fedora 4: Technical Working Group Assessment

Open Repositories 2015, Indianapolis IN

Esmé Cowles

Zhiwu Xie

Andrew Woods

Clustering

Zhiwu Xie

Virginia Tech

Current Scope of Clustering

- High availability
- Load balance
- Redundancy

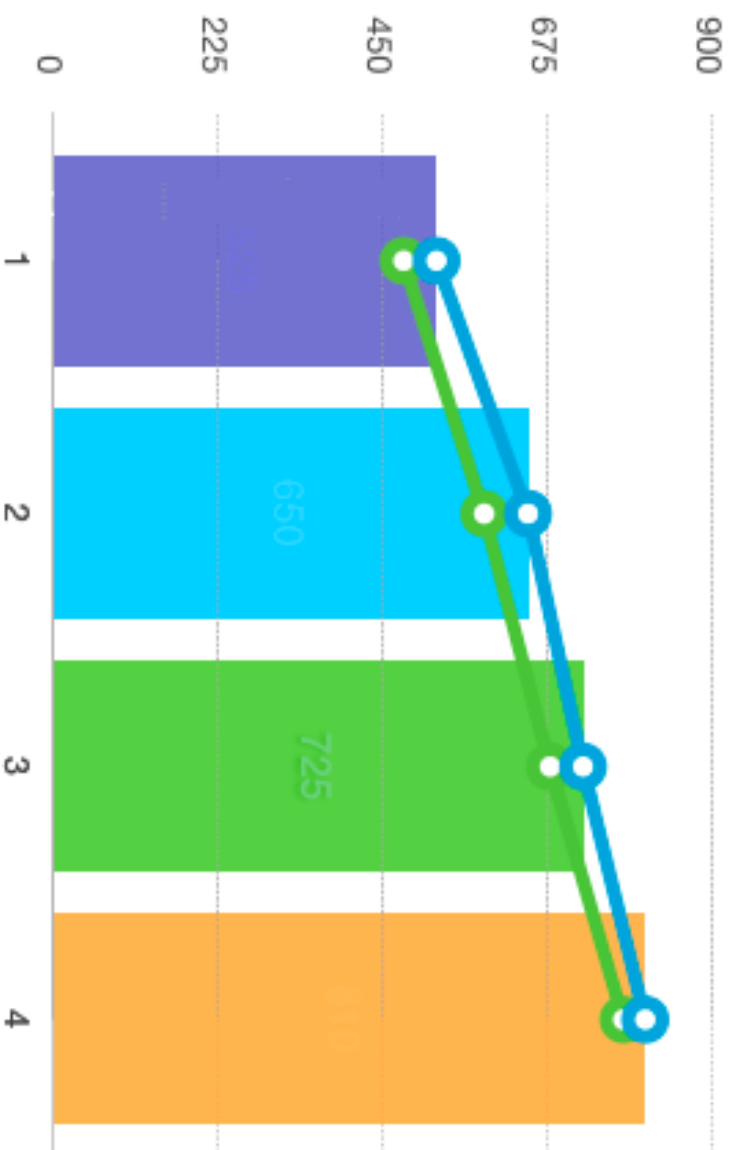
Objective: Gain better understanding on clustering performance and benefits

Clustering Test Configurations

- Test on AWS with EC2 nodes with a single load balancer
- Replicated nodes in the same local network
 - Read-any: any node can handle read requests without consulting other nodes
 - Write-all: any node can handle write request, but write is not completed until replicated to all the other nodes
- Test against read-only and read-dominant (5% write, 95% read) workloads
 - Write does not present as a bottleneck in our experiments

Results

- Read capacity increases linearly under both read-only and read-dominant workloads



Results

- When workload is within server capacity, read latency does not change significantly with the number of nodes
- When workload is within server capacity, write latency shows signs of increase when the number of replicated nodes increases
- Can add and remove nodes at will via load balancing

Summary

- Provide high confidence to use replicated Fedora 4 cluster for high availability and load balanced use cases