



NETFLIX

Netflix Data Benchmark (NDBench)


Circa end of 2015 AD

NETFLIX





NETFLIX



NETFLIX

Netflix Data Benchmark (NDBench)

Vinay Chella

Ioannis Papapanagiotou

Architects

Cloud Database Engineering @ Netflix

NETFLIX ORIGINAL
STRANGER THINGS

95% Match 2016 1 Season 4K Ultra HD 5.1

When a young boy vanishes, a small town uncovers a mystery involving secret experiments, terrifying supernatural forces and one strange little girl.

Winona Ryder, David Harbour, Matthew Modine
TV Shows, TV Sci-Fi & Fantasy, Teen TV Shows




Popular on Netflix



Recently Watched



Databases and Caches




Netflix Data Scale

- Tens of petabytes of data
- Trillions of ops / day
- Tens of Millions of replications / sec
- Tens of Thousands of servers
- Hundreds of microservice clients

Agenda

- Background
- Why NDBench?
- Architecture
- Auto Tuning
- Usage @ Netflix





Our Needs

- **Single benchmark** framework for all data stores/services
- **Dynamic configuration** while the test is running.
- Test a platform along with production microservices
- **Integrate** with other platform cloud services
- **Pluggable** patterns and loads
- **Auto tunable** framework
- Run **infinite horizon tests**

What is NDBench?

Netflix Data Benchmark (NDBench) is a **Pluggable cloud-enabled** benchmarking tool that can be used across **any data store system/ micro service.**



NDBench allows

- Run **infinite horizon** tests
- Test out Chaos enabled **failure scenarios**
- Performance test of **heavy processes**
- Coordinated **rate limiters**



NDBench Features

Side-by-Side comparison



NETFLIX

Varying Data Models

```
sunomia_cass_app_node_info_1
  appname : text
  compaction_schedule : text
  createdby : text
  customerddl : text
  customergroup : text
  customerprd : text
  dbengowner : text
  env : text
  externalbackuplocation : text
  externalbackupretention : text
  externalbackupschedule : text
  isactive : boolean
  isbatchcluster : boolean
  iscustomerfacing : boolean
  isexternalbackupenabled : boolean
  ismanaged : boolean
  ismonitored : boolean
  isprimarybackupenabled : boolean
  isredblack : boolean
  isrefreshed : boolean
  json_val : text
  mapped_to : set<text>
  messengerlink : text
  nodeinfojson : text
  nodeuid : text
  other_attr : map<text,text>
  primarybackuplocation : text
  primarybackupretention : text
  primarybackupschedule : text
  refreshed_from : set<text>
  refreshed_to : set<text>
  region : text
  repair_schedule : text
```

VS


```
sunomia_cass_app_node_info_1
  appname : text
  env : text
  region : text
  nodeuid : text
  compaction_schedule : text
  createdby : text
  customerddl : text
  customergroup : text
  customerprd : text
  dbengowner : text
  externalbackuplocation : text
  externalbackupretention : text
  externalbackupschedule : text
  isactive : boolean
  isbatchcluster : boolean
  iscustomerfacing : boolean
  isexternalbackupenabled : boolean
  ismanaged : boolean
  ismonitored : boolean
  isprimarybackupenabled : boolean
  isredblack : boolean
  isrefreshed : boolean
  mapped_to : set<text>
  messengerlink : text
  nodeinfojson : text
  other_attr : map<text,text>
  primarybackuplocation : text
  primarybackupretention : text
  primarybackupschedule : text
  refreshed_from : set<text>
  refreshed_to : set<text>
  repair_schedule : text
```

NETFLIX

Different instance types



Heterogeneous service comparison



cassandra




JanusGraph

Workload Patterns and Loads



NETFLIX

Different Client APIs



NETFLIX

Cloud Ecosystem integration

NETFLIX | OSS



Cloud Discovery



Properties



Metrics



NETFLIX

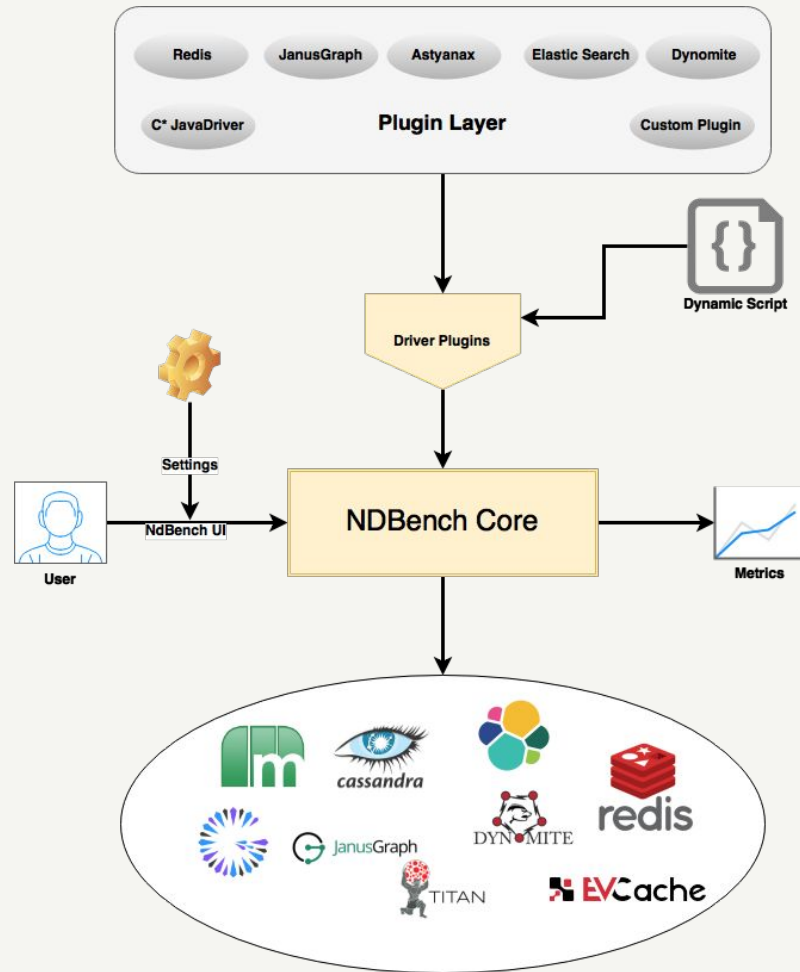
Agenda

- Background
- Why NDBench?
- **Architecture**
- Usage @ Netflix
- Auto Tuning



Architecture

- **Plugin:** Adding Client drivers
- **Core:** Workload generator
- **Web:** UI and the servlets



What is Pluggable?

- Client Plugins
- Load Patterns
- Configurations
- Metrics

Load Patterns

- Random
 - Random load pattern cannot exercise the caches.
- Sliding Window
 - Leverages cache and disk IOs

Configuring a cluster

TEST



ndb_qcon--useast1c

8 ▲ : 100%

US-EAST-1



N/A: Build: #67

8 ▲ : 100%



TEST



ndb_qcon--useast1d

8 ▲ : 100%

US-EAST-1



N/A: Build: #67

8 ▲ : 100%



TEST



ndb_qcon--useast1e


8 ▲ : 100%

US-EAST-1



N/A: Build: #67

8 ▲ : 100%



NETFLIX

Continuous Delivery Integration

Properties

Filter

+ Create Persisted Property

PROPERTIES

ROLLOUTS

Showing 12 results

| Property ▾ | Value ▾ | Env ▾ | Summary | Region ▾ | Stack ▾ | Upd |
|--------------------------------|---------------|-------|---------------|----------|---------|-----|
| ndbench.config.cass.cfname | test1 | TEST | App: ndb_qcon | | | 201 |
| ndbench.config.cass.cluster | cass_qcon1 | TEST | App: ndb_qcon | | | 201 |
| ndbench.config.cass.colsPerRow | 10 | TEST | App: ndb_qcon | | | 201 |
| ndbench.config.cass.keyspace | pappyperftest | TEST | App: ndb_qcon | | | 201 |
| ndbench.config.dataSize | 200 | TEST | App: ndb_qcon | | | 201 |
| ndbench.config.numBackfill | 100 | TEST | App: ndb_qcon | | | 201 |
| ndbench.config.numKeys | 1000000 | TEST | App: ndb_qcon | | | 201 |
| ndbench.config.numReaders | 4 | TEST | App: ndb_qcon | | | 201 |
| ndbench.config.numValues | 1000 | TEST | App: ndb_qcon | | | 201 |

1. Select a Cluster

Cluster
localhost

2. Connect a Driver

Driver
InMemoryTest
InMemoryTestPlugin - ConnectionInfo: InMemoryMap Key Count: 0

HIDE SETTINGS

DISCONNECT DRIVER

Initial Settings

These settings can only be changed before connecting a Driver.

backfillStartKey

dataSize

dataSizeLowerBound

dataSizeUpperBound

numBackfill

numKeys

numReaders

numValues

numWriters

readRateLimit

statsResetFreqSeconds

statsUpdateFreqSeconds

writeRateLimit

readEnabled

useStaticData

useVariableDataSize

writeEnabled

Runtime Settings

These settings can be changed at any time while a Load Test is running.

readRateLimit

writeRateLimit

SAVE SETTINGS

3. Run Load Tests


Load Pattern
RANDOM

| Instance | Backfill | Writes | Reads |
|----------------|----------|------------------|------------------|
| localhost:8080 | | | |
| | | RandomString_... | RandomString_... |

Instance Statistics

localhost:8080

— readLatAvg — readRPS — writeLatAvg — writeRPS



- cacheHitRatioInt
 cacheHits
 cacheMiss
 readFailure
- readLatAvg
 readLatP50
 readLatP95
 readLatP99
- readLatP995
 readLatP999
 readRPS
 readSuccess
- writeFailure
 writeLatAvg
 writeLatP50
 writeLatP95
- writeLatP99
 writeLatP995
 writeLatP999
 writeRPS
- writeSuccess

HIDE STATS JSON

```
{
  "cacheHitRatioInt": 29,
  "cacheHits": 2791,
  "cacheMiss": 6659,
  "readFailure": 0,
  "readLatAvg": 2,
  "readLatP50": 1,
  "readLatP95": 1,
  "readLatP99": 1,
  "readLatP995": 1,
  "readLatP999": 3,
  "readRPS": 25,
  "readSuccess": 9450,
  "writeFailure": 0,
  "writeLatAvg": 3,
  "writeLatP50": 3,
  "writeLatP95": 5,
  "writeLatP99": 6,
  "writeLatP995": 7,
  "writeLatP999": 14,
  "writeRPS": 15,
  "writeSuccess": 7267
}
```

1. Select a Cluster

Cluster
 NDB_QCON

2. Connect a Driver

Driver
 A6XGeneric
 Cluster Name - cass_qcon1 : Keyspace Name - pappypferfest : CF Name - test1 :: ColsPerRow - 10 :
 ReadCL - CL_LOCAL_ONE : WriteCL - CL_LOCAL_ONE

SHOW SETTINGS DISCONNECT DRIVER


3. Run Load Tests

Load Pattern
 RANDOM

| Instance | Backfill | Writes | Reads |
|-------------------------------------|----------|--------|-------|
| ...compute-1.amazonaws.com:7001 | ▶ ◻ | ▶ ◻ | ▶ ◻ |
| ec2-...compute-1.amazonaws.com:7001 | ▶ | ◻ | ▶ |
| ec2-...compute-1.amazonaws.com:7001 | ▶ | ▶ | ▶ |
| ec2-...compute-1.amazonaws.com:7001 | ▶ | ▶ | ▶ |
| ec2-...compute-1.amazonaws.com:7001 | ▶ | ▶ | ▶ |
| ec2-...compute-1.amazonaws.com:7001 | ▶ | ▶ | ▶ |
| ec2-...compute-1.amazonaws.com:7001 | ▶ | ▶ | ▶ |
| ec2-...compute-1.amazonaws.com:7001 | ▶ | ▶ | ▶ |

Instance Statistics

ec2-34-203-200-0.compute-1.amazonaws.com:7001



- readSuccess
- readLatP50
- readLatP95
- readLatP99
- readLatP995
- readLatP999
- readFailure
- readLatP95
- readLatP99
- readLatP995
- readLatP999
- writeSuccess
- writeLatP99
- writeLatP995
- writeLatP999
- cacheHits
- cacheMiss
- cacheHitRatioInt
- writeFailure
- writeLatP95
- writeLatP99
- writeLatP995
- status
- readRPS
- writeRPS
- stats
- writeLatP99
- readLatAvg
- writeLatP995

SHOW STATS JSON

selecting a cluster

1. Select a Cluster

Cluster
NDB_QCON

2. Connect a Driver

Driver
A6XGeneric
Cluster Name - cass_qcon1 :: Keyspace Name - pappyperftest :: CF Name - test1 :: ColsPerRow - 10 ::
ReadCL - CL_LOCAL_ONE :: WriteCL - CL_LOCAL_ONE

SHOW SETTINGS DISCONNECT DRIVER

3. Run Load Tests

Load Pattern
RANDOM

| Instance | Backfill | Writes | Reads |
|---|-------------------------------|-------------------------------|-------------------------------|
| ec2-34-203-200-0.compute-1.amazonaws.com:7001 | ▶ ⏸ | ▶ ⏸ | ▶ ⏸ |
| ec2-34-203-200-0.compute-1.amazonaws.com:7001 | ▶ | RandomStri... ⏸ | RandomStri... ⏸ |
| ec2-34-203-200-0.compute-1.amazonaws.com:7001 | ▶ | ▶ | ▶ |
| ec2-34-203-200-0.compute-1.amazonaws.com:7001 | ▶ | ▶ | ▶ |
| ec2-34-203-200-0.compute-1.amazonaws.com:7001 | ▶ | ▶ | ▶ |
| ec2-34-203-200-0.compute-1.amazonaws.com:7001 | ▶ | ▶ | ▶ |
| ec2-34-203-200-0.compute-1.amazonaws.com:7001 | ▶ | ▶ | ▶ |
| ec2-34-203-200-0.compute-1.amazonaws.com:7001 | ▶ | ▶ | ▶ |

Instance Statistics

ec2-34-203-200-0.compute-1.amazonaws.com:7001

readSuccess readFailure writeSuccess writeFailure readRPS writeRPS



- readSuccess
- readLatP50
- readLatP99
- readLatP995
- readLatP999
- readFailure
- readLatP50
- readLatP99
- readLatP995
- readLatP999
- writeSuccess
- writeLatP50
- writeLatP99
- writeLatP995
- writeLatP999
- writeFailure
- writeLatP50
- writeLatP99
- writeLatP995
- writeLatP999
- cacheHits
- cacheMiss
- readRPS
- writeRPS
- stats
- readLatAvg
- writeLatAvg
- writeLatP99
- writeLatP95
- writeLatP99
- writeLatP95
- status
- documentation
- status

SHOW STATS JSON

2. Connect a Driver

- Driver
- DynoPlugin
 - ES
 - ElassandraCassJavaDriverPlugin
 - InMemoryTest
 - MaterializedView
 - Custom Script

Backfill Writes Reads

1. Select a Cluster

Cluster
NDB_QCON

2. Connect a Driver

Driver
A6XGeneric
Cluster Name - cass_qcon1 : Keyspace Name - pappyperftest : CF Name - test1 :: ColsPerRow - 10 :
ReadCL - CL_LOCAL_ONE : WriteCL - CL_LOCAL_ONE

SHOW SETTINGS DISCONNECT DRIVER


3. Run Load Tests

Load Pattern
RANDOM

| Instance | Backfill | Writes | Reads |
|---|-------------------------------------|-------------------------------------|-------------------------------------|
| ec2-34-203-200-0.compute-1.amazonaws.com:7001 | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> |
| ec2-34-203-200-0.compute-1.amazonaws.com:7001 | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> |
| ec2-34-203-200-0.compute-1.amazonaws.com:7001 | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> |
| ec2-34-203-200-0.compute-1.amazonaws.com:7001 | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> |
| ec2-34-203-200-0.compute-1.amazonaws.com:7001 | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> |
| ec2-34-203-200-0.compute-1.amazonaws.com:7001 | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> |
| ec2-34-203-200-0.compute-1.amazonaws.com:7001 | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> |
| ec2-34-203-200-0.compute-1.amazonaws.com:7001 | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> |

Instance Statistics

ec2-34-203-200-0.compute-1.amazonaws.com:7001



- readFailure
- readLatP50
- writeLatP999
- cacheMiss
- readLatP95
- cacheHitRatioInt
- readLatP99
- documentation
- readLatP995
- status
- readRPS
- writeLatP50
- writeLatP95
- stats
- writeRPS
- writeLatP99
- readLatAvg
- writeLatP995

SHOW STATS JSON

1. selecting a load pattern

2. starting the workload

1. Select a Cluster

Cluster
NDB_QCON



Monitor

2. Connect a Driver

Driver
A6XGeneric
Cluster Name - cass_qcon1 : Keyspace Name - pappyperfest : CF Name - test1 ::: ColsPerRow - 10 :
ReadCL - CL_LOCAL_ONE : WriteCL - CL_LOCAL_ONE
SHOW SETTINGS **DISCONNECT DRIVER**


3. Run Load Tests

Load Pattern
RANDOM

| Instance | Backfill | Writes | Reads |
|-------------------------------------|----------|--------|-------|
| compute-1.amazonaws.com:7001 | | | |
| ec2-...compute-1.amazonaws.com:7001 | | | |
| ec2-...compute-1.amazonaws.com:7001 | | | |
| ec2-...compute-1.amazonaws.com:7001 | | | |
| ec2-...compute-1.amazonaws.com:7001 | | | |
| ec2-...compute-1.amazonaws.com:7001 | | | |
| ec2-...compute-1.amazonaws.com:7001 | | | |
| ec2-...compute-1.amazonaws.com:7001 | | | |

Instance Statistics

ec2-34-203-200-0.compute-1.amazonaws.com:7001



NDBench exemplar uses





CentOS

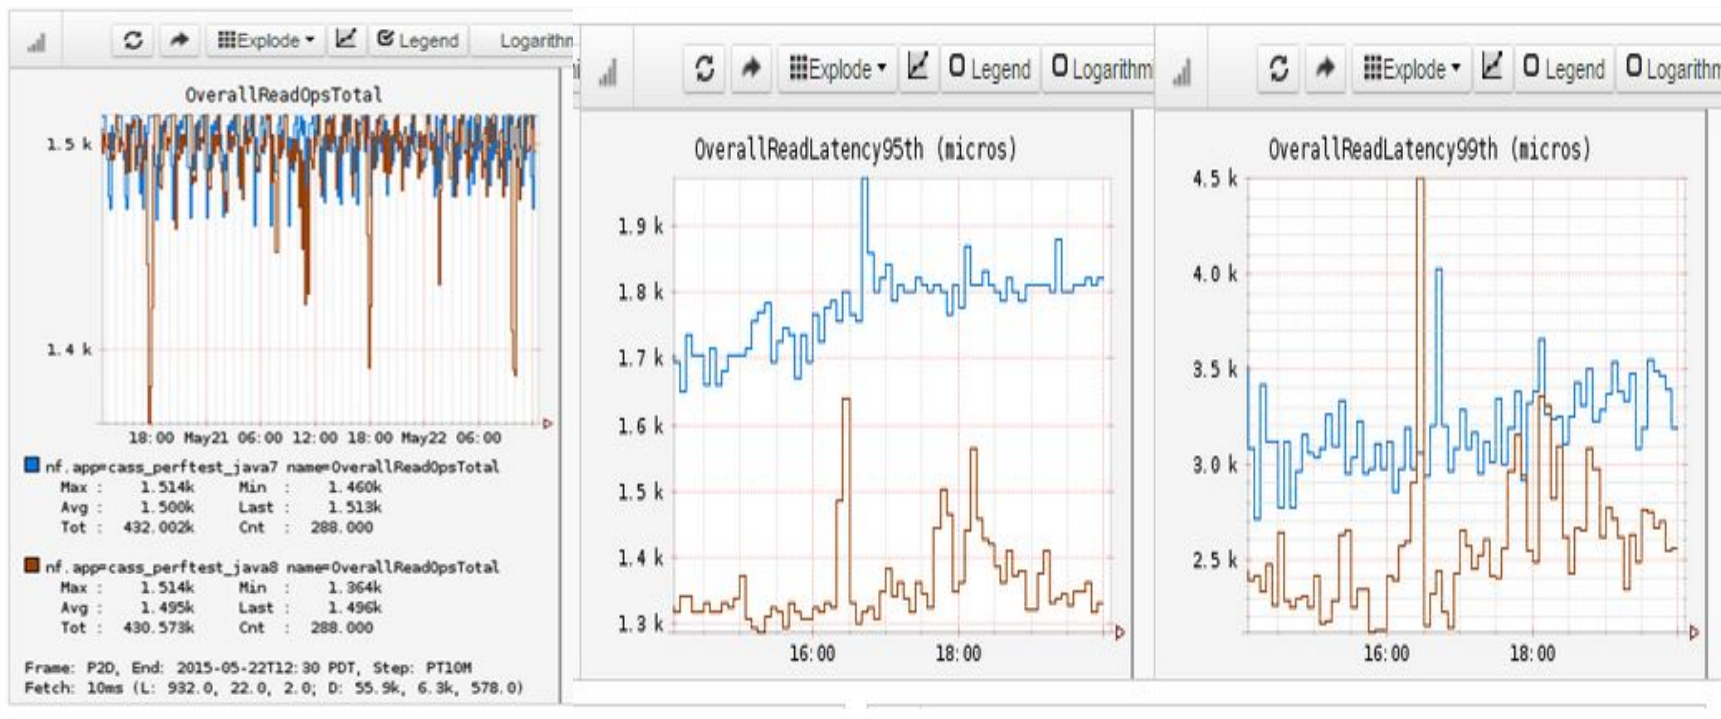


CASS | cass_perfest_trustyupg




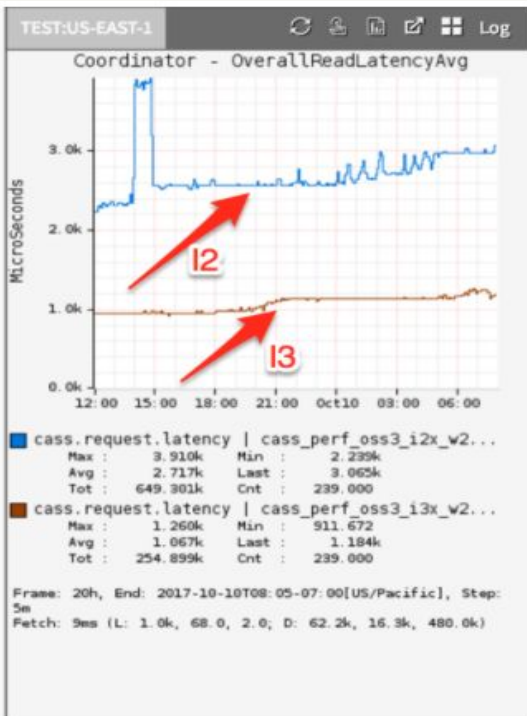
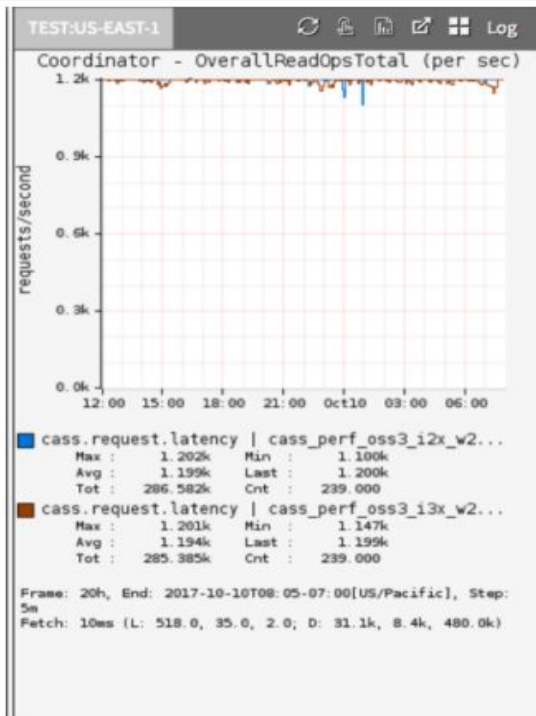
NETFLIX

C* 7 to 8 (reads)

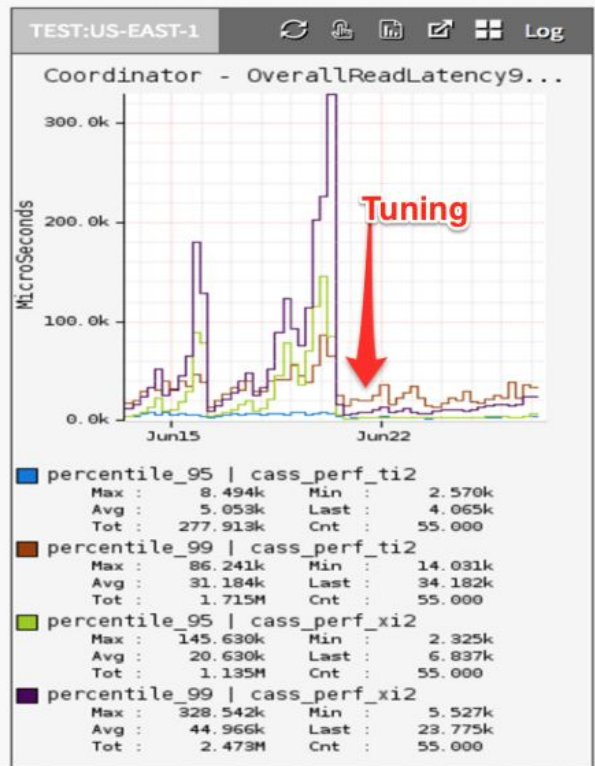
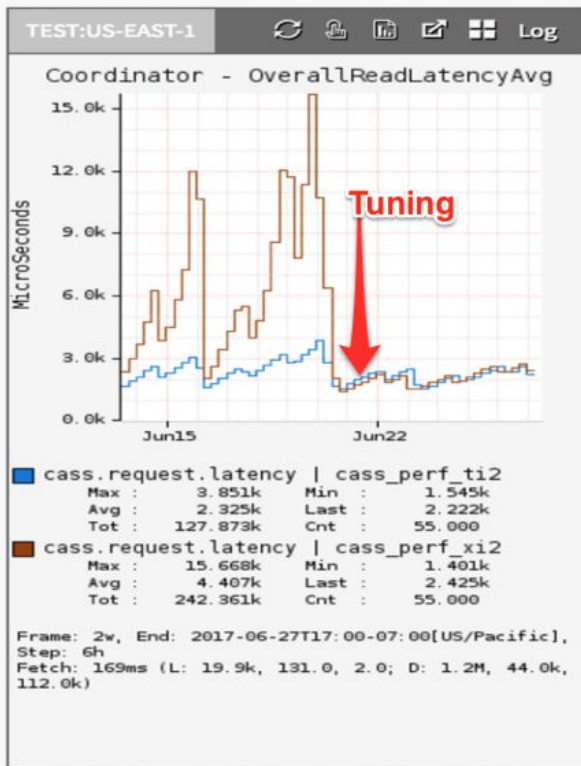
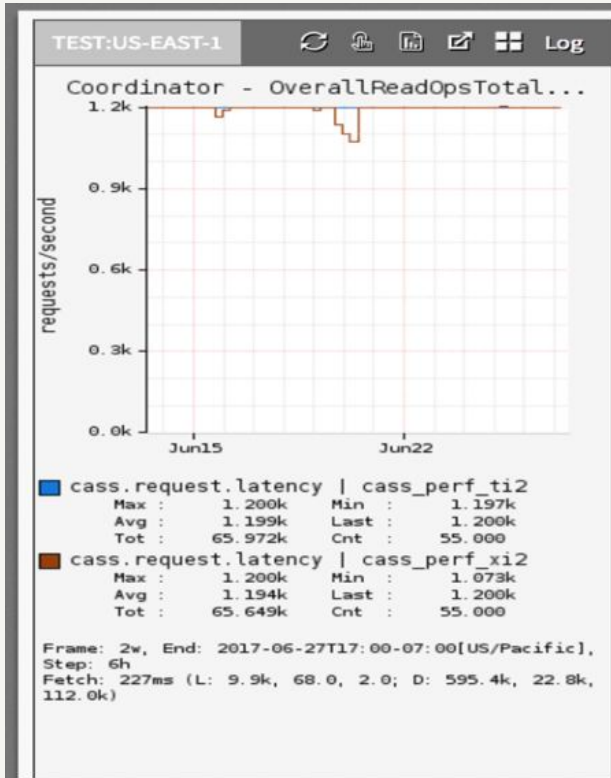


NETFLIX

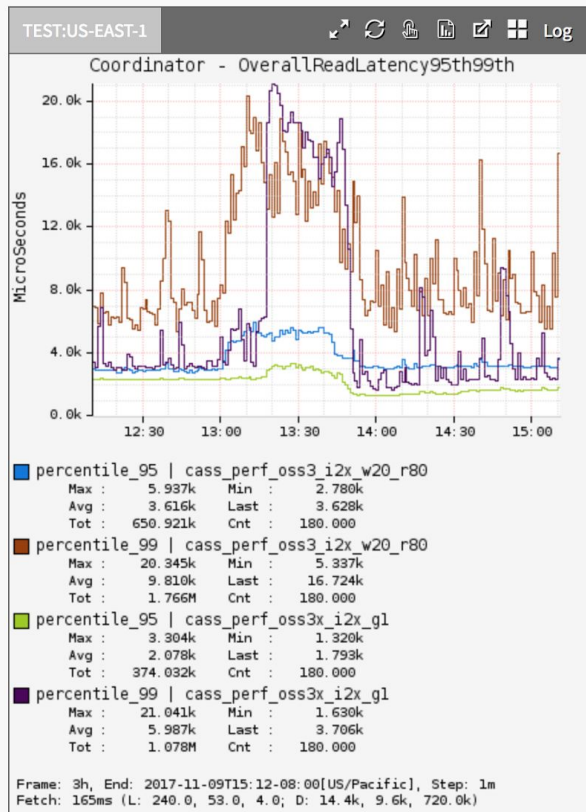


i2 vs i3 instance type



SSD Tuning



CMS vs G1GC for Cassandra



AMI Certification Process

BUILD #129

TRIGGERED BUILD
cdeamicass48
8 days ago
pipelineengine: "cass48"

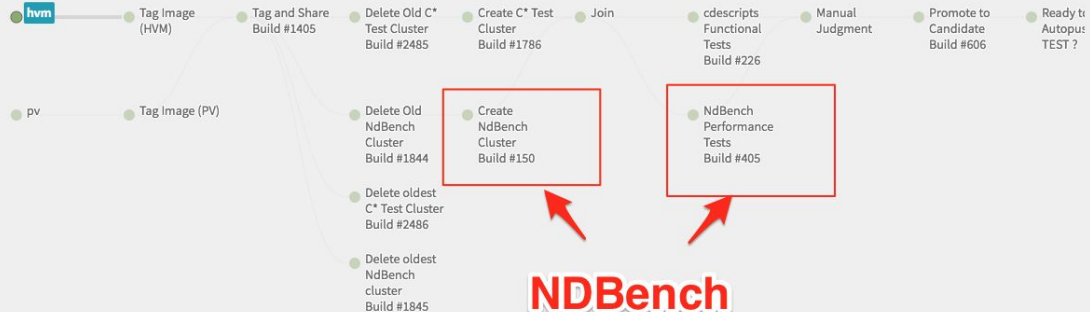
32:15 31:43 01:51 02:19 02:40 09:17 07:44 00:54 00:12 26:06 21:39 00:15 02:16 01:34 00:11 00:11 01:01 06:26 25:25

00:15

Status: **SUCCEEDED** Restarted by [redacted] 2017-10-31 09:59:18 PDT

Duration: 33:33

Details



NDBench

STAGE DETAILS: HVM

Duration: 32:15

| Step | Started | Duration | Status |
|-------------------|-------------------------|----------|-----------|
| Bake in us-east-1 | 2017-10-26 13:55:52 PDT | | SUCCEEDED |
| Bake in us-west-1 | 2017-10-26 13:55:52 PDT | | SUCCEEDED |
| Bake in us-west-2 | 2017-10-26 13:55:52 PDT | | SUCCEEDED |
| Bake in eu-west-1 | 2017-10-26 13:55:52 PDT | | SUCCEEDED |
| Hvm | 2017-10-26 13:55:51 PDT | | SUCCEEDED |

Comments

Pinning it to unstable version to get user account fix, Please change it back to stable

BAKE IN US-EAST-1

Bake Config

Task Status

Provider Amazon
Image [redacted]
Region us-east-1
Package cdeamicass48
Base OS trusty
VM Type HVM
Store Type S3
Label release

Image:

[redacted]:86_64-20171026205551-
trusty-hvm-sriov-s3
(ami-177bdd6d)

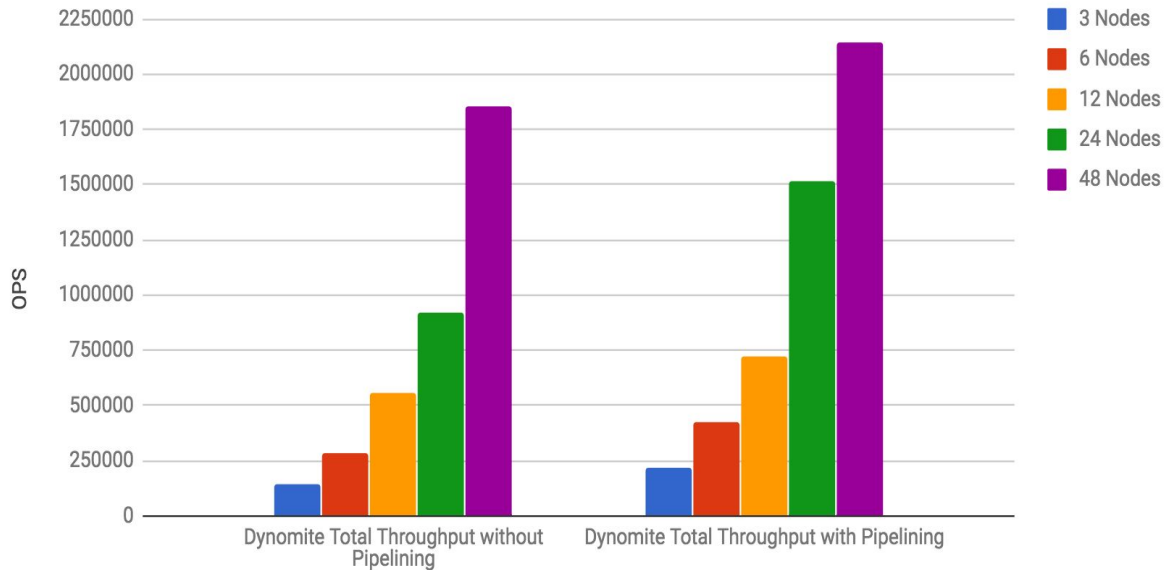
[View Bakery Details](#)

[Source](#) | [Permalink](#)

Throughput Scalability



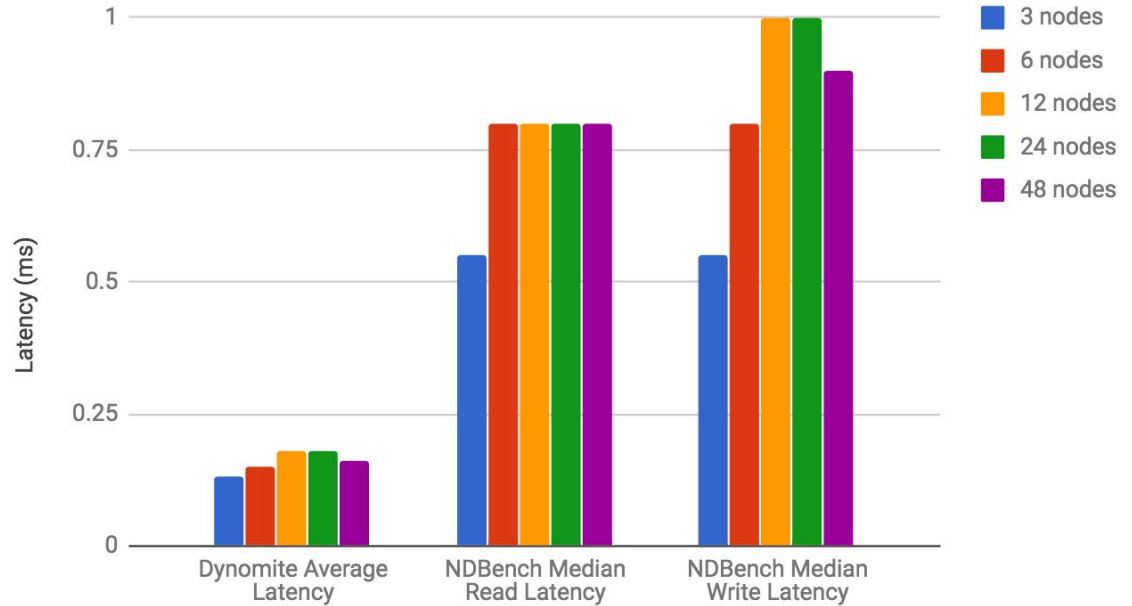
NDBench Generated Traffic (DC_ONE mode)



Latency



Dynomite & NDBench Latency



NDBench @ Netflix as ...

- Benchmarking Tool
- Integration Tests
- Deployment Validation

Auto-tuning

- Find the appropriate capacity that the cluster can sustain
 - RPS/WPS
 - Number of documents indexed/sec
- After X% of SLA violations, NDBench stops stepping up the load
 - Exponential backoff





NDBENCH

Github: <https://github.com/netflix/ndbench>

Talk to us: <https://gitter.im/Netflix/ndbench>

NETFLIX

Take away

*“Benchmark your micro services and data stores in
Cloud ecosystem at **Scale** using **NDBench**”*

Q & A

Vinay Chella - [@vinaykchella](#)

Ioannis Papapanagiotou - [@ipapapa](#)



NETFLIX

NETFLIX

Architecture

- **Core:** Workload generator
- **API:** Adding plugins
- **Web:** UI and the servlets

