# Benchmarking Suite: Update

C. Cordeiro, D. Giordano
CERN-IT

*pre-GDB Benchmarking*
February 2017

# This is not a performance talk

> *Performance is a key criterion in the design, procurement, and use of computer systems [...] to get the highest performance for a given cost.* [1]

- Performance Measurement is essential
  - Deal with the intrinsic variability and inhomogeneity
  - Compare the presumed and perceived performance
  - Identify performance issues

- Standard procedure during the procurement process

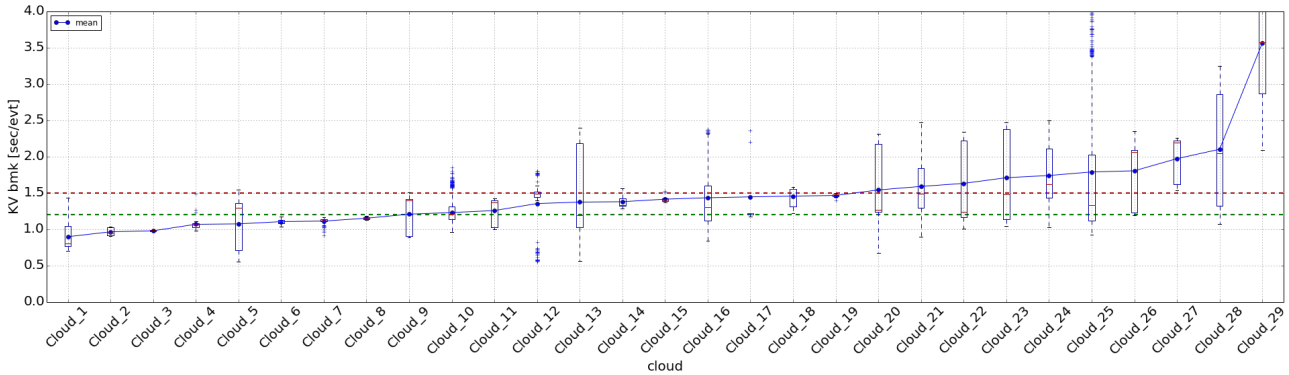


[1] *Art of Computer Systems Performance Analysis Techniques For Experimental Design Measurements Simulation And Modeling* by Raj Jain , Wiley Computer Publishing, John Wiley & Sons, Inc


**AUTOMATE**

**EXECUTION**

&

**EASE SYSTEMATIC**

**DATA COLLECTION**

# Benchmarking Model

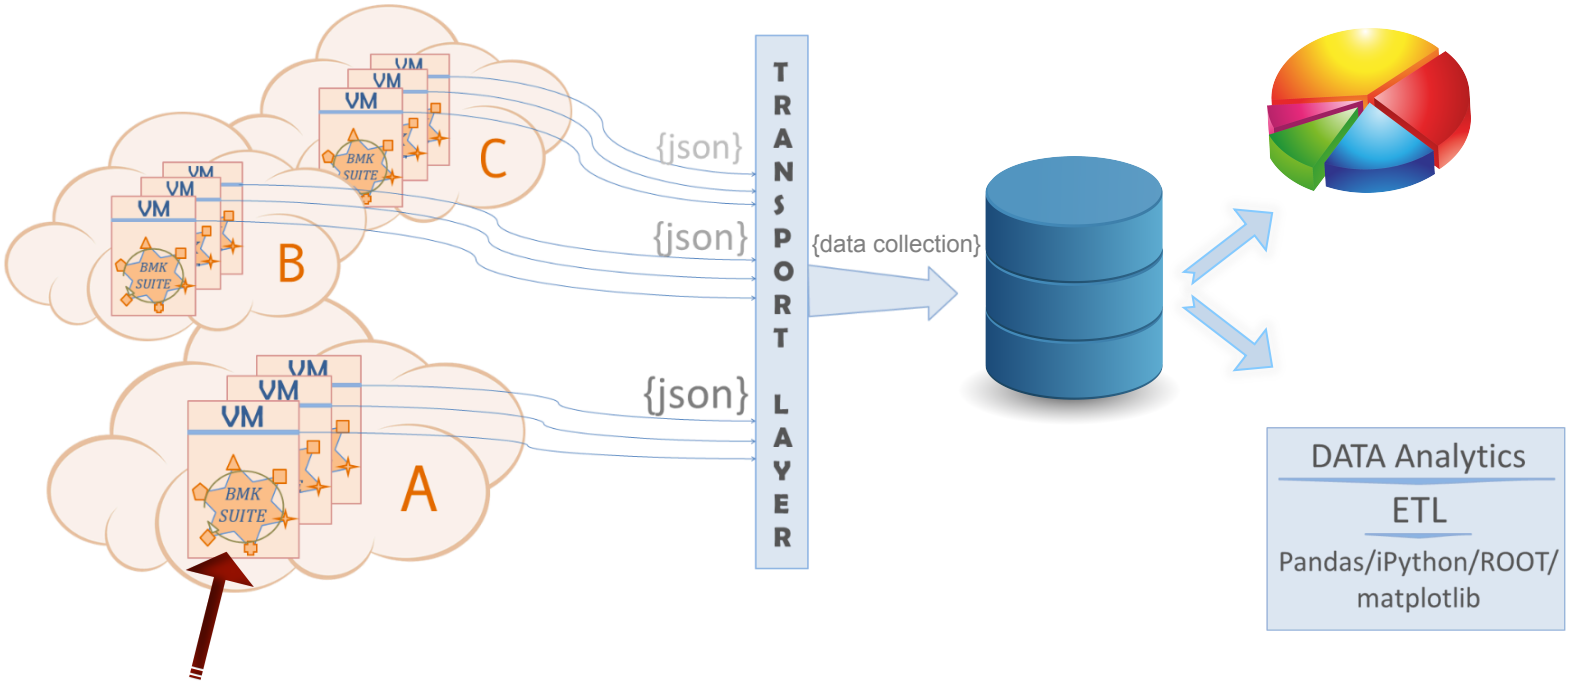


* *a toolkit done for internal usage. Adoption also by stakeholders of the HEPiX Benchmark working group*

# Benchmark Suite

CONFIGURABLE SEQUENCE OF BENCHMARKS

RUN IN ANY IAAS

AUTOMATE AND GENERALIZE BENCHMARK EXECUTION

COMMON OUTPUT STRUCTURE

PUBLISH RESULTS

# This is a toolkit, not a service

# Benchmarking Suite

## How it works

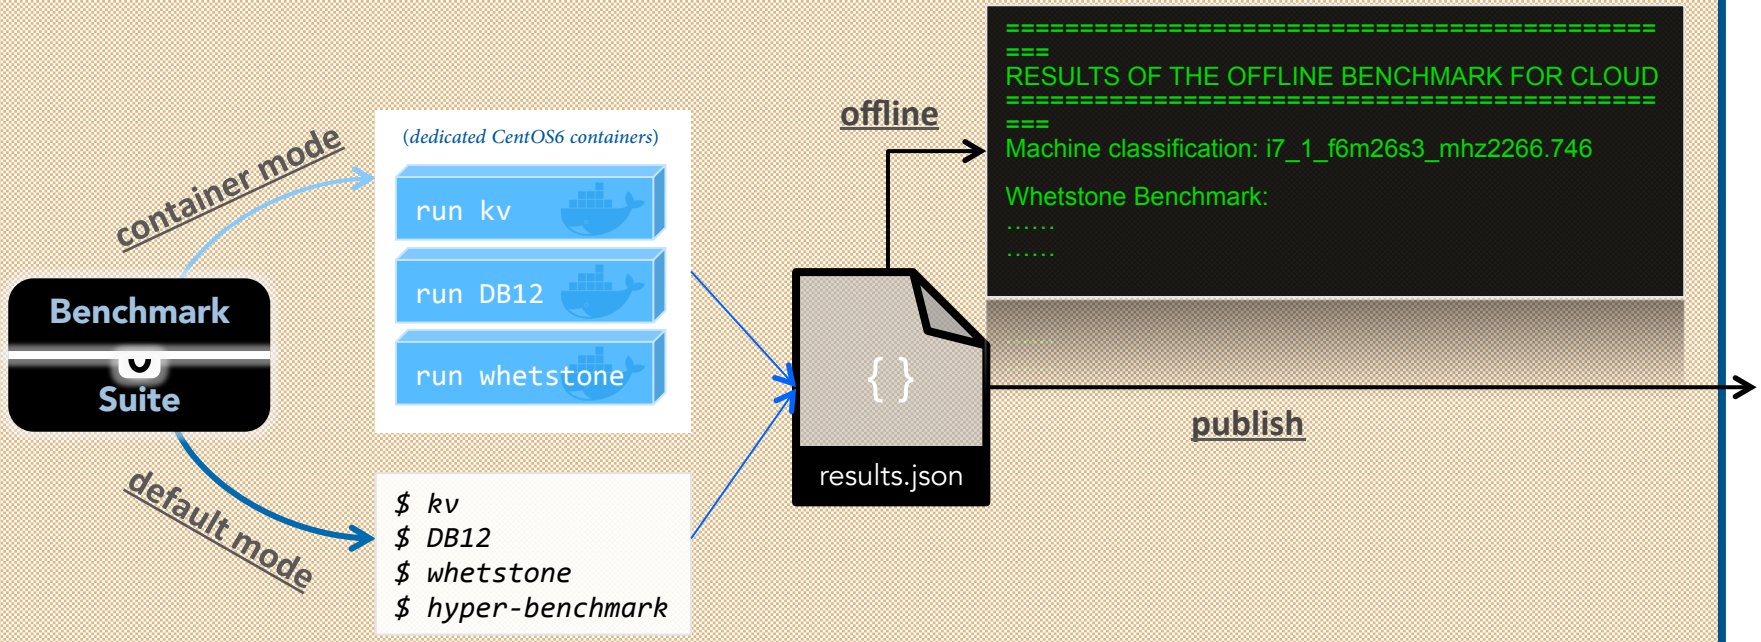

```
Usage:
 $0 [OPTIONS]

OPTIONS:

 -q
           Quiet mode. Do not prompt user
 -o
           Offline mode. Do not publish results
 -i
           Solves/checks the general and unique dependencies for the specified --benchmarks.
           If not used, assumes all the dependencies are already installed and configured. NOTE: should run as root
 --benchmarks=<bmk1;bmk2>
           Semi-colon separated list of benchmarks to run. Available benchmarks are:
                     - kv
                     - whetstone
                     - DB12
                     - hyper-benchmark (*)
 --mp_num=#
           Number of concurrent processes (usually cores) to run
 --kv_xml=<xmlFile>
           Input file for the KV benchmark. If not provided, SingleMuonGenerator is default
 --uid=<id>
           Unique identifier for the host running this script
 --public_ip=<ip>
           Public IP address of the host running this script
 --cloud=<cloudName>
           Cloud name to identify the results - if not specified, CLOUD=test
 --vo=<VO>
           Name of the VO responsible for the underlying resource
 --pnode=<physicalNode>
           (Optional) Name of the hypervisor machine hosting the VM
 --queue_port=<portNumber>
           Port number of the ActiveMQ broker where to send the benchmarking results
 --queue_host=<hostname>
           Hostname with the ActiveMQ broker where to send the benchmarking results
 --username=<username>
           Username to access the ActiveMQ broker where to send the benchmarking results
 --password=<password>
           User password to access ActiveMQ broker where to send the benchmarking results
 --amq_key=<path_to_key>
           Key file for the AMQ authentication, without passphrase. Expects --amq_cert
 --amq_cert=<path_to_cert>
           Certificate for the AMQ authentication. Expects --amq_key
 --topic=<topicName>
           Topic (or Queue) name used in the ActiveMQ broker
 --freetext=<string>
           Any additional free text to add to the output JSON
```

```
Usage:
 $0 [OPTIONS]

OPTIONS:

 -q
        Quiet mode. Do not prompt user
 -o
        Offline mode. Do not publish results
 -i
        Solves/checks the general and unique dependencies for the specified --benchmarks.
        If not used, assumes all the dependencies are already installed and configured. NOTE: should run as root
--benchmarks=<bmk1;bmk2>
          Semi-colon separated list of benchmarks to run. Available benchmarks are:
                    - kv
                    - whetstone
                    - DB12
                    - hyper-benchmark (*)
--mp_num=#
          Number of concurrent processes (usually cores) to run
--kv_xml=<xmlFile>
          Input file for the KV benchmark. If not provided, SingleMuonGenerator is default
--uid=<id>
          Unique identifier for the host running this script
--public_ip=<ip>
          Public IP address of the host running this script
--cloud=<cloudName>
          Cloud name to identify the results - if not specified, CLOUD=test
--vo=<VO>
          Name of the VO responsible for the underlying resource
--pnode=<physicalNode>
          (Optional) Name of the hypervisor machine hosting the VM
--queue_port=<portNumber>
          Port number of the ActiveMQ broker where to send the benchmarking results
--queue_host=<hostname>
          Hostname with the ActiveMQ broker where to send the benchmarking results
--username=<username>
          Username to access the ActiveMQ broker where to send the benchmarking results
--password=<password>
          User password to access ActiveMQ broker where to send the benchmarking results
--amq_key=<path_to_key>
          Key file for the AMQ authentication, without passphrase. Expects --amq_cert
--amq_cert=<path_to_cert>
          Certificate for the AMQ authentication. Expects --amq_key
--topic=<topicName>
          Topic (or Queue) name used in the ActiveMQ broker
--freetext=<string>
          Any additional free text to add to the output JSON
```

```
Usage:
 $0 [OPTIONS]

OPTIONS:

 -q
          Quiet mode. Do not prompt user
 -o
          Offline mode. Do not publish results
 -i
          Solves/checks the general and unique dependencies for the specified --benchmarks.
          If not used, assumes all the dependencies are already installed and configured. NOTE: should run as root
```

**--benchmarks=<bmk1;bmk2>**

**Semi-colon separated list of benchmarks to run. Available benchmarks are:**

- **kv**
- **whetstone**
- **DB12**
- **hyper-benchmark (*)**

```
--mp_num=#
          Number of concurrent processes (usually cores) to run
--kv_xml=<xmlFile>
          Input file for the KV benchmark. If not provided, SingleMuonGenerator is default
--uid=<id>
          Unique identifier for the host running this script
--public_ip=<ip>
          Public IP address of the host running this script
--cloud=<cloudName>
          Cloud name to identify the results - if not specified, CLOUD=test
--vo=<VO>
          Name of the VO responsible for the underlying resource
--pnode=<physicalNode>
          (Optional) Name of the hypervisor machine hosting the VM
--queue_port=<portNumber>
          Port number of the ActiveMQ broker where to send the benchmarking results
--queue_host=<hostname>
          Hostname with the ActiveMQ broker where to send the benchmarking results
--username=<username>
          Username to access the ActiveMQ broker where to send the benchmarking results
--password=<password>
          User password to access ActiveMQ broker where to send the benchmarking results
--amq_key=<path_to_key>
          Key file for the AMQ authentication, without passphrase. Expects --amq_cert
--amq_cert=<path_to_cert>
          Certificate for the AMQ authentication. Expects --amq_key
--topic=<topicName>
          Topic (or Queue) name used in the ActiveMQ broker
--freetext=<string>
          Any additional free text to add to the output JSON
```

```
Usage:
 $0 [OPTIONS]

OPTIONS:

 -q
          Quiet mode. Do not prompt user
 -o
          Offline mode. Do not publish results
 -i
          Solves/checks the general and unique dependencies for the specified --benchmarks.
          If not used, assumes all the dependencies are already installed and configured. NOTE: should run as root
 --benchmarks=<bmk1;bmk2>
          Semi-colon separated list of benchmarks to run. Available benchmarks are:
                    - kv
                    - whetstone
                    - DB12
                    - hyper-benchmark (*)
```

**--mp_num=#**

**Number of concurrent processes (usually cores) to run**

```
 --kv_xml=<xmlFile>
          Input file for the KV benchmark. If not provided, SingleMuonGenerator is default
 --uid=<id>
          Unique identifier for the host running this script
 --public_ip=<ip>
          Public IP address of the host running this script
 --cloud=<cloudName>
          Cloud name to identify the results - if not specified, CLOUD=test
 --vo=<VO>
          Name of the VO responsible for the underlying resource
 --pnode=<physicalNode>
          (Optional) Name of the hypervisor machine hosting the VM
 --queue_port=<portNumber>
          Port number of the ActiveMQ broker where to send the benchmarking results
 --queue_host=<hostname>
          Hostname with the ActiveMQ broker where to send the benchmarking results
 --username=<username>
          Username to access the ActiveMQ broker where to send the benchmarking results
 --password=<password>
          User password to access ActiveMQ broker where to send the benchmarking results
 --amq_key=<path_to_key>
          Key file for the AMQ authentication, without passphrase. Expects --amq_cert
 --amq_cert=<path_to_cert>
          Certificate for the AMQ authentication. Expects --amq_key
 --topic=<topicName>
          Topic (or Queue) name used in the ActiveMQ broker
 --freetext=<string>
          Any additional free text to add to the output JSON
```

```
Usage:
 $0 [OPTIONS]

OPTIONS:

 -q
          Quiet mode. Do not prompt user
 -o
          Offline mode. Do not publish results
 -i
          Solves/checks the general and unique dependencies for the specified --benchmarks.
          If not used, assumes all the dependencies are already installed and configured. NOTE: should run as root
 --benchmarks=<bmk1;bmk2>
          Semi-colon separated list of benchmarks to run. Available benchmarks are:
                    - kv
                    - whetstone
                    - DB12
                    - hyper-benchmark (*)
 --mp_num=#
          Number of concurrent processes (usually cores) to run
```

**--kv_xml=<xmlFile>**

**Input file for the KV benchmark. If not provided, SingleMuonGenerator is default**

```
 --uid=<id>
          Unique identifier for the host running this script
 --public_ip=<ip>
          Public IP address of the host running this script
 --cloud=<cloudName>
          Cloud name to identify the results - if not specified, CLOUD=test
 --vo=<VO>
          Name of the VO responsible for the underlying resource
 --pnode=<physicalNode>
          (Optional) Name of the hypervisor machine hosting the VM
 --queue_port=<portNumber>
          Port number of the ActiveMQ broker where to send the benchmarking results
 --queue_host=<hostname>
          Hostname with the ActiveMQ broker where to send the benchmarking results
 --username=<username>
          Username to access the ActiveMQ broker where to send the benchmarking results
 --password=<password>
          User password to access ActiveMQ broker where to send the benchmarking results
 --amq_key=<path_to_key>
          Key file for the AMQ authentication, without passphrase. Expects --amq_cert
 --amq_cert=<path_to_cert>
          Certificate for the AMQ authentication. Expects --amq_key
 --topic=<topicName>
          Topic (or Queue) name used in the ActiveMQ broker
 --freetext=<string>
          Any additional free text to add to the output JSON
```

```
Usage:
 $0 [OPTIONS]

OPTIONS:

 -q
          Quiet mode. Do not prompt user
 -o
          Offline mode. Do not publish results
 -i
          Solves/checks the general and unique dependencies for the specified --benchmarks.
          If not used, assumes all the dependencies are already installed and configured. NOTE: should run as root
 --benchmarks=<bmk1;bmk2>
          Semi-colon separated list of benchmarks to run. Available benchmarks are:
                    - kv
                    - whetstone
                    - DB12
                    - hyper-benchmark (*)
 --mp_num=#
          Number of concurrent processes (usually cores) to run
 --kv_xml=<xmlFile>
          Input file for the KV benchmark. If not provided, SingleMuonGenerator is default
 --uid=<id>
       Unique identifier for the host running this script
 --public_ip=<ip>
       Public IP address of the host running this script
 --cloud=<cloudName>
       Cloud name to identify the results - if not specified, CLOUD=test
 --vo=<VO>
       Name of the VO responsible for the underlying resource
 --pnode=<physicalNode>
       (Optional) Name of the hypervisor machine hosting the VM
 --queue_port=<portNumber>
          Port number of the ActiveMQ broker where to send the benchmarking results
 --queue_host=<hostname>
          Hostname with the ActiveMQ broker where to send the benchmarking results
 --username=<username>
          Username to access the ActiveMQ broker where to send the benchmarking results
 --password=<password>
          User password to access ActiveMQ broker where to send the benchmarking results
 --amq_key=<path_to_key>
          Key file for the AMQ authentication, without passphrase. Expects --amq_cert
 --amq_cert=<path_to_cert>
          Certificate for the AMQ authentication. Expects --amq_key
 --topic=<topicName>
          Topic (or Queue) name used in the ActiveMQ broker
 --freetext=<string>
          Any additional free text to add to the output JSON
```

```
Usage:
 $0 [OPTIONS]

OPTIONS:

 -q
          Quiet mode. Do not prompt user
 -o
          Offline mode. Do not publish results
 -i
          Solves/checks the general and unique dependencies for the specified --benchmarks.
          If not used, assumes all the dependencies are already installed and configured. NOTE: should run as root
 --benchmarks=<bmk1;bmk2>
          Semi-colon separated list of benchmarks to run. Available benchmarks are:
                    - kv
                    - whetstone
                    - DB12
                    - hyper-benchmark (*)
 --mp_num=#
          Number of concurrent processes (usually cores) to run
 --kv_xml=<xmlFile>
          Input file for the KV benchmark. If not provided, SingleMuonGenerator is default
 --uid=<id>
          Unique identifier for the host running this script
 --public_ip=<ip>
          Public IP address of the host running this script
 --cloud=<cloudName>
          Cloud name to identify the results - if not specified, CLOUD=test
 --vo=<VO>
          Name of the VO responsible for the underlying resource
 --pnode=<physicalNode>
          (Optional) Name of the hypervisor machine hosting the VM
 --queue_port=<portNumber>
          Port number of the ActiveMQ broker where to send the benchmarking results
 --queue_host=<hostname>
          Hostname with the ActiveMQ broker where to send the benchmarking results
 --username=<username>
          Username to access the ActiveMQ broker where to send the benchmarking results
 --password=<password>
          User password to access ActiveMQ broker where to send the benchmarking results
 --amq_key=<path_to_key>
          Key file for the AMQ authentication, without passphrase. Expects --amq_cert
 --amq_cert=<path_to_cert>
          Certificate for the AMQ authentication. Expects --amq_key
 --topic=<topicName>
          Topic (or Queue) name used in the ActiveMQ broker
 --message=<string>
          Any additional free text to add to the output JSON
```

```
Usage:
 $0 [OPTIONS]

OPTIONS:

 -q
          Quiet mode. Do not prompt user
 -o
          Offline mode. Do not publish results
 -i
          Solves/checks the general and unique dependencies for the specified --benchmarks.
          If not used, assumes all the dependencies are already installed and configured. NOTE: should run as root
 --benchmarks=<bmk1;bmk2>
          Semi-colon separated list of benchmarks to run. Available benchmarks are:
                    - kv
                    - whetstone
                    - DB12
                    - hyper-benchmark (*)
 --mp_num=#
          Number of concurrent processes (usually cores) to run
 --kv_xml=<xmlFile>
          Input file for the KV benchmark. If not provided, SingleMuonGenerator is default
 --uid=<id>
          Unique identifier for the host running this script
 --public_ip=<ip>
          Public IP address of the host running this script
 --cloud=<cloudName>
          Cloud name to identify the results - if not specified, CLOUD=test
 --vo=<VO>
          Name of the VO responsible for the underlying resource
 --pnode=<physicalNode>
          (Optional) Name of the hypervisor machine hosting the VM
 --queue_port=<portNumber>
          Port number of the ActiveMQ broker where to send the benchmarking results
 --queue_host=<hostname>
          Hostname with the ActiveMQ broker where to send the benchmarking results
 --username=<username>
          Username to access the ActiveMQ broker where to send the benchmarking results
 --password=<password>
          User password to access ActiveMQ broker where to send the benchmarking results
 --amq_key=<path_to_key>
          Key file for the AMQ authentication, without passphrase. Expects --amq_cert
 --amq_cert=<path_to_cert>
          Certificate for the AMQ authentication. Expects --amq_key
 --topic=<topicName>
          Topic (or Queue) name used in the ActiveMQ broker
```

**--freetext=<string>**

**Any additional free text to add to the output JSON**

# Benchmarking Suite

## Running examples*


running Whetstone quickly, offline

```
cern-benchmark --benchmarks="whetstone" -q -o   # CLOUD name is set to test by default
```

running ATLAS Kit Validation, DIRAC Benchmark and Whetstone, and publish to AMQ with user:pass authentication

```
cern-benchmark --benchmarks="kv;DB12;whetstone" --cloud=CERN --queue_host=AMQ_MB_SERVER.domain
--queue_port=PORT_NUMBER --username=yourUser --password=`cat /pwdfile` --topic=topicORqueueName
```

pushing to AMQ using certificate authentication

```
cern-benchmark --benchmarks="whetstone" -q --cloud=test --vo=test --freetext="test"
--queue_host=AMQ_MB_SERVER.domain --queue_port=PORT_NUMBER_SSL --username=yourUser
--amq_key=/yourkey.key --amq_cert=/yourcert.crt --topic=topicORqueueName
```

* for more example please read the documentation at http://bmkwg.web.cern.ch/bmkwg/docs/HowToRun.html

# Benchmarking Suite in Action: examples

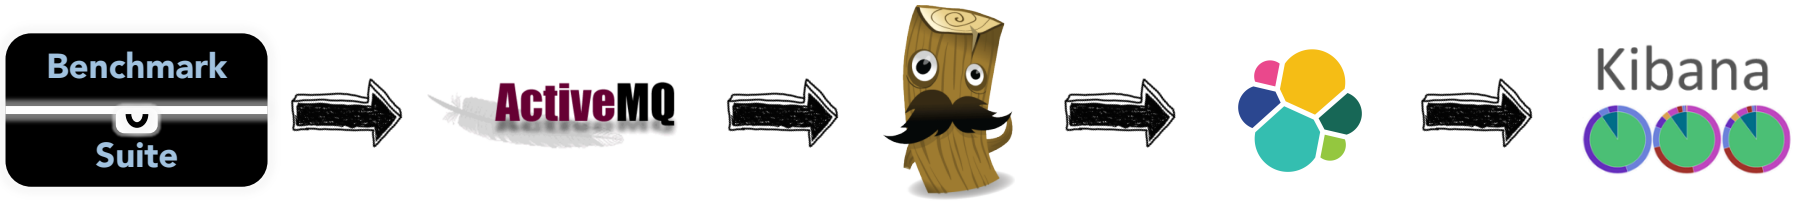


- For direct providers' performance comparison

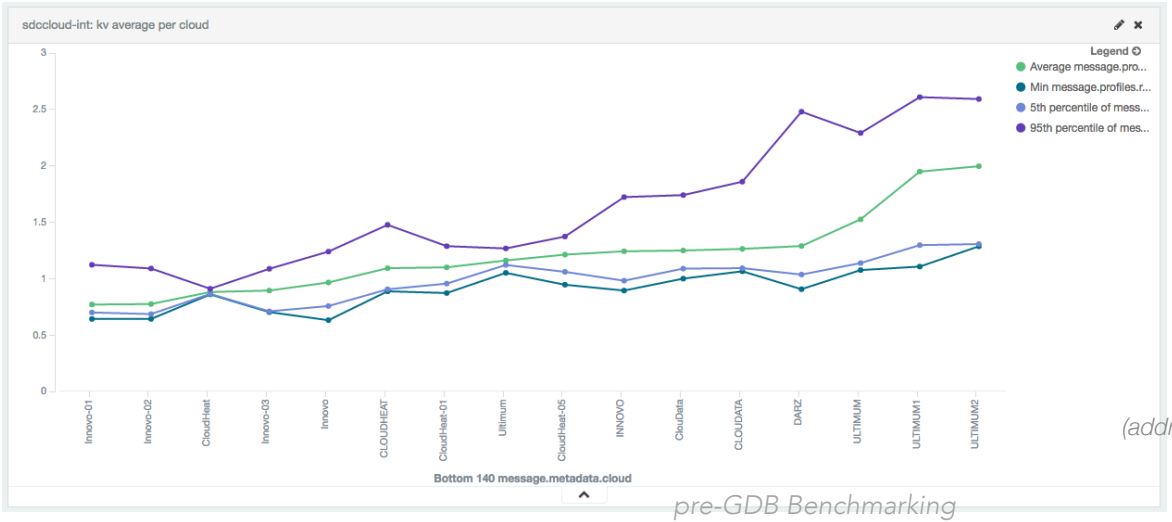


*(addressed by D. Giordano in the previous talk)*

# Benchmarking Suite in Action: examples


- For scrutiny analysis of new hardware


```
{
   "_index":"sdccloud-int",
   "_type":"vmspec",
   "_id":"the240s-901bfaac-d358-405a-b943-0a4accf56d24_7b9cc47e-6b51-48cd-94c2e3943856d941_2016-07-21T15:05:02Z",
   "_score":null,
   "_source":{
      "message":{
         "_timestamp":"2016-07-21T15:05:02Z",
         "_id":"the240s-901bfaac-d358-405a-b943-0a4accf56d24_7b9cc47e-6b51-48cd-94c2e3943856d941_2016-07-21T15:05:02Z",
         "profiles":{
            "fastBmk":{
               "value":"8.01781737194"
            }
         },
         "metadata":{
            "benchmark_target":"machine",
            "meminfo":1876284,
            "UID":"the240s-901bfaac-d358-405a-b943-0a4accf56d24_7b9cc47e-6b51-48cd-94c2-e3943856d941",
            "classification":"i6_1_f6m63s2_mhz2394.452",
            "freetext":"VM-1",
            "cpuname":"Intel(R) Xeon(R) CPU E5-2630 v3 @ 2.40GHz",
            "ip":"188.185.184.233",
            "osdist":"Scientific Linux CERN SLC release 6.8 (Carbon)",
            "bogomips":4788.9,
            "VO":null,
            "cpunum":1,
            "pnode":"P06253971G07182",
            "mp_num":"1",
            "pyver":"2.6.6",
            "cloud":"CERN-wig_project_011"
```


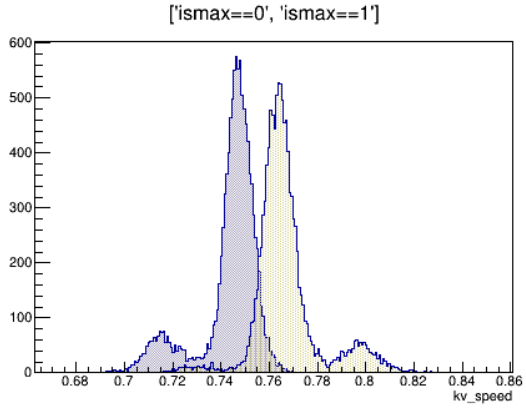


*(addressed in the previous talks)*

# BENCHMARKING SUITE BIOGRAPHY

**4M** benchmarks collected since August 6th 2015 17:57:49

**16+** data centers (4 grid sites and over 12 clouds)

**125** different CPU architectures

**16** different releases from 5 different OS

**>26k** unique IPs


**16k** benchmarks collected from BOINC


**Pilot users**. Testing and providing the code at /cvmfs/atlas.cern.ch/repo/benchmarks/cern/current


**Benchmark Working Group**
Steering, testing and coordinating the activity

# Highlights

- Additional benchmarks can be easily plugged in
  - Same for publishing support for other transportation layers
- Need to refine packaging and distribution strategy
  - "Grid" adopters prefer to find it in a common CVMFS area
    - *KV needs atlas.cern.ch though…*
- Prototyped dependency-free and isolated benchmarks with containers

- Wider toolkit adoption ➔ more manpower
  - Need support from HEPiX community...

# Temporary Brochure

## DOCUMENTATION

http://bmkwg.web.cern.ch/bmkwg/

## CODE

https://gitlab.cern.ch/cloud-infrastructure/cloud-benchmark-suite

## HEPiX BENCHMARKING WORKING GROUP

hepix-cpu-benchmark@HEPIX.ORG

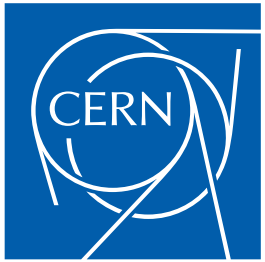
CERN