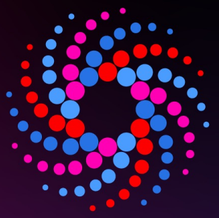

SC22
Dallas, TX | hpc accelerates.

pyp2pcluster: A cluster discovery tool

Mr Robert Tracey

Dr Mobayode O. Akinsolu

Prof Vadim Elisseev

Dr Sultan Shoaib

IBM Research Europe

Robert Tracey
pyp2pcluster: A cluster discovery tool

@wob86

Dr Mobayode O. Akinsolu
Prof Vadim Elisseev
Dr Sultan Shoaib

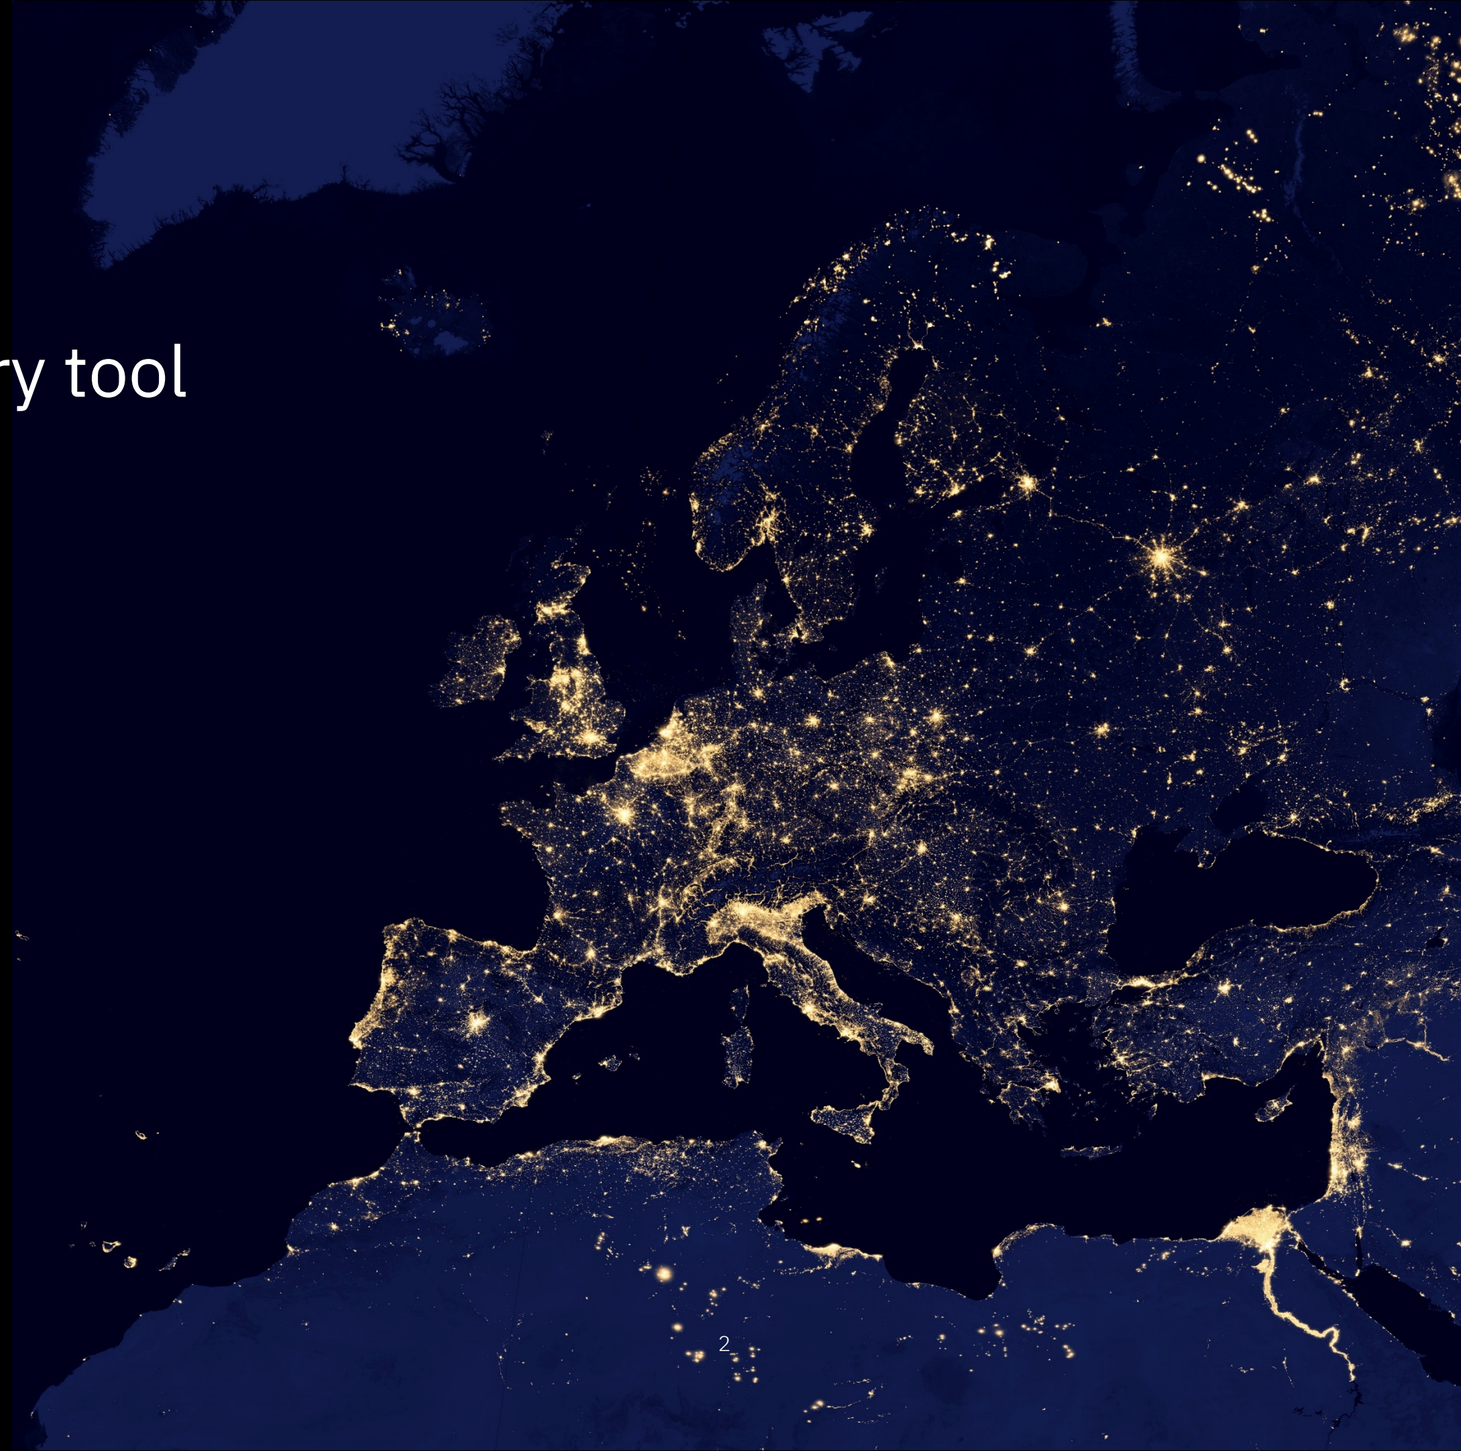

© 2022 IBM Corporation

Outline

- Motivation
- Requirements
- Results
- Next Steps

© 2022 IBM Corporation

Motivation

Phd: Dynamic routing of data based on resource availability and benefit

What is the best cluster or single node (compute resources) to run your job on?

Data needed from compute resources to be sent to fuzzy logic

Compute resources are then classified on how ‘good’ they are

This is then passed to swarm algorithms to find the ‘best’ cluster to use

© 2022 IBM Corporation

Library requirements

Easily adaptable for our project

Be able to gather information from different types of clusters

Be able to work with single nodes

To consumes as few resources as possible

Able to obtain network information between resources

Still be being supported

© 2022 IBM Corporation

Why P2P?

No central system needed

Increased resilience

Doesn’t reserve bandwidth

Can send and receive easily between other nodes

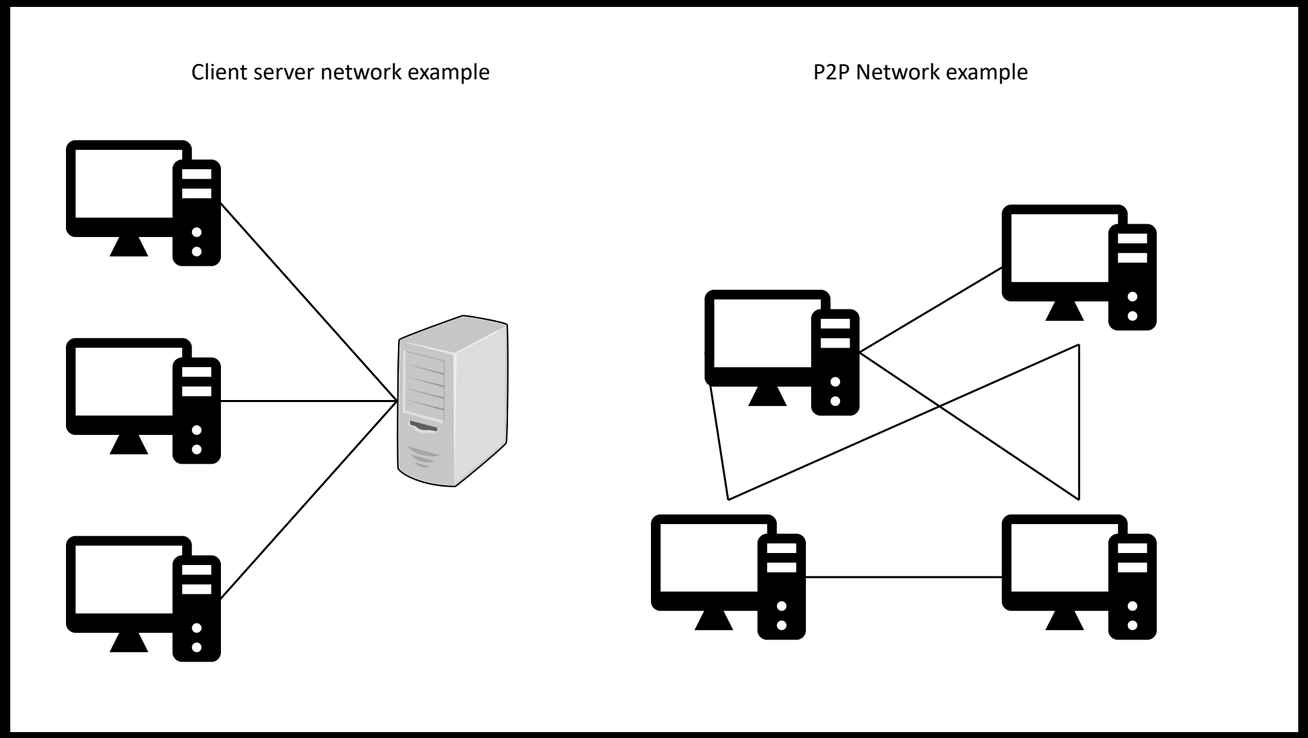

© 2022 IBM Corporation

Overview of communication

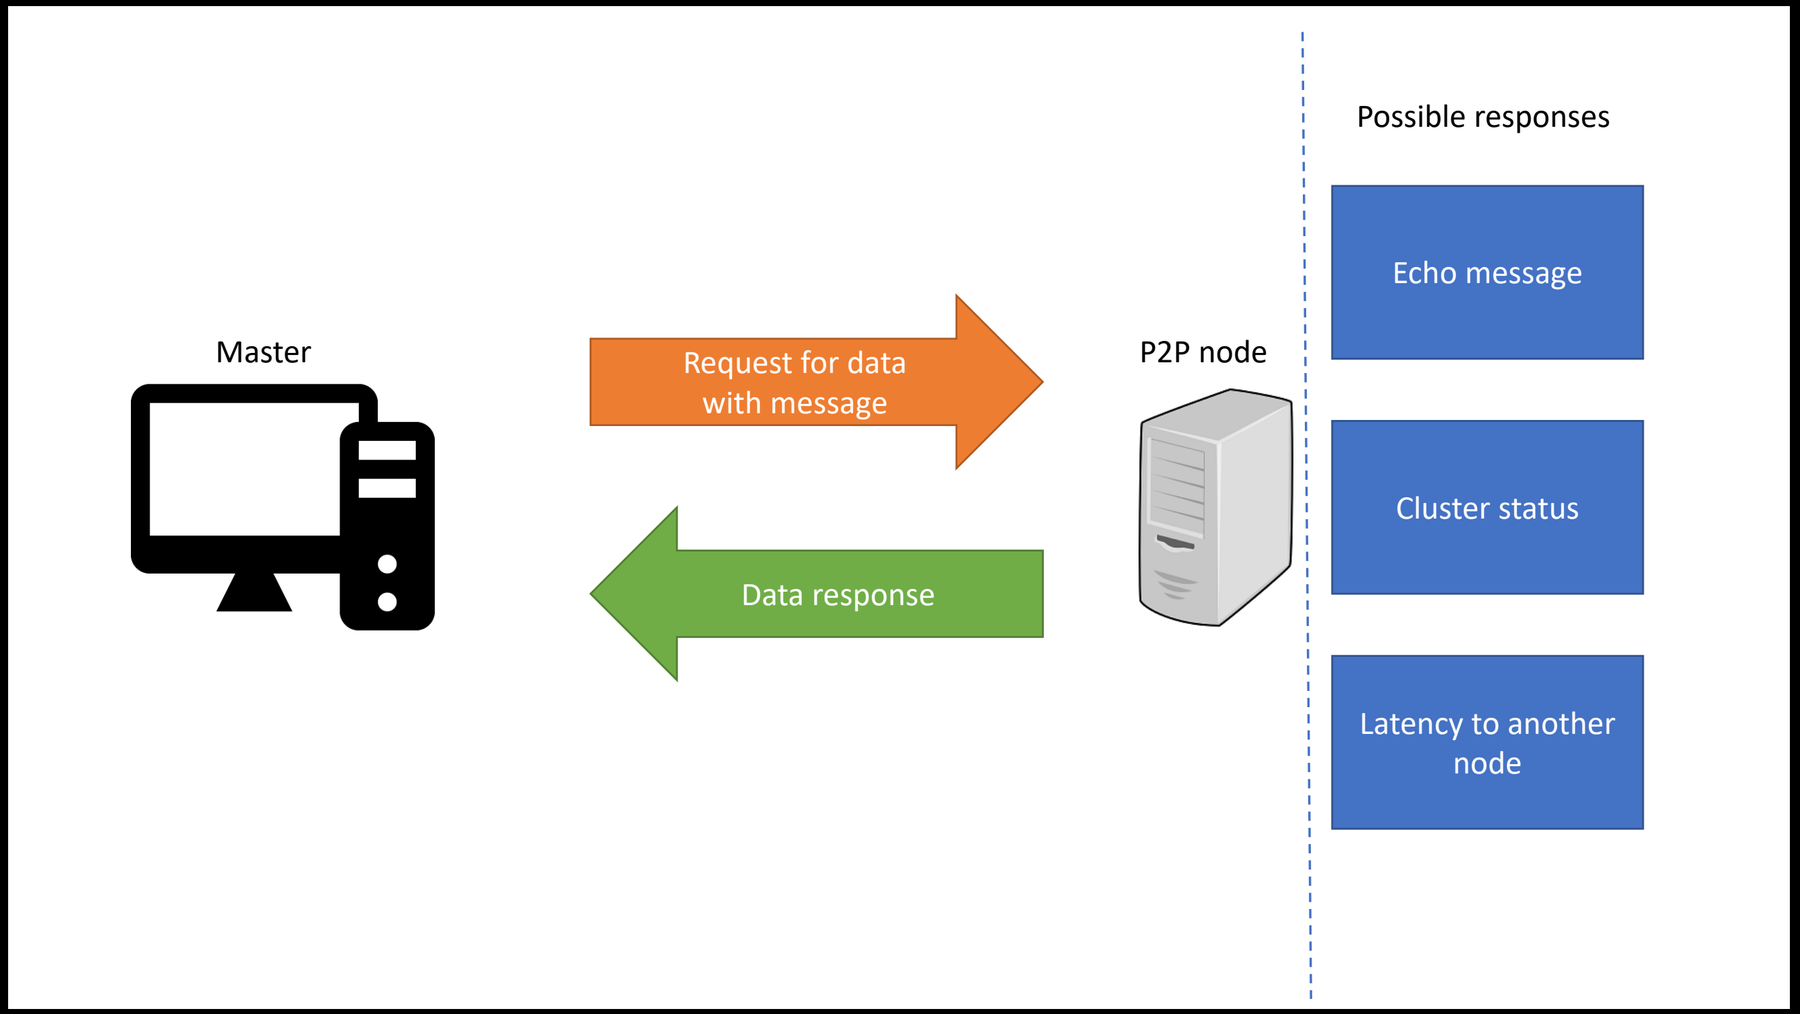

© 2022 IBM Corporation

Examples of usage

Node creation

```
from pyp2pcluster import pyp2pconn as pyp2p
import sys

port = int(sys.argv[2])
sysid = sys.argv[1]

loop = pyp2p(port,sysid)
loop.mainloop()
```

© 2022 IBM Corporation

Examples of usage

Gathering data from LSF

```
import socket
import pickle
def clustertest(host, port, cluster,mode):
    message = ["getcluster",cluster,mode]
    with socket.socket(socket.AF_INET, \
     socket.SOCK_STREAM) as s:
        s.connect((host, port))
        messageb = pickle.dumps(message)
        s.sendall(messageb)
        data = s.recv(1024)
        outy = pickle.loads(data)
    return outy
clustertest("192.168.0.125", 44444,"lsf","table")
```

© 2022 IBM Corporation

Cluster results

Hostname	CPU	Mem	GPU
node1	1	196	0
node2	1	199	0
node3	1	199	0
node4	1	200	0

Hostname	CPU	Mem	GPU
loginnode	128	64G	0
host1	32	1.1T	4
host2	32	1.1T	4
host3	32	1.1T	4
host4	20	1.1T	4
host5	20	512G	4
host6	20	512G	4
host7	20	512G	4
host8	20	512G	4
host9	20	1T	0
host10	20	1T	0
host11	20	1T	0
host12	20	1T	0
host13	20	1T	0
host14	20	1T	0
host15	20	1T	0
host16	20	1T	0
host17	20	1T	0
host18	20	1T	0
host19	20	1T	0
host20	20	1T	0

© 2022 IBM Corporation

Latency results

Source	Destination	Latency(ms)
master	flux	2.800941
master	nova	3.694057
master	sandbox	1.065969
flux	nova	2.099037
flux	dtfi4	11.276484
nova	dtfi4	11.115551

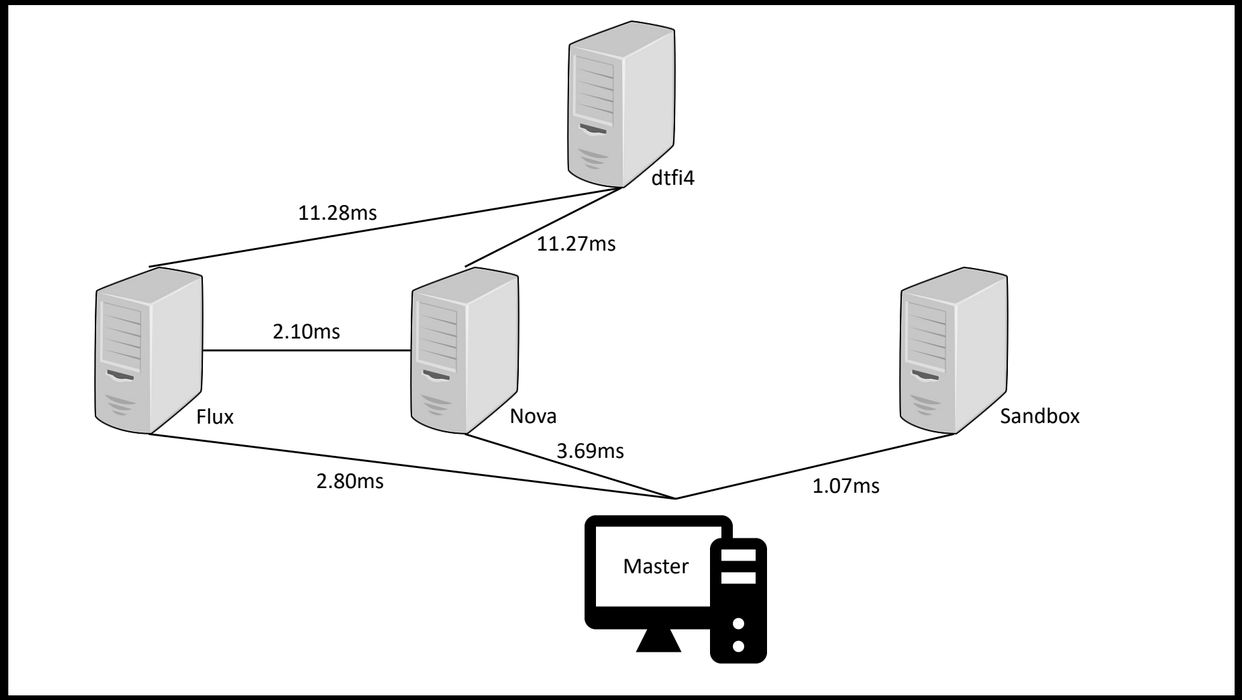

© 2022 IBM Corporation

Next Steps

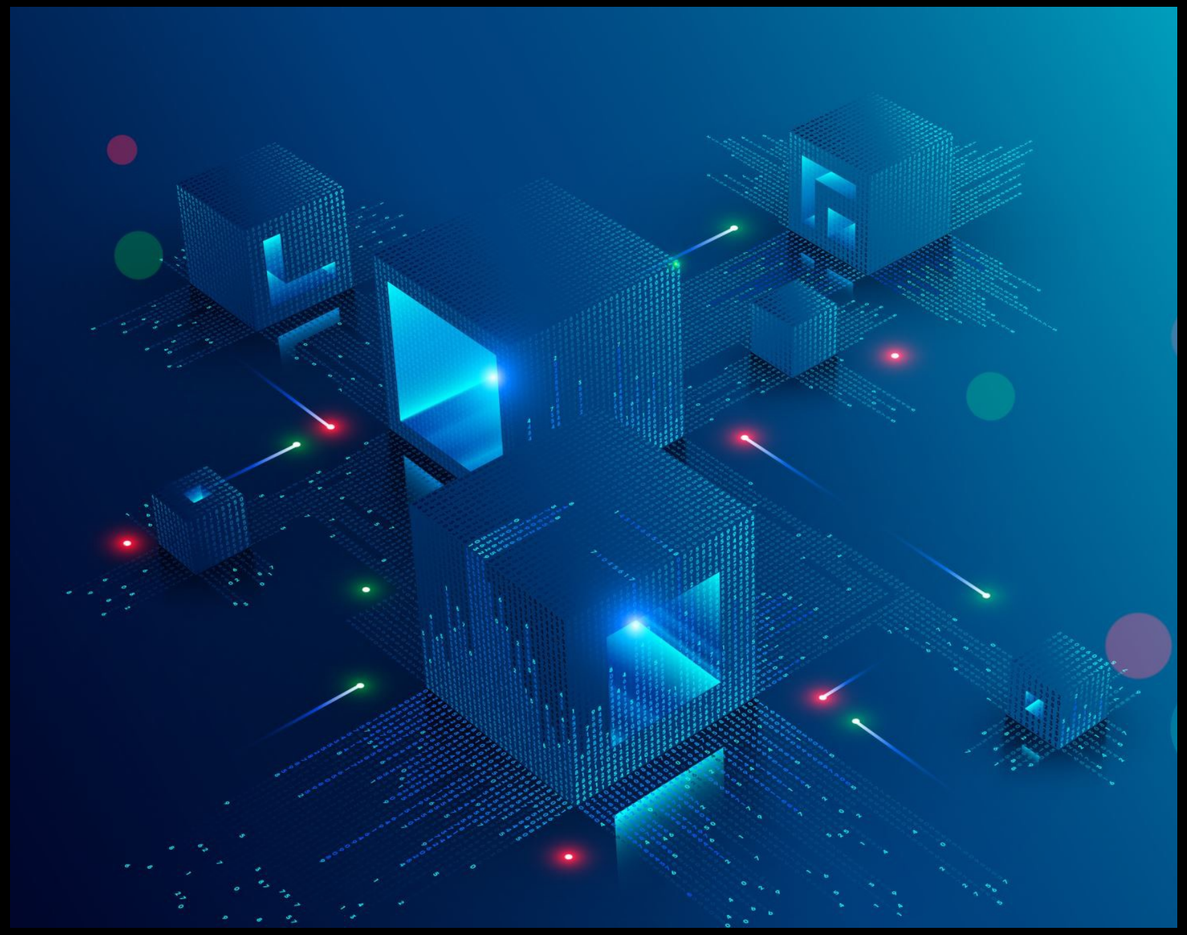

Expand data received from cluster: GPU type, CPU type, queue size

Add in support for cluster management software and OS package management

Release a beta version

© 2022 IBM Corporation

Any Questions?

© 2022 IBM Corporation