Testing of Grid Worker Nodes Integration in OpenStack

ALEXANDRU IOAN CUZA UNIVERSITY of IAŞI www.uaic.ro

Condegrid & RaaS-IS

Ciprian Pinzaru, Valeriu Vraciu, Paul Gasner, Octavian Rusu

ciprian.pinzaru@uaic.ro

Content

- *Hardware resource of RAAS-IS*
- *Ro-16-UAIC grid site*
- *Cloud orchestration for new WN*
- *Testing in production*
- *Conclusion*

Power computation of RAAS-IS

ALEXANDRU IOAN CUZA UNIVERSITY of IAŞI

- For computation:
- 2 x blade system HPE Synergy 12000 Frame
  - Full redundancy: management blade system, fan, power, networking and fibre channel
- 16 x compute node HPE Synergy 480 Gen10:
- 2 x Intel Xeon-Gold 6240 (2.6GHz/18-core)
- 128 GB RAM @ 2933 MHz
- 2 x 300GB SAS 12G Enterprise 10K in RAID 1
- 2 x 25 Gbps Ethernet and 2 x 32 Gbps FC
- For management:
- 4 x rack server HPE ProLiant DL360
- 2 x Intel Xeon-Gold 6240 (2.6GHz/18-core)
- 128 GB RAM @ 2933 MHz
- 2 x 300GB SAS 12G Enterprise 10K in RAID 1
- 2 x 25 Gbps Ethernet and 4 x 1 Gbps Ethernet

Storage of RAAS-IS

ALEXANDRU IOAN CUZA UNIVERSITY of IAŞI

- Enterprise storage system HPE 3PAR 8440 2N
- hybrid system with SSD and HDD
- flash optimization architecture with 64 GB Cache
- 18 x 920GB SAS SFF (2.5in) SSD in RAID 5
- 96 x 8TB SAS 7.2K LFF (3.5in) HDD in RAID 6
- tier configuration for optimization SSD-> HDD (with statistics of data write/read)
- Performance
- 162K IOPS Mixed (read/write), 110K write
- 3.5 GB/s read, 2.4 GB/s write
- Capacity 760 TB RAW, 550 TB USABLE
- IO 4 x 32 Gbps FC and 4 x 16 Gbps

Network and fiber channel of RAAS-IS

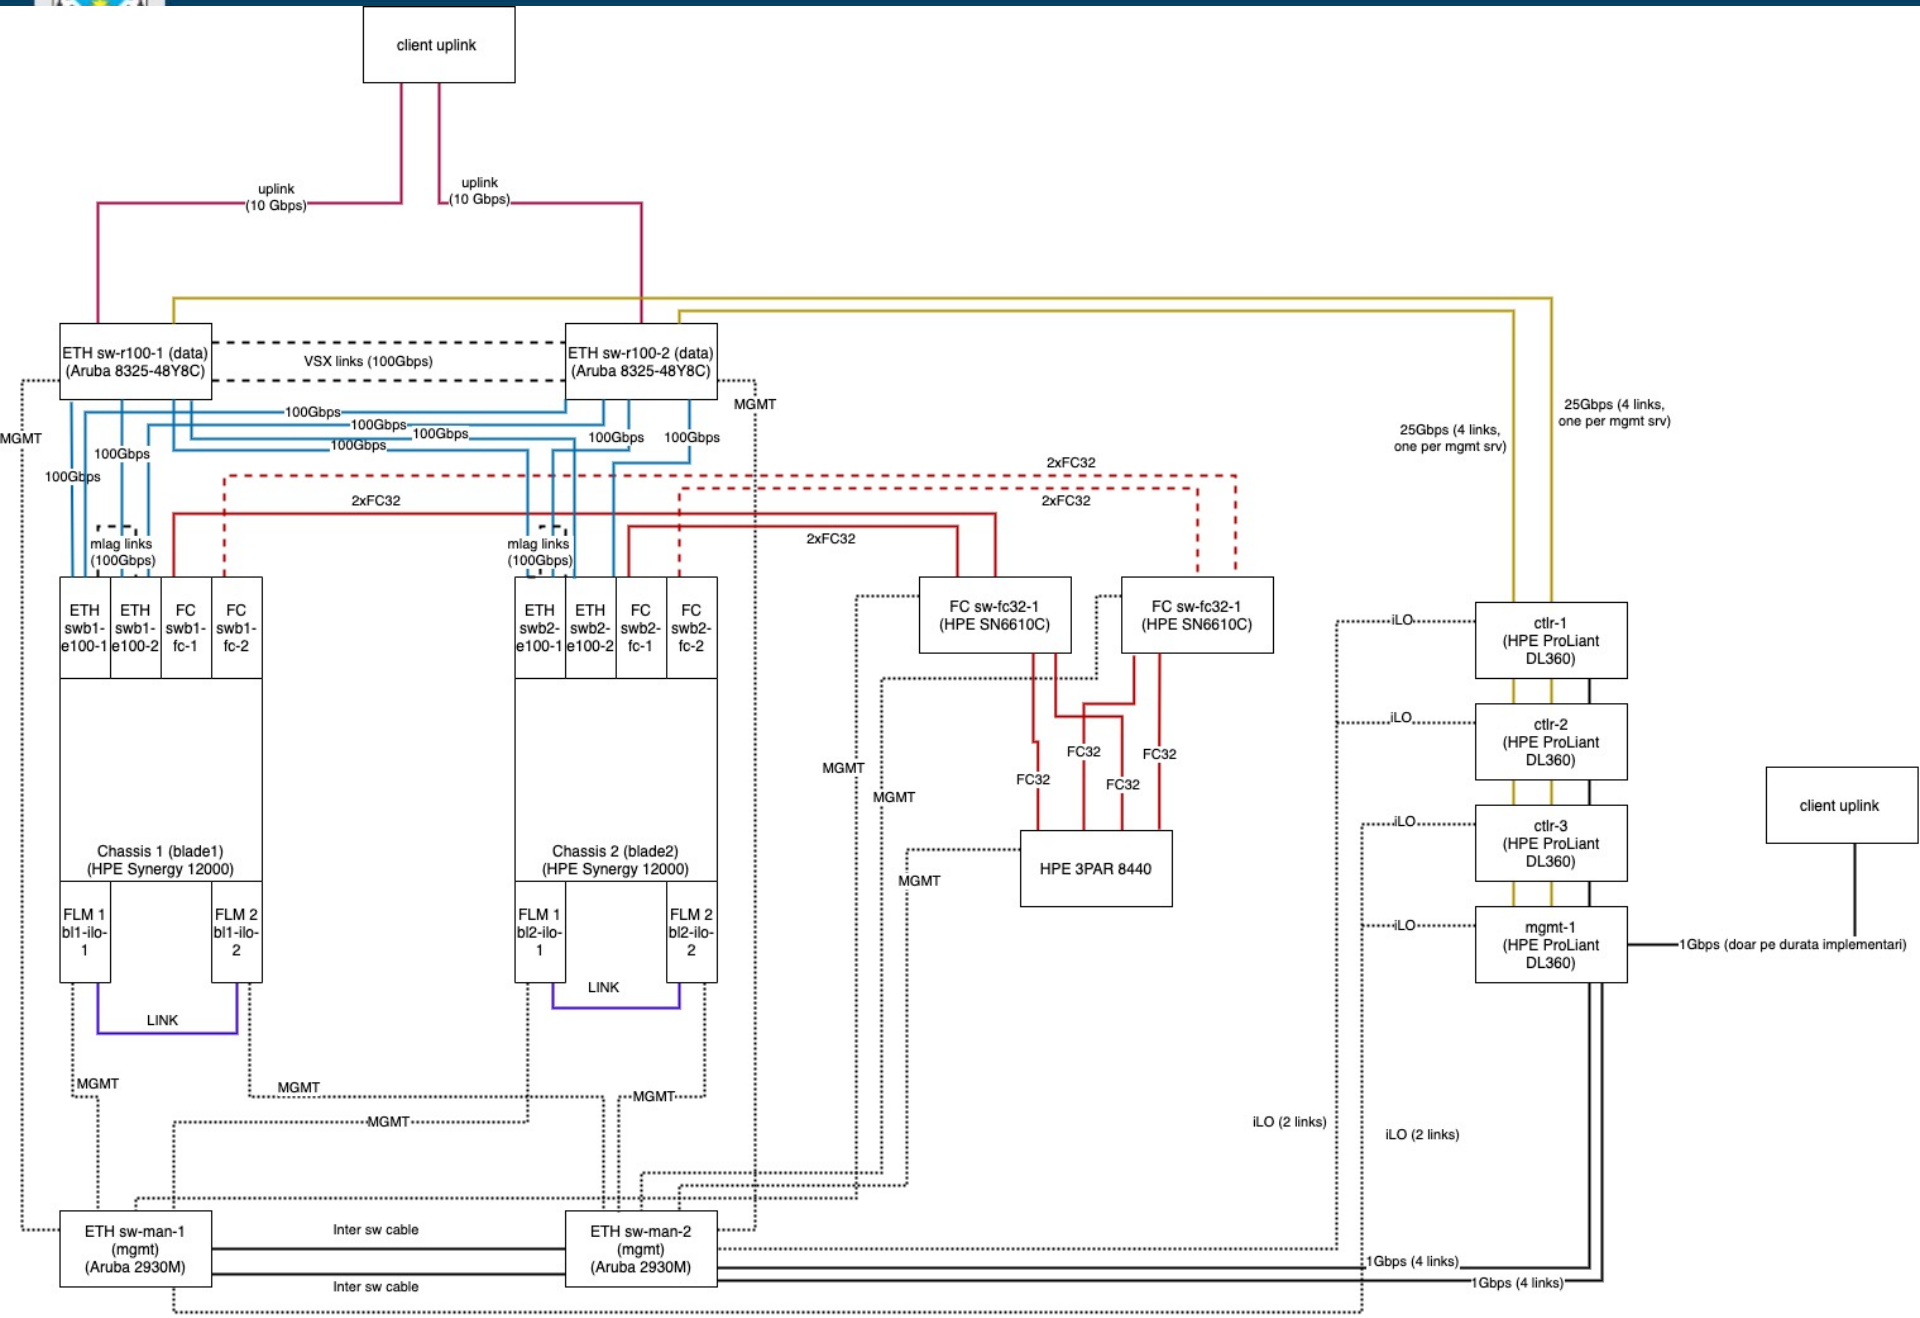

Hardware management and OS of RAAS-IS

- Metal-As-A-Service (MAAS) provisioning with Linux (Ubuntu)
- Infrastructure monitoring and discovery
- Ansible, Chef, Puppet, SALT, Juju integration
- Disk and network configuration
- API-driven DHCP, DNS, PXE, IPAM
- Hardware testing and commissioning

FQDN | MAC IP	POWER	STATUS	OWNER TAGS	POOL NOTE	ZONE SPACES	FABRIC VLAN	CORES ARCH	RAM	DISKS	STORAGE
Deployed 23 machines										—
cpu-sr-b1-1.raas.uaic.ro 192.168.250.11	On Ipmi	Ubuntu 20.04 LTS	admin compute, nova...	default	B1 os-mgmt	fabric-1 Default VL...	72 amd64	127.8 GiB	1	300 GB
cpu-sr-b1-2.raas.uaic.ro 192.168.250.12	On Ipmi	Ubuntu 20.04 LTS	admin compute, nova...	default	B1 os-mgmt	fabric-1 Default VL...	72 amd64	127.8 GiB	1	300 GB
cpu-sr-b1-3.raas.uaic.ro 192.168.250.13	On Ipmi	Ubuntu 20.04 LTS	admin compute, nova...	default	B1 os-mgmt	fabric-1 Default VL...	72 amd64	127.8 GiB	1	300 GB
cpu-sr-b1-4.raas.uaic.ro 192.168.250.14	On Ipmi	Ubuntu 20.04 LTS	admin compute, nova...	default	B1 os-mgmt	fabric-1 Default VL...	72 amd64	127.8 GiB	1	300 GB
cpu-sr-b1-5.raas.uaic.ro 192.168.250.15	On Ipmi	Ubuntu 20.04 LTS	admin compute, nova...	default	B1 os-mgmt	fabric-1 Default VL...	72 amd64	127.8 GiB	1	300 GB
cpu-sr-b1-6.raas.uaic.ro 192.168.250.16	On Ipmi	Ubuntu 20.04 LTS	admin compute, nova...	default	B1 os-mgmt	fabric-1 Default VL...	72 amd64	127.8 GiB	1	300 GB
cpu-sr-b1-7.raas.uaic.ro 192.168.250.17	On Ipmi	Ubuntu 20.04 LTS	admin compute, nova...	default	B1 os-mgmt	fabric-1 Default VL...	72 amd64	127.8 GiB	1	300 GB

Cloud implementation RAAS-IS

All Hypervisors

Hypervisor Summary

VCPU Usage

Used 700 of 1,152

Memory Usage

Used 1.9TB of 2TB

Hypervisor | Compute Host

Displaying 16 items

Hostname	Type	VCPUs (used)	VCPUs (total)	RAM (used)	RAM (total)	Local Storage
cpu-sr-b1-1.raas.uaic.ro	QEMU	42	72	113GB	125.6GB	84GB
cpu-sr-b1-2.raas.uaic.ro	QEMU	30	72	119GB	125.6GB	44GB
cpu-sr-b1-3.raas.uaic.ro	QEMU	78	72	165GB	125.6GB	26GB
cpu-sr-b1-4.raas.uaic.ro	QEMU	34	72	117GB	125.6GB	36GB
cpu-sr-b1-5.raas.uaic.ro	QEMU	54	72	121GB	125.6GB	37GB
cpu-sr-b1-6.raas.uaic.ro	QEMU	34	72	117GB	125.6GB	36GB
cpu-sr-b1-7.raas.uaic.ro	QEMU	27	72	101GB	125.6GB	70GB
cpu-sr-b1-8.raas.uaic.ro	QEMU	44	72	122GB	125.6GB	56GB
cpu-sr-b2-1.raas.uaic.ro	QEMU	32	72	108GB	125.6GB	64GB

Main Services in cloud

Cinder Heat Nova **Swift**

RO-16-UAIC site grid

ALEXANDRU IOAN CUZA UNIVERSITY of IAŞI

- For computation:

- 2 x blade system M1000e and MX7000 Dell
- Full redundancy: management blade system, fan, power, networking and fibre channel
- 17 x server blade from multiple generations:
- 404 core processors (808 hyper threads) and 1.7 TB of RAM running production jobs from VO ATLAS

- For management:

- 2 x rack server R720 Dell
- 2 x Intel Xeon E5649 (2.4GHz/6-core)
- 32 GB RAM and 600GB in Raid 5
- Used for management of grid site in virtualization environment

Cloud orchestration for new WN

- define the YAML template file on Heat

```
description: Simple template to deploy WN
parameters:
  network:
    type: string
    default: magnum-net
  subnet:
    type: string
    default: magnum-subnet
  image:
    type: string
    default: centos7
  flavor:
    type: string
    default: m1.small
```

```
resources:
--------------------
  head_network_port:
    type: OS::Neutron::Port
  properties:
    network: {get_param: network}
  fixed_ips:
     - subnet_id: {get_param: subnet}
```

Cloud orchestration for new WN


```
head:
  type: OS::Nova::Server
  properties:
    key_name: ciprian_pinzaru
    image: {get_param: image}
    flavor: {get_param: flavor}
    networks:
    - port: { get_resource: head_network_port }
    user_data_format: RAW
    user_data:
      str_replace:
        template: |
         #!/bin/bash
        echo "Running boot script"
        setenforce 0
        yum -y install epel-release
        yum –y install wget
        wget https://object-
store.raas.uaic.ro:4433/swift/v1/AUTH_d7b5384cc4e84112aae4343a88bd46a3/
GRID/instal-wn.sh
       chmod 755 install-wn.sh
       ./install-wn.sh
```

Cloud orchestration for new WN

ALEXANDRU IOAN CUZA UNIVERSITY of IAŞI www.uaic.ro

Displaying 1 item

Stack Name	Created	Updated	Status	Actions
wn-grid	2 weeks	Never	Create Complete	Check Stack

Displaying 1 item

Displaying 4 items

Instance Name	Image Name	IP Address	Flavor	Key Pair	Status	Availability Zone	Task	Power State	Age	Actions
wn-grid-node3-ell5iqhw6a6q	centos 7	172.16.16.236	m1.small	ciprian_pinzaru	Active	nova	None	Running	2 weeks	Create Snapshot
wn-grid-node2-ogkqptowztkn	centos 7	172.16.16.244	m1.small	ciprian_pinzaru	Active	nova	None	Running	2 weeks	Create Snapshot
wn-grid-node1-m6ike7pbvosn	centos 7	172.16.16.16	m1.small	ciprian_pinzaru	Active	nova	None	Running	2 weeks	Create Snapshot
wn-grid-head-gpvxqmkzjtm3	centos 7	172.16.16.8, 192.168.254.195	m1.small	ciprian_pinzaru	Active	nova	None	Running	2 weeks	Create Snapshot

Example with one Head and three nodes

Testing in production

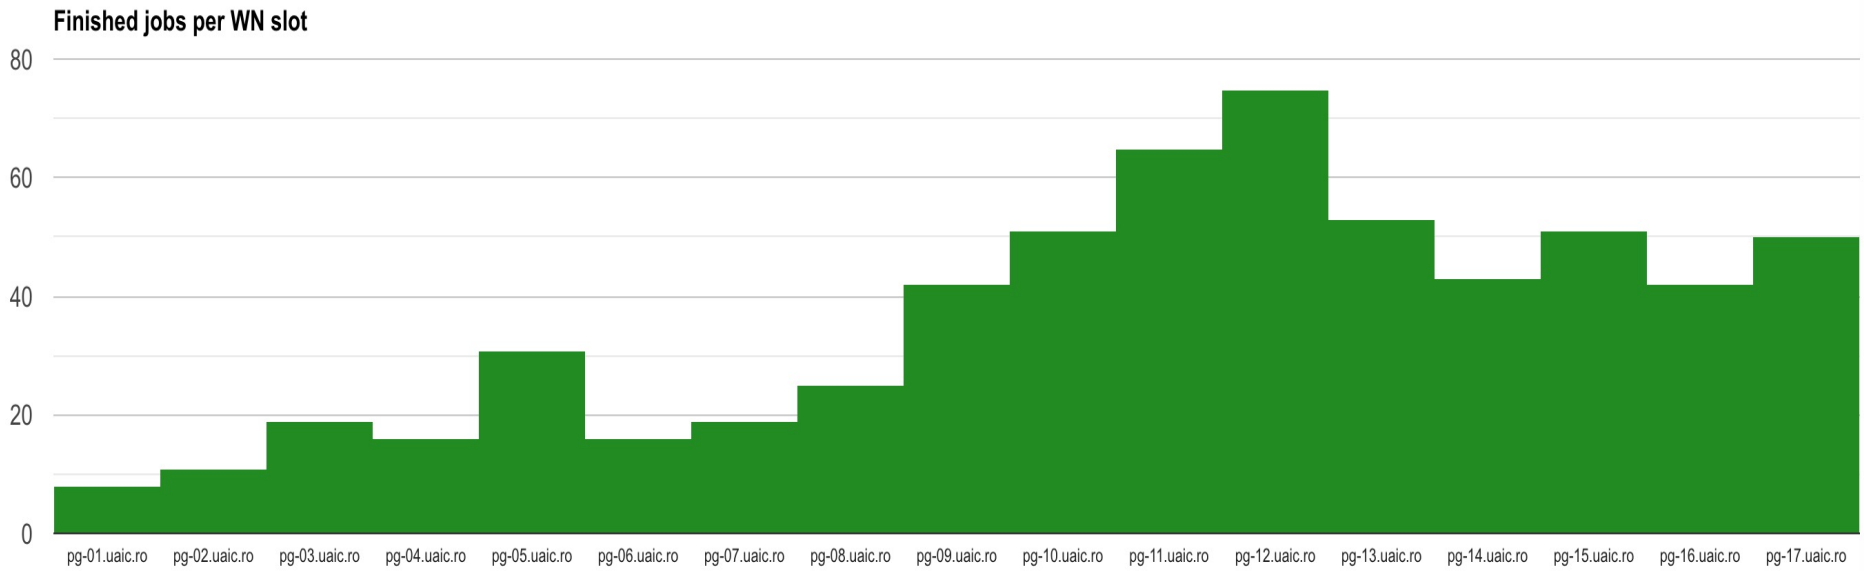

WN	N slots	N running cores	N jobs	defined	waiting	throttled	sent	starting	running	holding	merging	transferring	finished	failed	cancelled	closed	% failed	outliers
pg-17.uaic.ro	1	35	88	-	-	-	-	-	28	1	-	8	50	1	-	-	1	

RO-16-UAIC: 565 k jobs in VO ATLAS, equivalent to 48 millions of sum CPU Work HS06, which represents 15 % from RO-LCG normalized CPU times for this VO

Conclusion

- *Heat service of OpenStack – useful in automatization process of defining grid worker node;*
- *Reduce time and manpower in process of adding WN;*
- *The advantages of a cloud infrastructure will be considered for the development of the RO-16-UAIC grid site;*

ACKNOWLEDGMENT

This work was supported by Romanian National Authority for Scientific Research, project number 7/2020, named "National contribution to the development of the LCG computing grid elementary particle physics" and by the Competitiveness Operational Programme Romania, under project SMIS 124759 - RaaS-IS (Research as a Service Iasi).