

# ESCAPE WP5 Progress Meeting

26-27th October 2020

S. Luna, S. Sánchez-Expósito, J. Garrido, J. Moldón, L. Verdes-Montenegro  
IAA-CSIC



# Outline

- Prototype of SKA Regional Centre at IAA-CSIC
- Activities in the context of ESCAPE's WP5

# Prototype of SKA Regional Centre at IAA-CSIC

- Activities
  - Computing infrastructure procurement and setup
    - Face to face meetings with researchers at IAA-CSIC
    - Group A: 80% of researchers
      - low number of CPU cores ( $\leq 16$ )
      - decent amount of RAM memory per CPU core ( $\leq 8$  GB/core)
      - access to fast and large storage ( $\sim 10$  TB/project)
    - Group B: 20% of researchers
      - high number of CPU cores ( $> 100$ )
      - high amount of RAM per CPU core ( $> 10$  GB/core)
      - access to fast but relatively small storage ( $\leq 1$  TB).
    - Heterogeneous software requirements
      - Most of them transitioning to Python, but still with Fortran + IDL code
  - Deployment of software and services
    - Focus on Open and Reproducible Science
    - Integration with the Virtual Observatory ecosystem
    - Offer support for the SKA Data Challenge 2 (SDC2)


# Prototype of SKA Regional Centre at IAA-CSIC

- The computing infrastructure
  - OpenStack controllers: 3 nodes
  - OpenStack hypervisors: 5 nodes
  - Ceph Monitors: 3 nodes
  - Ceph OSDs: 8 nodes
  - Interconnect: 100 Gbps Ethernet
  - Totals:
    - 200 CPU cores for hypervisors, 384 GB RAM each, except one with 1 TB
    - ~600 TB usable storage capacity, ~400 TB for file, ~200 TB for block



# Activities in the context of ESCAPE WP5

- Open and reproducible paper
  - Evolution of compact groups from intermediate to final stages. A case study of the HI content of Hickson Compact Group 16



- Candidate project to deploy in ESCAPE?
- <https://github.com/AMIGA-IAA/hcg-16>



# Activities in the context of ESCAPE WP5

- JupyterHub deployment: work in progress
  - Thanks to Rohini and Stelios for sharing their examples!
  - Virtualized Kubernetes cluster with OpenStack Magnum
  - Used Zero to JupyterHub Helm chart
  - Currently under testing
    - Checked ESCAPE IAM authentication
    - Using OpenStack Cinder for persistent storage



# Activities in the context of ESCAPE WP5

- JupyterHub deployment: our questions
  - Hardware and software requirements for integration with ESCAPE
    - Integration with rucio?
  - Can we have multiple authentication sources for the same JupyterHub?
  - What authorization options are available?
    - Can you limit access to subgroups of users?
    - Can you assign different resource limits to different groups of users?
  - Pending issue: long timeout logging in, why?
  - Do you recommend block or file storage for persistent storage in JHub?
  - Can you mount file shares into user sessions?
  - Can you mount extra additional block storage into user sessions?

# What's next?

- Interests for the future
  - We are open to collaborate with other centres and SRC prototypes
  - Find synergies with other ESCAPE members
    - Best practices for operating a private cloud with OpenStack, Ceph
    - Deploying common interest tools like JupyterHub
    - Federated user authentication and authorization