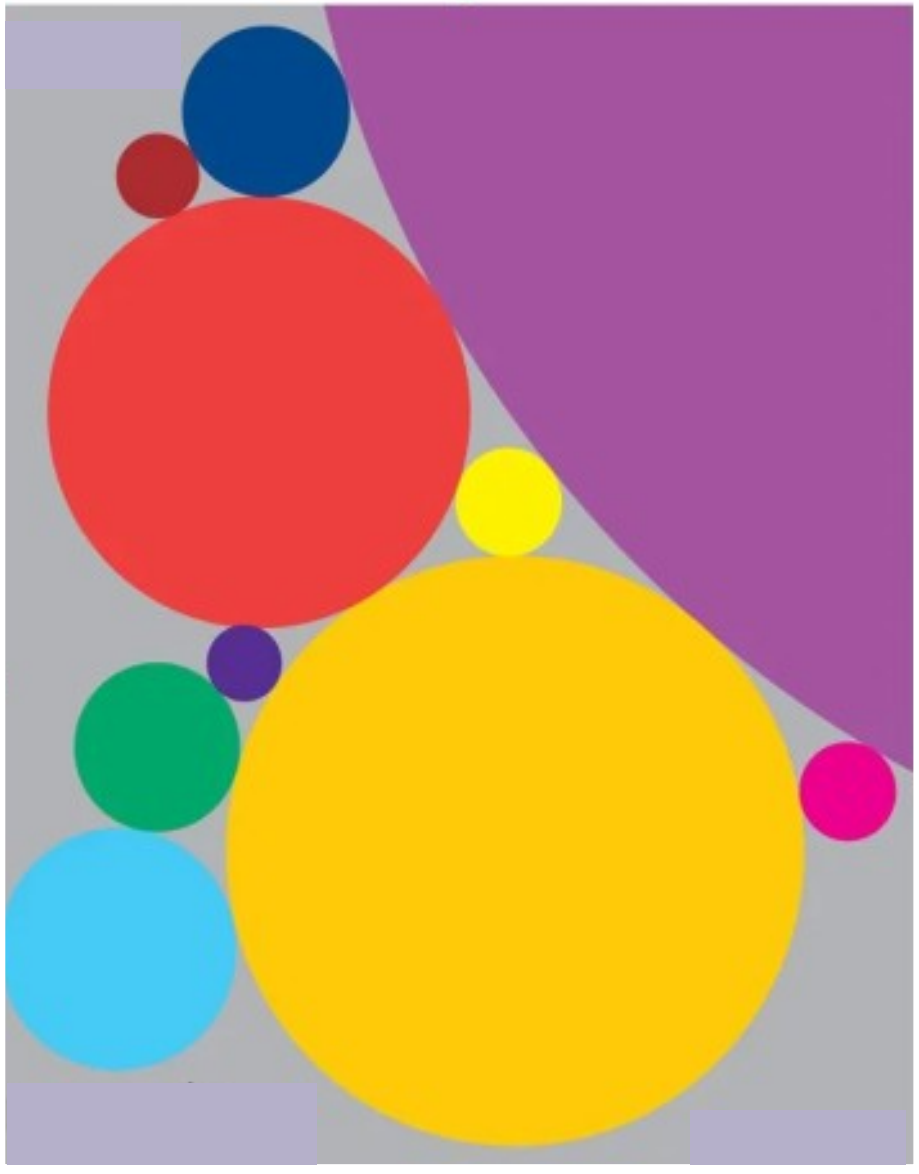

Archives and Data Management Systems in the Big Data Era

26–28 Feb 2025 CNR Bologna

The IVOA ExecutionBroker, an overview

D.Morris, S.Bertocco

sara.bertocco@inaf.it
dave.morris@manchester.ac.uk

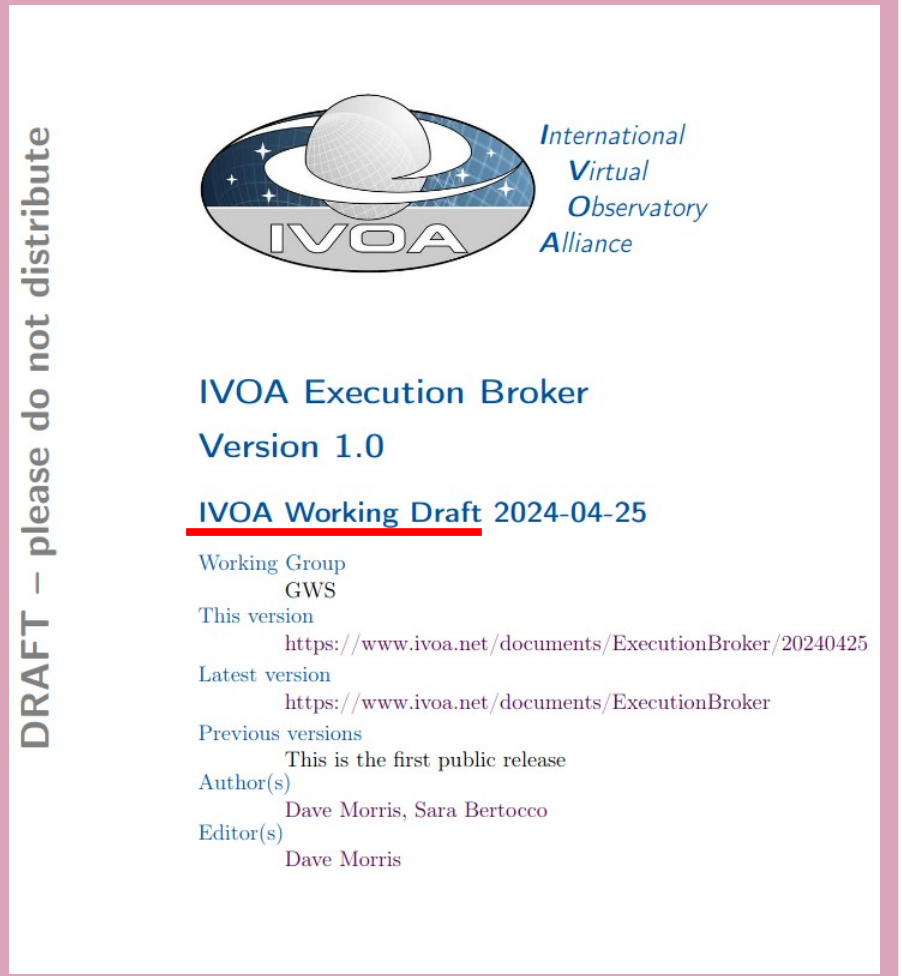

https://github.com/ivoa-std/ExecutionBroker
The IVOA Execution Broker web-service and data model specification.

IVOA ExecutionBroker

- What is it ?
- Why ?
- Why should we use it ?

IVOA

The International Virtual Observatory Alliance (IVOA) is an organisation that debates and agrees the technical standards that are needed to make the VO possible

VO

The Virtual Observatory (VO) is the vision that astronomical datasets and other resources and services should work as a seamless whole

IVOA Documents lifecycle

1. working group prepares working draft (version ≥1.0) and submits to document coordinator for posting in the ivoa documents.
2. working group reviews the working draft. two reference implementations of any associated software and validation tools are expected.
3. the chair of the working group, with consent of the wg, promotes the document to a proposed recommendation.
4. the chair of the working group issues a formal request for comments (rfc) to the e-mail distribution list interop@ivoa.net. . The chairs and vice chairs of other working groups are required to examine proposed recommendations during the rfc period
5. the working group chair and/or editor(s) respond to comments , if comments lead to significant changes to the document, the wg chair may decide to revert the document status to working draft (back to step 1).
6. if comments are addressed to the satisfaction of the wg chair and wg members, the wg chair requests a final vote, if successful, the document is sent to the executive committee for approval.
7. the executive committee is polled by the ivoa chair to ascertain if there is consensus for promotion to recommendation.
8. if yes, the ivoa chair reports on approval to the tcg and wg chairs and asks the document coordinator to update the document status to recommendation. if no, the concerns of the ivoa executive need to be resolved and a new poll taken, or if serious revisions are required, the document would revert to step 1

https://www.ivoa.net/documents/

```
# OpenAPI schema
DockerContainer:
  description: |
    A Docker or OCI container.
    See https://opencontainers.org/
  type: object
  title: DockerContainer
  allOf:
    - $ref: 'AbstractExecutable'
    - type: object
      properties:
        repository:
          description: >
            The image respository URL.
          type: string
        image:
          description: >
            The image name within the repository.
          type: string
        tag:
          type: string
          description: >
            The image version tag.
        digest:
          description: >
            The image digest checksum, used for verification.
```

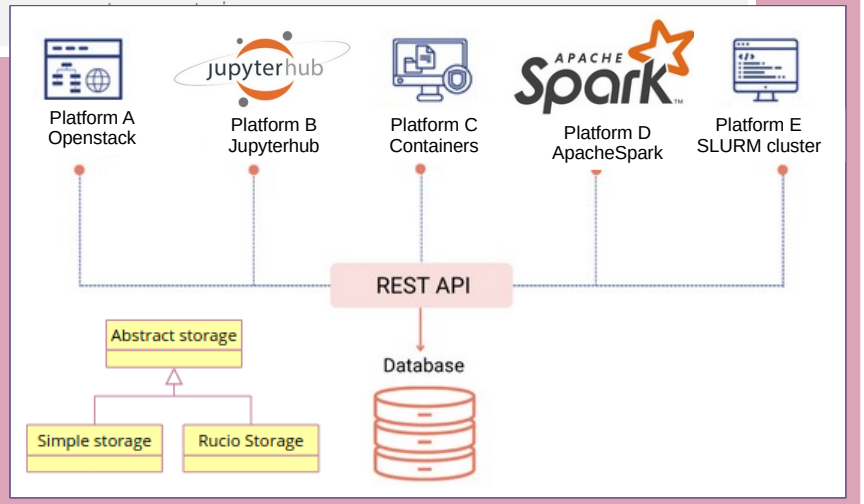

What is the ExecutionBroker ?

- a data model for describing ‘executable things’
- a web service:

  a REST based web service to find execution platforms, allocate resources and schedule execution sessions.

IVOA ExecutionBroker

- What is it ?
- **Why ?**
- Why should we use it ?

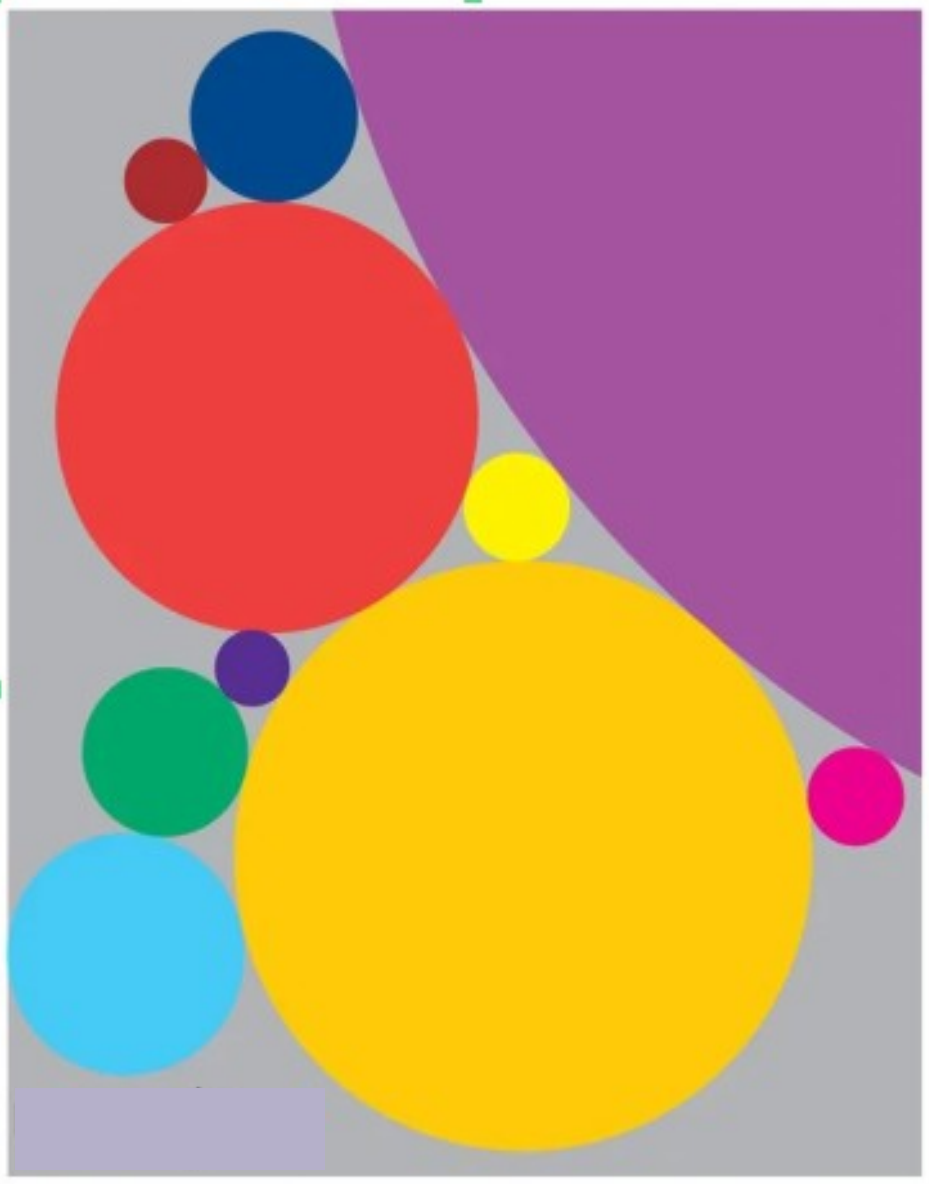

Data growth, Big Data

Challenges

- Volume
- Velocity
- Variety

Comparing to the World Outside

ATLAS
Computing Model Evolution

Borut Paul Kersevan
Simone Campana

Data Management
Where is HEP in Big Data Terms?

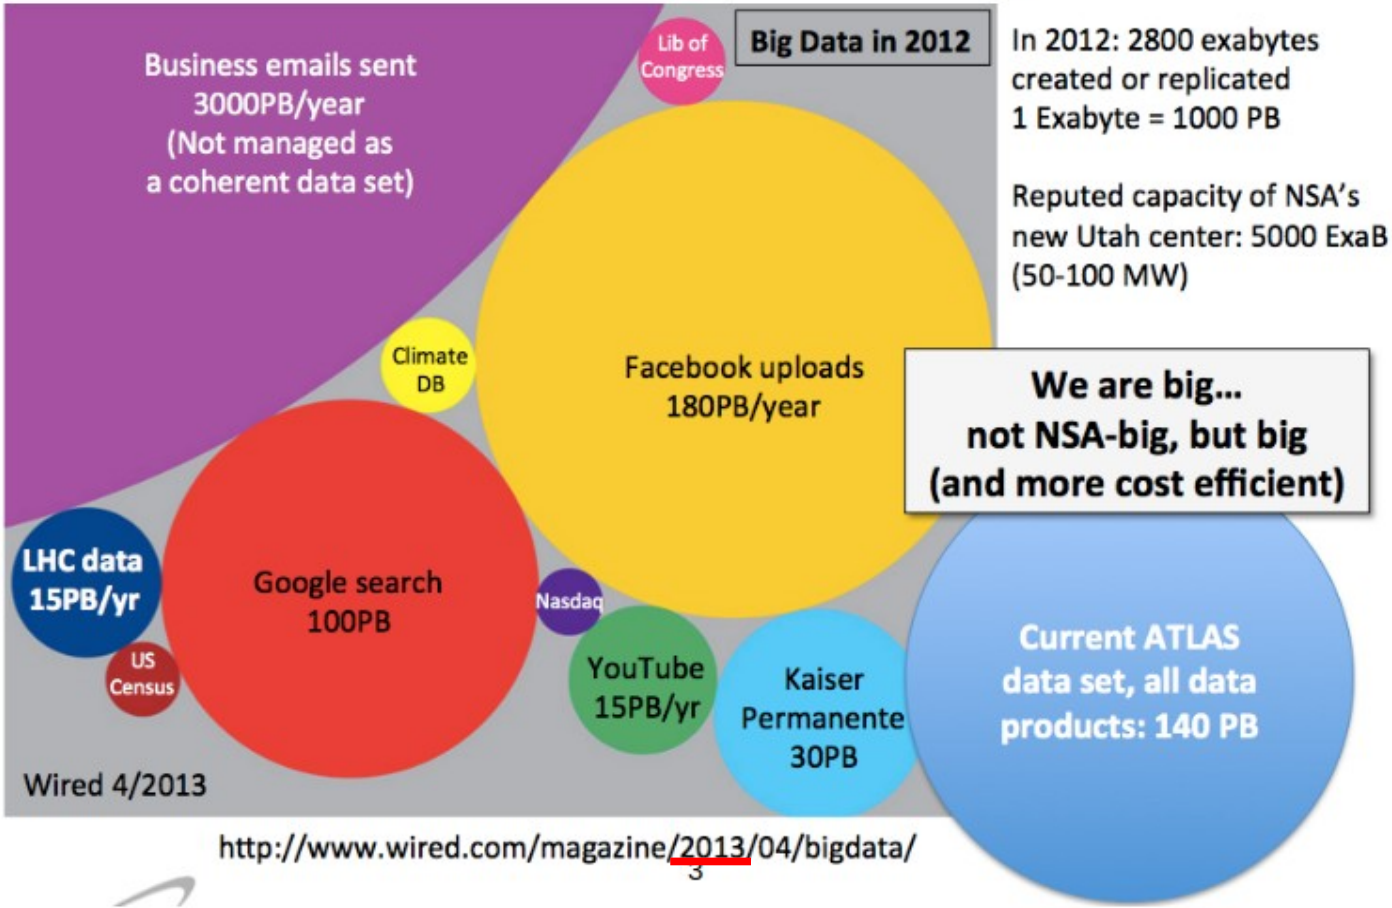

http://www.wired.com/magazine/2013/04/bigdata/

Comparison of yearly data volumes of current and future projects

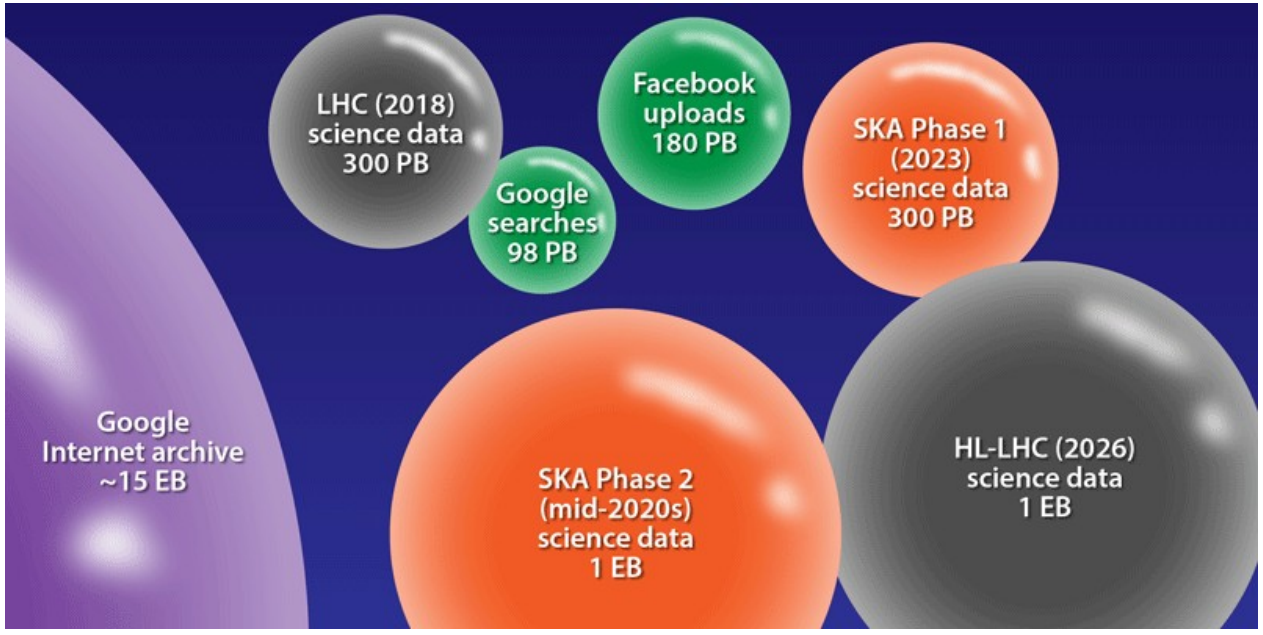

Article: “Facing a Downpour of Data, Scientists Look to the Cloud” February 3, 2020• Physics 13, 14 https://physics.aps.org/articles/v13/14

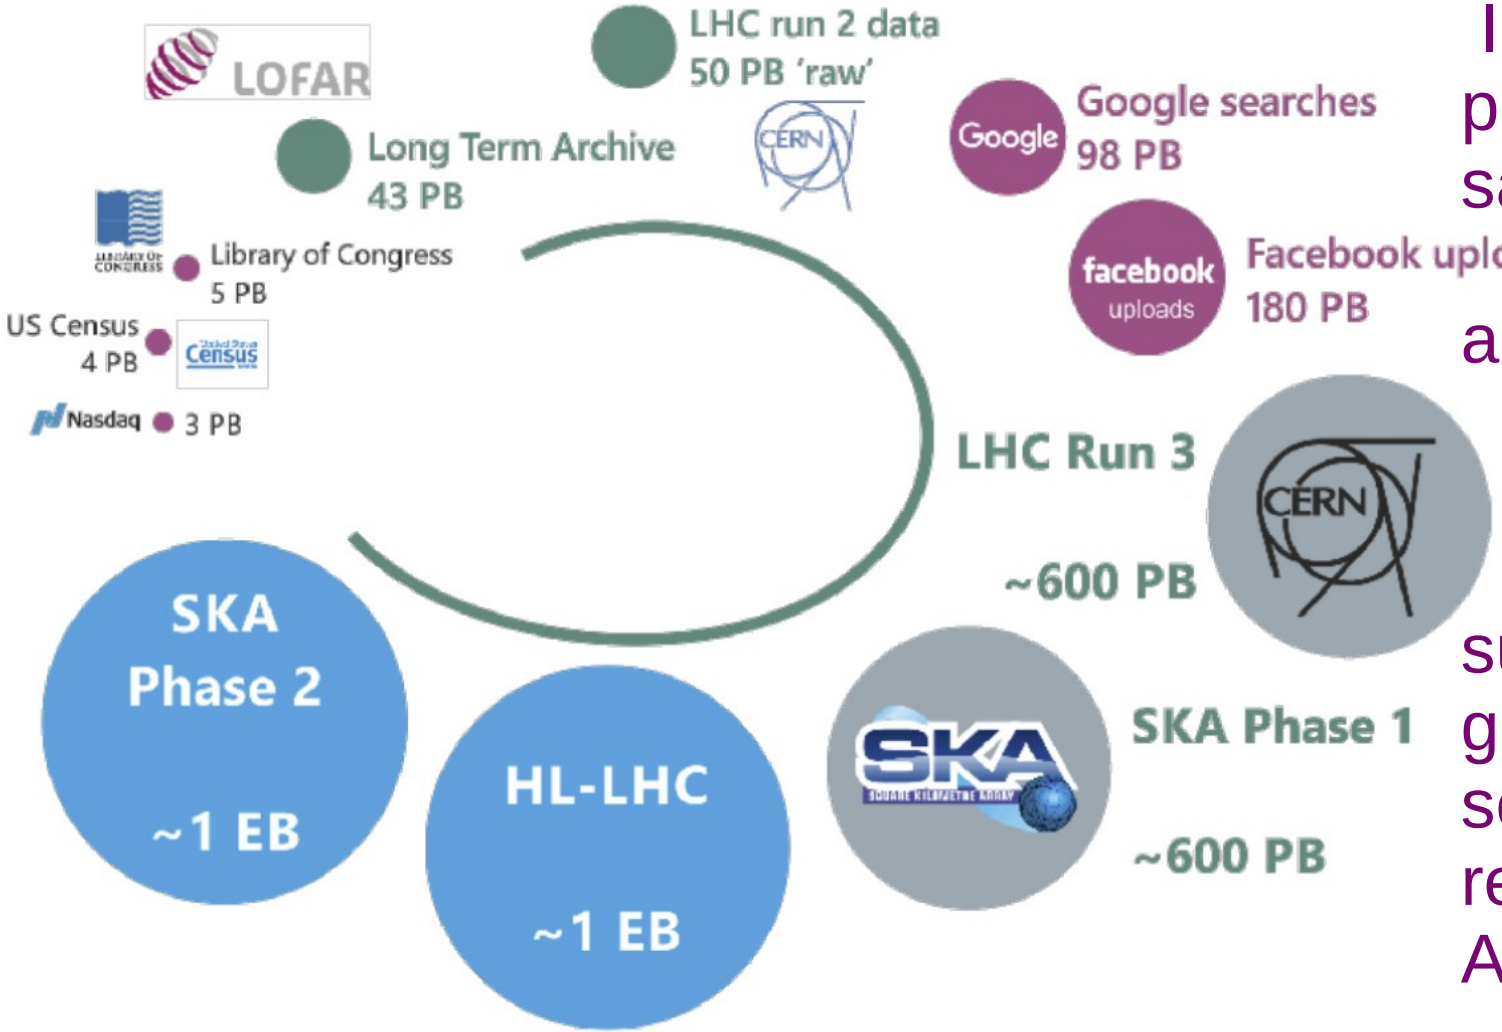

Il traffico di dati prodotto da SKA sarà più di 3 volte superiore all'odierno traffico generato su Facebook e supererà il più grande progetto scientifico ad oggi realizzato, il Large Adron Collider

https://www.oact.inaf.it/2023/09/19/radioastronomia-galattica-con-ska-e-precursori/

Data analysis through computing resources heterogeneous and spread around the world

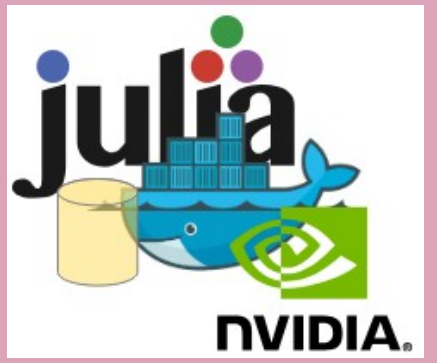

for data analyses

- lots of different types of software
  - different requirements
  - different interfaces

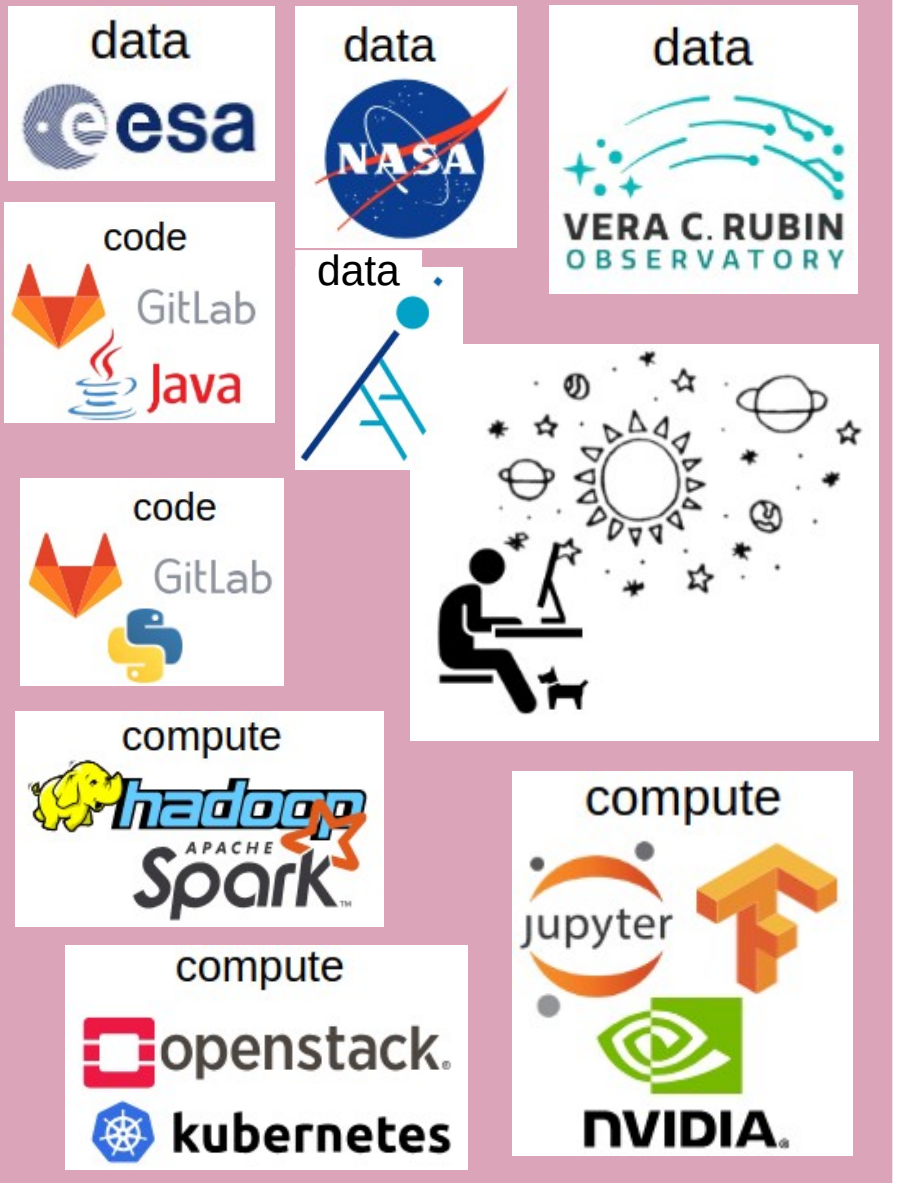

- Data spread around the world
- Heterogeneous code (analyses software)
- Heterogeneous computing resources geographically distributed

Everyone is slightly different

Thanks to Dave Morris

IVOA ExecutionBroker

- What is it ?
- Why ?
- **Why should we use it ?**

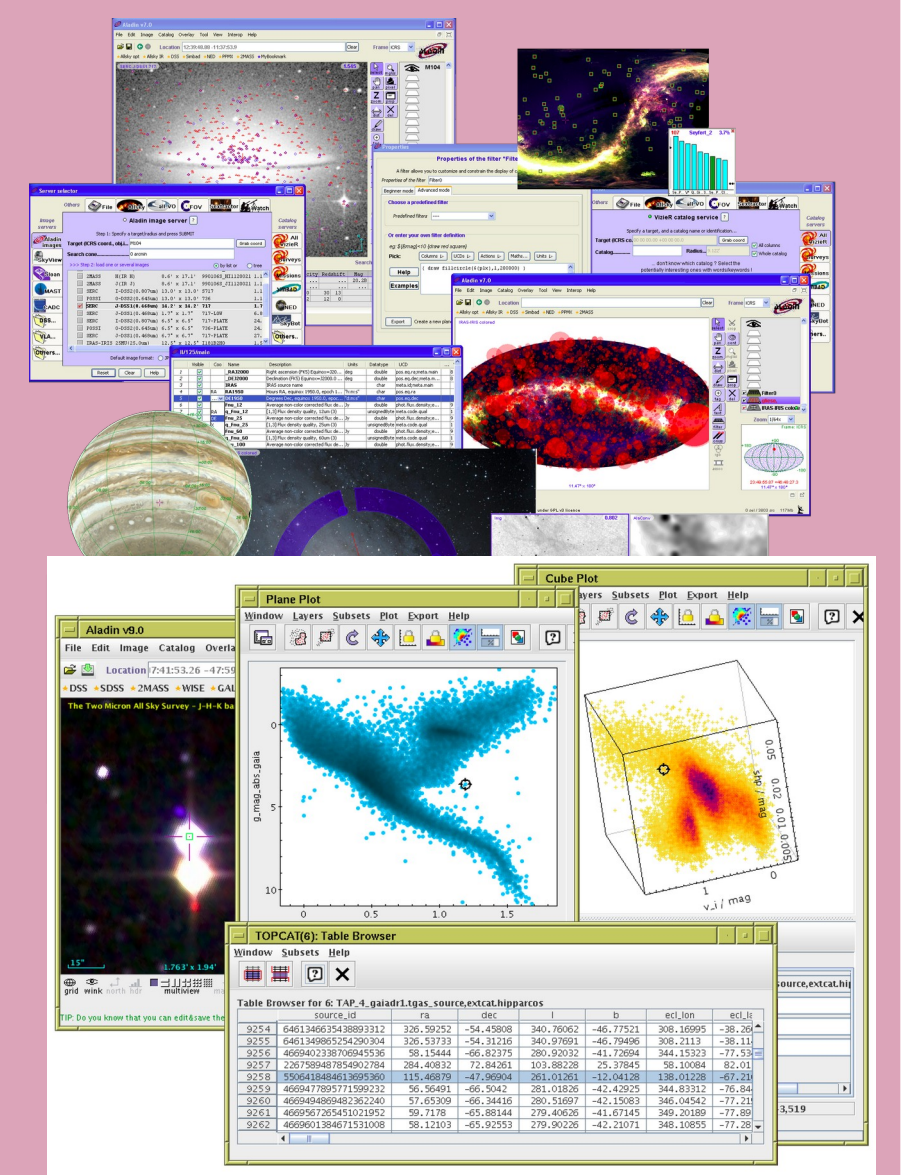

The Virtual Observatory is a solution for data:
it provides standards (data models, dictionaries, data access protocols, etc.) enabling tools and services to access data geographically distributed

Execution Broker: The REST interface

is a common interface for different science platforms

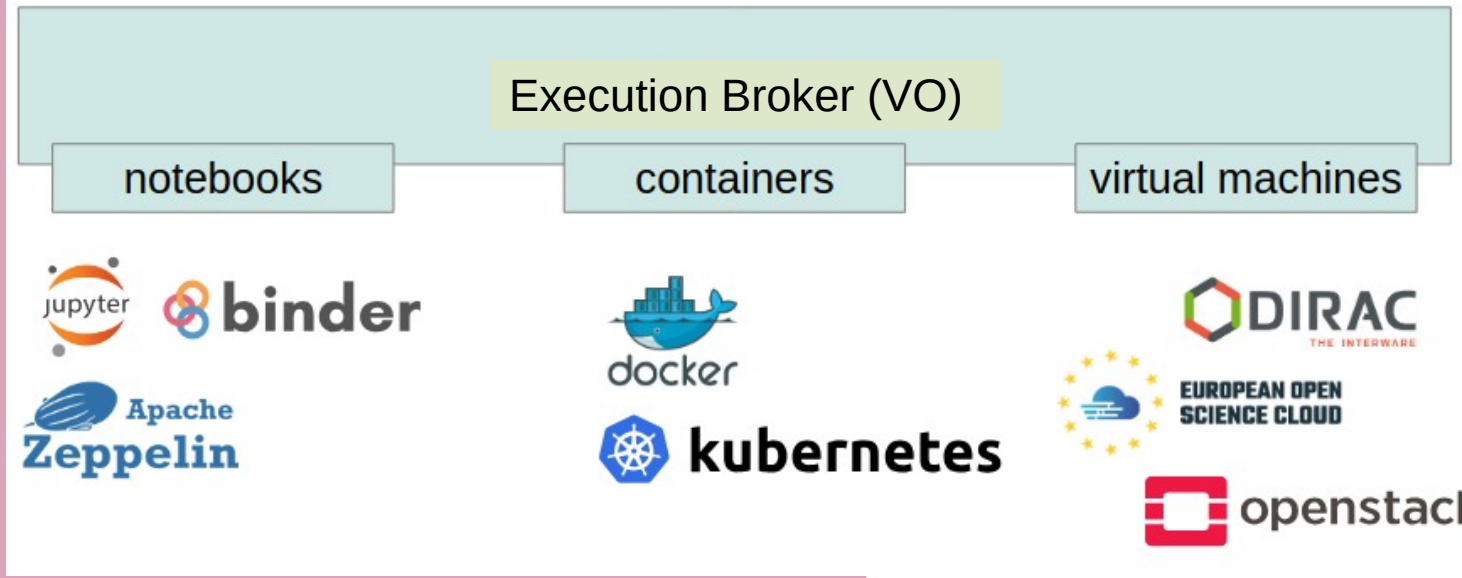

Different science platforms use different technologies.

Each platform only needs to understand the technologies it provides. If a platform doesn't understand the question, it can just say no.

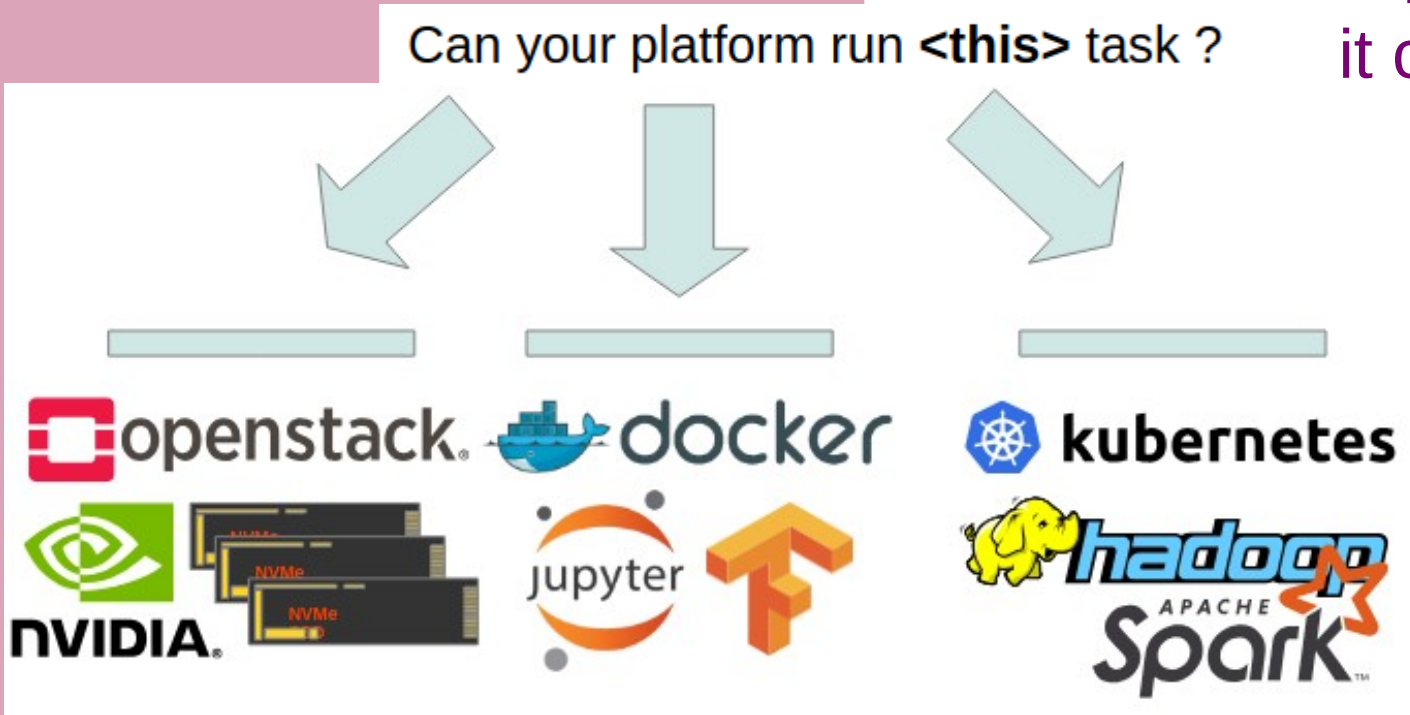

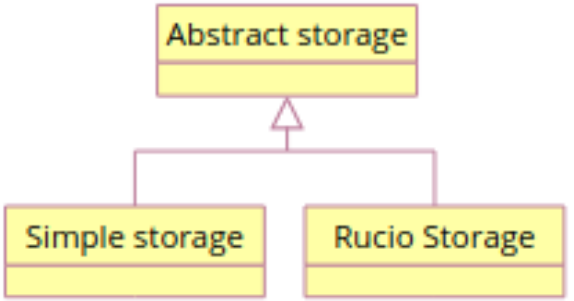

Plugin architecture helps

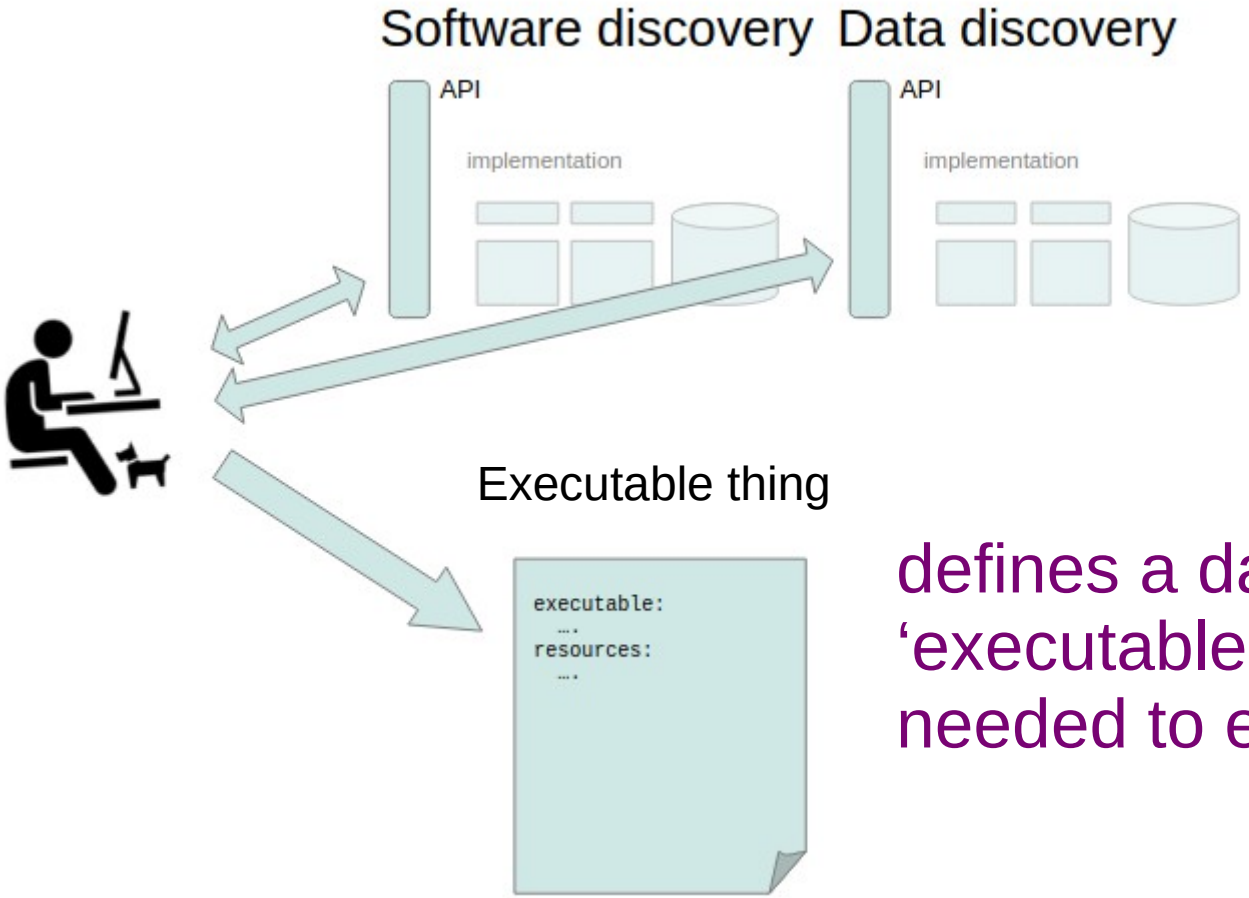

ExecutionBroker: data model

defines a data model for describing ‘executable things’ and the resources needed to execute them

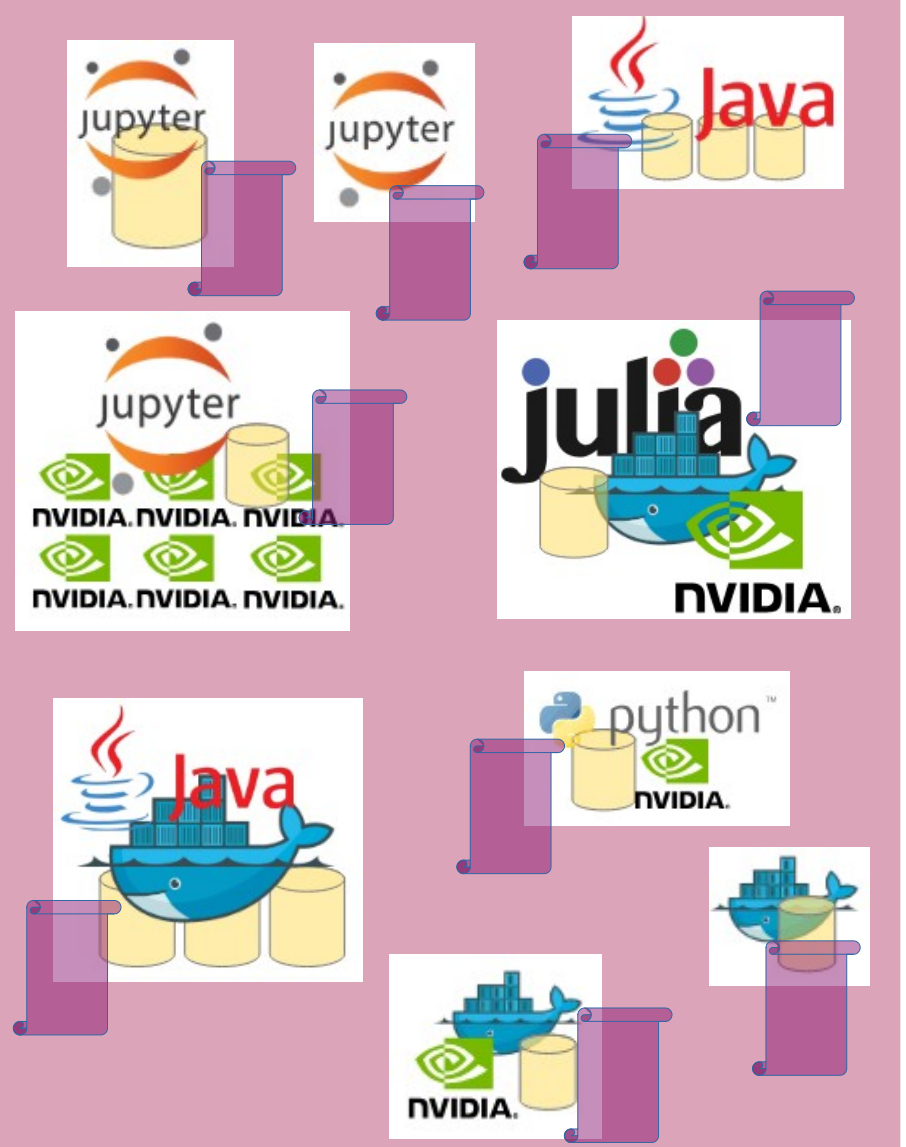

Use a common data model to describe executable things

Openapi data model description

```yaml
# OpenAPI schema
AbstractExecutable:
  type: object
  discriminator:
    propertyName: type
    mapping:
      "uri:java-program-1.0":       'JavaProgram'
      "uri:python-program-1.0":     'PythonProgram'
      "uri:docker-container-1.0":   'DockerContainer'
      "uri:singular-container-1.0": 'SingularContainer'
      "uri:jupyter-notebook-1.0":   'JupyterNotebook'
      ....
  properties:
    name:
      description: >
        A human readable name, assigned by the client.
      type: string
    uuid:
      description: >
        A machine readable UUID, assigned by the server.
      type: string
      format: uuid
    type:
      description: >
        The type discriminator.
      type: string
      format: uri
```

```yaml
# OpenAPI schema
DockerContainer:
  description: |
    A Docker or OCI container.
    See https://opencontainers.org/
  type: object
  title: DockerContainer
  allOf:
    - $ref: 'AbstractExecutable'
    - type: object
      properties:
        repository:
          description: >
            The image respository URL.
          type: string
        image:
          description: >
            The image name within the repository.
          type: string
        tag:
          type: string
          description: >
            The image version tag.
        digest:
          description: >
            The image digest checksum, used for verification.
          type: string
          ....
```

OpenAPI is an interoperable, machine-readable, and human-friendly specification format used to define HTTP APIs. It relies on JSON Schema to describe the API's underlying data.

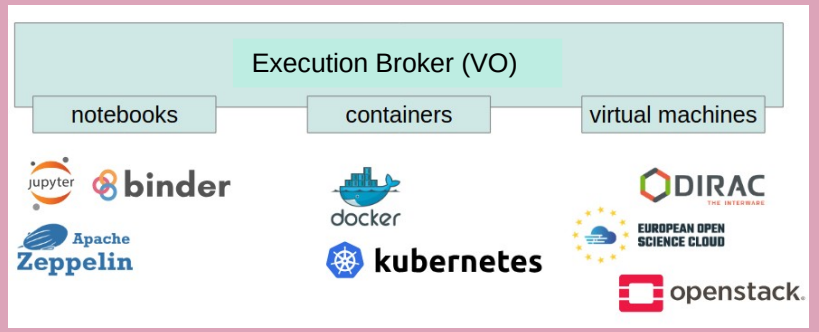

ExecutionBroker API:

The IVOA Execution Broker is designed to address a specific question: given an executable thing, e.g. a Python program or Jupyter notebook, in which facilities can I run it?

Pass a common data-model description to a common interface

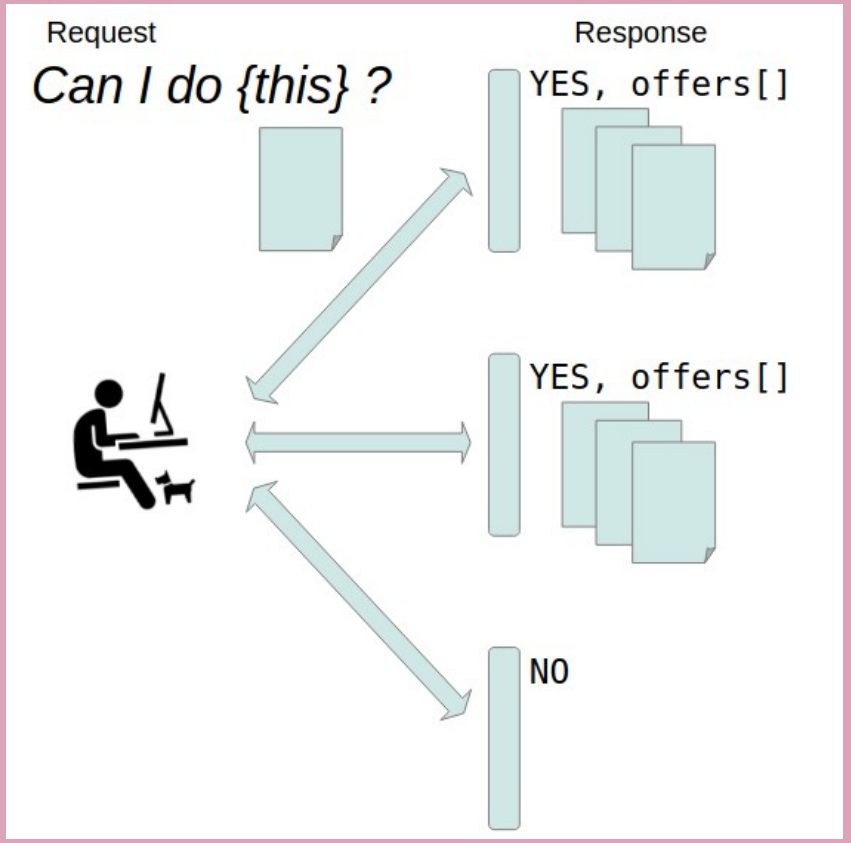

ExecutionBroker API

- Query one or more ExecutionBroker services.
- Each ExecutionBroker replies with one or more offers. The offers will include details of the available resources and estimated cost.

Client Request

```
# ExecutionBroker offerset request.
executable:
  # Common fields from AbstractExecutable
  name: Experiment two
  type: uri:jupyter-notebook-1.0

  # The type specific details for a Jupyter notebook.
  notebook:
    file: https://.../example.jpnb
  requirements:
    file: https://.../requirements.txt
resources:
  compute:
    - type: uri:generic-compute
      cores:
        requested:
          min: 4
      memory:
        requested:
          min: 16
```

```
# ExecutionBroker offerset response.
executable:
  type: uri:jupyter-notebook-1.0
  notebook:
    file: https://.../example.jpnb
resources:
  compute:
   - type: uri:generic-compute
      cores:
        requested:
          min: 4
        offered:
          min: 8
      memory:
        requested:
          min: 16
        offered:
          min: 32
```

Service Response

Acknoledgment and rights:

- Thanks Dave Morris! For images, schemas, content

- https://www.vecteezy.com/free-vector/man-thinking"
  Man Thinking Vectors by Vecteezy