



The CANFAR Science Platform

Brian Major, Canadian Astronomy Data Centre
IVOA Interop, May 2021

Agenda



- Quick overview, what's new
- Interoperability and Standard Potential
 - User Contributed Software
 - Groups for Allocations and Discovery

CANFAR Science Platform Overview



CADC: Canadian Astronomy Data Centre (Victoria, BC)

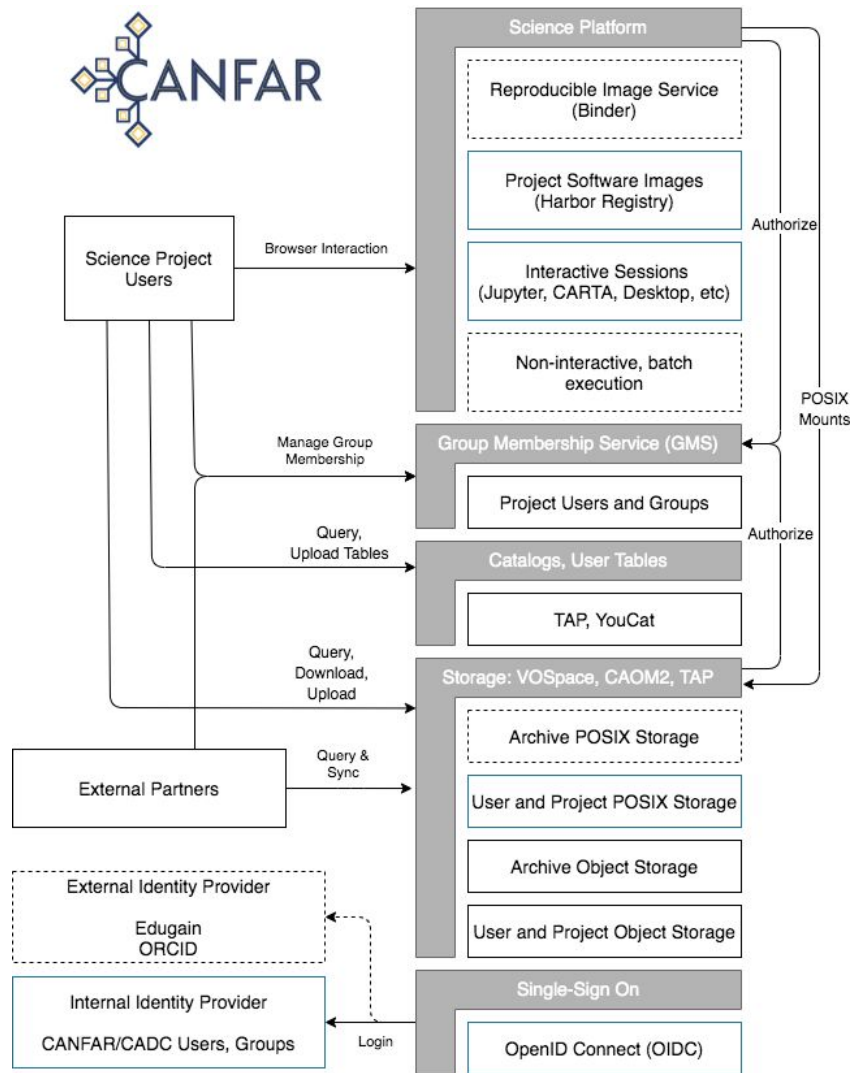
CANFAR: Canadian Advanced Network for Astronomical Research

- A lot of work over the last 18 months:
 - ARCADE Development Study - ALMA Reduction in the CANFAR Data Environment
 - CANARIE Research Software Program - Jupyter Notebooks and Image Registry
- CANFAR:
 - VMs on OpenStack - multi purpose and batch
 - **New: 100% Container-based, higher-level Science Platform**
- Production deployment in place, scaling up users cautiously
- Kubernetes cluster provided by Compute Canada / UVic Research Computing.

CANFAR Architecture & Design

- Be guided by a science problem that someone wants solved.
- Look to solve the general problem rather than the specific issue (if possible)
- Accept that science users will see the general approach as longer to implement, but we will satisfy more users, eventually and create sustainable infrastructure.
- Use an IVOA (on top of whatever the latest shiny is) and help develop standards where new ones are needed.

<https://github.com/opencadc/skaha>



Platform API

skaha is our platform API that pulls together:

- Authentication (SSO, OIDC Client)
- Groups and Authorization (GMS)
- POSIX VOSpace (cavern)
- User-controlled software container/image registry
- Kubernetes

<https://github.com/opencadc/skaha>

Images

GET

/image

Resource Contexts

GET

/context

Session Management

GET

/session

POST

/session

DELETE

/session/{sessionID}

POST

/session/{sessionID}/app

[BASE URL: /skaha , API VERSION: 0.3]

CANFAR Portal



[Documentation](#) [Services](#) [About](#) [Open Source](#) [Support](#)

[Brian Major](#)

Canadian Advanced Network for Astronomical Research

[Take the tour](#)



Storage Management



Group Management



Data Publication



Open Stack



CADC Search



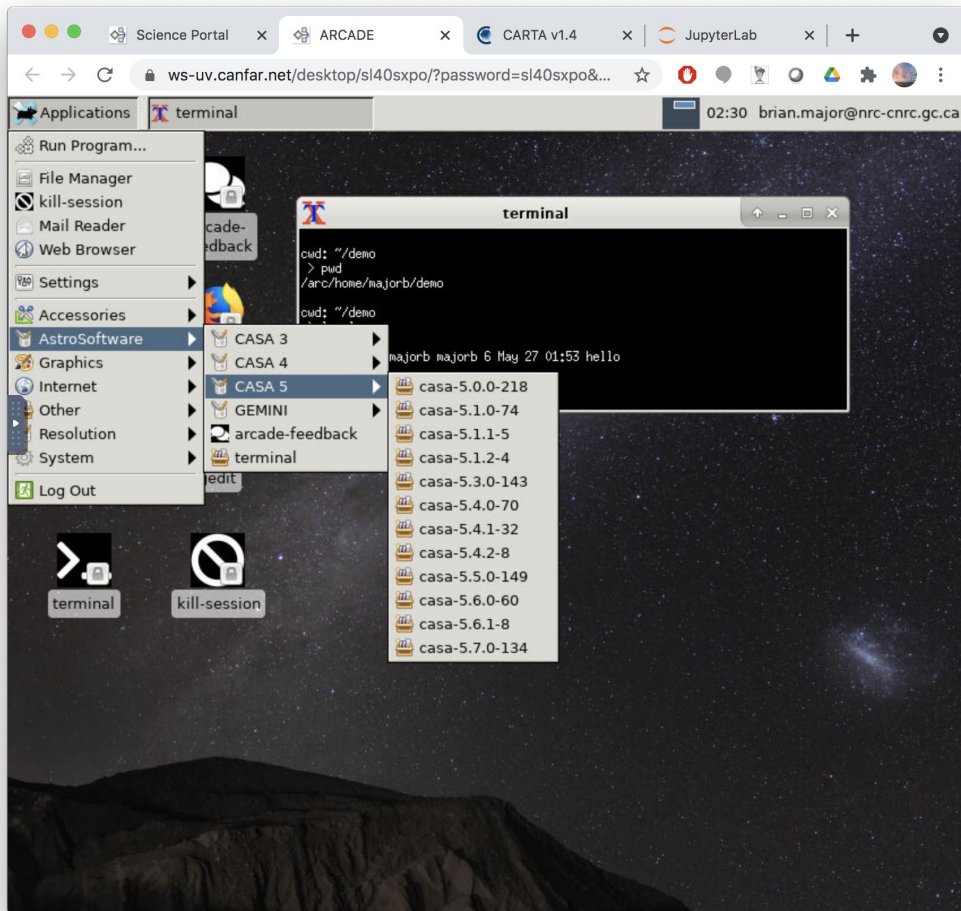
Science Portal

CANFAR Science Portal

The screenshot displays the CANFAR Science Portal interface. At the top, there are browser tabs for Science Portal, ARCADE, CARTA v1.4, and JupyterLab. The address bar shows the URL `canfar.net/science-portal/`. The main header features the CANFAR logo and navigation links: Documentation, Services, About, Open Source, Support, and a user profile for Brian Major.

The "Active Sessions" section shows three active sessions: `carta1` (carta), `desktop1` (desktop), and `notebook1` (notebook). Below this is a "Launch" section with a "Session Type" dropdown menu. A tooltip is visible over the dropdown, stating: "Select from the list of supported session types (notebook, carta, desktop.)". The dropdown is currently set to "notebook". Other configuration options include "container image" (set to `images.canfar.net/highz-alma/notebook-astrophy:0.1`), "memory" (set to 16), and "# cores" (set to 2). There are "Launch" and "Reset" buttons at the bottom of the form.

Desktop Session



CARTA

Session

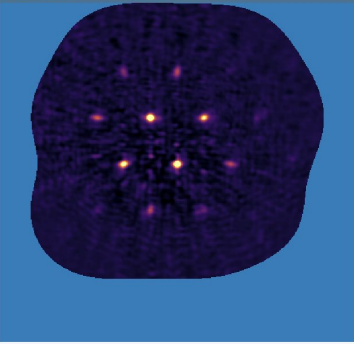
Science Portal x ARCADE x CARTA v1.4 x JupyterLab x

ws-uv.canfar.net/carta/http/bjc3ggtg/?socketUrl=wss://ws-...

File View Widgets Help

BHR71_cont_p1.image

WCS: (12:01:27.16, -65:08:14.8); Image: (501, 426); Value: 1.49326e-1 Jy/beam*

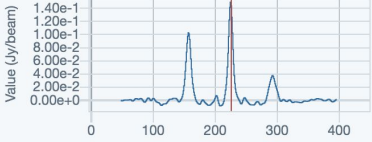


Declination

Right ascension

X Profile: Cursor

12:01:43 12:01:39 12:01:36 12:01:33 12:01:30




Value (Jy/beam)

X coordinate

Data: (WCS: 12:01:27.16, Image: 501 px, 1.49326e-1)

Y Profile: Cursor

-65:09:40 -65:09:20 -65:09:00 -65:08:40 -65:08:20




Value (Jy/beam)

Y coordinate

Data: (WCS: -65:08:14.8, Image: 426 px, 1.49326e-1)

Render Configuration

Clip Percentile 99.9%



Scaling Linear

Color map

Invert color map

Clip Min -0.0089254

Clip Max 0.0753930

Value (Jy/beam)

Image List

Image	Layers	Matchin
0 BHR71_cont_p1.image	R	XY

Notebook Session

The screenshot displays a JupyterLab interface within a browser window. The browser tabs include Science Portal, ARCADE, CARTA v1.4, and JupyterLab. The address bar shows the URL `ws-uv.canfar.net/notebook/cmri0jtp/lab`. The JupyterLab interface features a top menu bar with options: File, Edit, View, Run, Kernel, Tabs, Settings, and Help. On the left, a file browser shows a directory structure with folders like `demo`, `lab`, `nbconvert`, `pip`, and `yarn`, and files like `a`, `ASTR_407.507.Assig...`, `casa-20210420-161...`, `false`, `hello`, `migrated`, `Untitled.ipynb`, `untitled.md`, `untitled.txt`, `Untitled1.ipynb`, and `untitled1.txt`. The main area contains a terminal window titled "Terminal 1" with the following output:

```
majorb@skaha-notebook-majorb-cmri0jtp-df65b:~$ cd
majorb@skaha-notebook-majorb-cmri0jtp-df65b:~$ id
uid=20001(majorb) gid=20001(majorb) groups=20001(majorb),100(users),30093(ABC),30127(CANFAR-Staff),30497(caom2TestGroupWrite),34241(CADC),34337(CAOM2),35030(W-CADC),35039(W-ALL),35124(cadcstats),35130(cadc-dev),35219(INAF-Group),35440(MAST-RW),36002(arbutus-cloud-users),36003(CADC-DEVELOPMENT),36227(jao-cadc),1025424273(skaha-users),1240980498(rc-harbor),1477619040(ARCADE-Users),1623998838(skaha-admins)
majorb@skaha-notebook-majorb-cmri0jtp-df65b:~$
majorb@skaha-notebook-majorb-cmri0jtp-df65b:~$
majorb@skaha-notebook-majorb-cmri0jtp-df65b:~$ cd demo
majorb@skaha-notebook-majorb-cmri0jtp-df65b:~/demo$ ls -l
total 1
-rw-r--r-- 1 majorb majorb 6 May 27 01:53 hello
majorb@skaha-notebook-majorb-cmri0jtp-df65b:~/demo$
```

Cavern Storage

- Cavern Storage, accessible through:
 - POSIX mounts through all containers
 - VOSpace API
 - \$HOME and \$PROJECT dirs
- Each container **'runs as the user'** (with the **user's groups**) so plain POSIX permissions used in cavern

```
cwd: ~/demo
> pwd
/arc/home/majorb/demo

cwd: ~/demo
> ls -l
total 1
-rw-r--r-- 1 majorb majorb 6 May 27 01:53 hello

cwd: ~/demo
> █
```

```
Default
majorb@NRC-009195 ~ $ vls -l arc:/home/majorb/demo
-rw----r-- majorb          NONE          NONE
                               6 May 26 18:53 hello
majorb@NRC-009195 ~ $ █
```

```
Terminal 1
majorb@skaha-notebook-majorb-cmri0jtp-df65b:~$ id
uid=20001(majorb) gid=20001(majorb) groups=20001(majorb),100
97(caom2TestGroupWrite),34241(CADC),34337(CAOM2),35030(W-CAD
c-dev),35219(INAF-Group),35440(MAST-RW),36002(arbutus-cloud-
adc),1025424273(skaha-users),1240980498(rc-harbor),147761904
majorb@skaha-notebook-majorb-cmri0jtp-df65b:~$
majorb@skaha-notebook-majorb-cmri0jtp-df65b:~$ echo $HOME
/arc/home/majorb
```

```
/home/majorb/demo
Search Name... Home Up Root + Add - Download - Move
0.00 B remaining of 0.00 B (Request more from support@canfar.net)
Showing 1 to 1 of 1 entries
Name Size Last Modified (UTC) Read/Write Read
hello 6.00 B 2021-05-27 - 01:53:42 Public
```

Future



- Reproducibility with Binder: from Dockerfile to Container to paper.
- Cooperating containers: visualizers and notebooks exchanging data. Container-to-container SAMP?
- GPU integration into cluster (Nvidia vGUP V100D-8C)
- Batch scheduling



User contributed software

Harbor image registry

- > `docker login images.canfar.net`
- > `docker push images.canfar.net/skaha/notebook-scipy:0.2`

Projects

PROJECTS 7 PRIVATE 3 PUBLIC 10 TOTAL

REPOSITORIES 17 PRIVATE 5 PUBLIC 22 TOTAL

166 1500GB Limit STORAGE

+ NEW PROJECT X DELETE

<input type="checkbox"/>	Project Name	Access Level	Role	Repositories Count
<input type="checkbox"/>	arcade	Public	-	1
<input type="checkbox"/>	canucs	Private	-	4
<input type="checkbox"/>	chimefrb	Private	-	1
<input type="checkbox"/>	gbrammer	Private	-	1
<input type="checkbox"/>	highz-alma	Private	-	2
<input type="checkbox"/>	k-pop	Private	-	6
<input type="checkbox"/>	new-earth	Private	-	0
<input type="checkbox"/>	new-new-earth	Private	-	3
<input type="checkbox"/>	skaha	Public	-	4
<input type="checkbox"/>	uvickbos	Public	-	0

Projects < skaha

notebook-scipy

Info Artifacts

SCAN ACTIONS

Pull Command	Tags	Size	Vulnerabilities	Annotations	Labels	Push Time
	0.2	925.86MB	Not Scanned		noteb...	3/10/21, 9:22 AM

Harbor images you are allowed to use



Science Portal

Active Sessions

carta1 desktop1 notebook1

+ ↻

Launch Session

name [?](#)
notebook2

type [?](#)
notebook

container image [?](#)
images.canfar.net/highz-agma/notebook-astrophy:0.1

memory [?](#)
16

cores [?](#)
2

Launch Reset

- select stack
- images.canfar.net/canucs/canucs:0.6
 - images.canfar.net/canucs/canucs:0.5
 - images.canfar.net/canucs/canucs:0.4
 - images.canfar.net/gbrammer/canucs-copy:0.1
 - ✓ images.canfar.net/highz-agma/notebook-astrophy:0.1
 - images.canfar.net/highz-agma/notebook-scipy:0.2
 - images.canfar.net/k-pop/astromamba-notebook:0.2
 - images.canfar.net/k-pop/kpop-notebook:1.0
 - images.canfar.net/k-pop/kpop-notebook:1.1
 - images.canfar.net/k-pop/isochrones-notebook:1.0
 - images.canfar.net/k-pop/galpyorbits-notebook:1.0
 - images.canfar.net/k-pop/graces-notebook:1.1
 - images.canfar.net/k-pop/graces-notebook:1.2
 - images.canfar.net/k-pop/graces-notebook:1.3
 - images.canfar.net/new-new-earth/nesim:0.1.3
 - images.canfar.net/new-new-earth/snap:0.5.0
 - images.canfar.net/new-new-earth/snap:0.4.2
 - images.canfar.net/new-new-earth/snap:0.4.1
 - images.canfar.net/skaha/notebook-scipy:0.2

Projects bring their containers

- Each project creates their own containers or use public ones already available
- Projects maintain, update, version, and curate their software (containers)
- Agreement between skaha and containers
 - Must be able to run as **non-root** user
 - Must fit one of the **Session Types**
 - Aiming to reduce the number of types (generalize), or;
 - Allow users to define new session types
- Aim to for little or no skaha-specific container requirements.

Session/Container Types

Browser Sessions - HTML5/websocket containers

- Jupyter Notebooks, Desktop (NoVNC), CARTA Vis, Pluto, ...
- Connect to session canfar.net domain with port 443
- Each uses different **ports**
- Path rewriting usually required and supported by browser apps

Non-browser containers

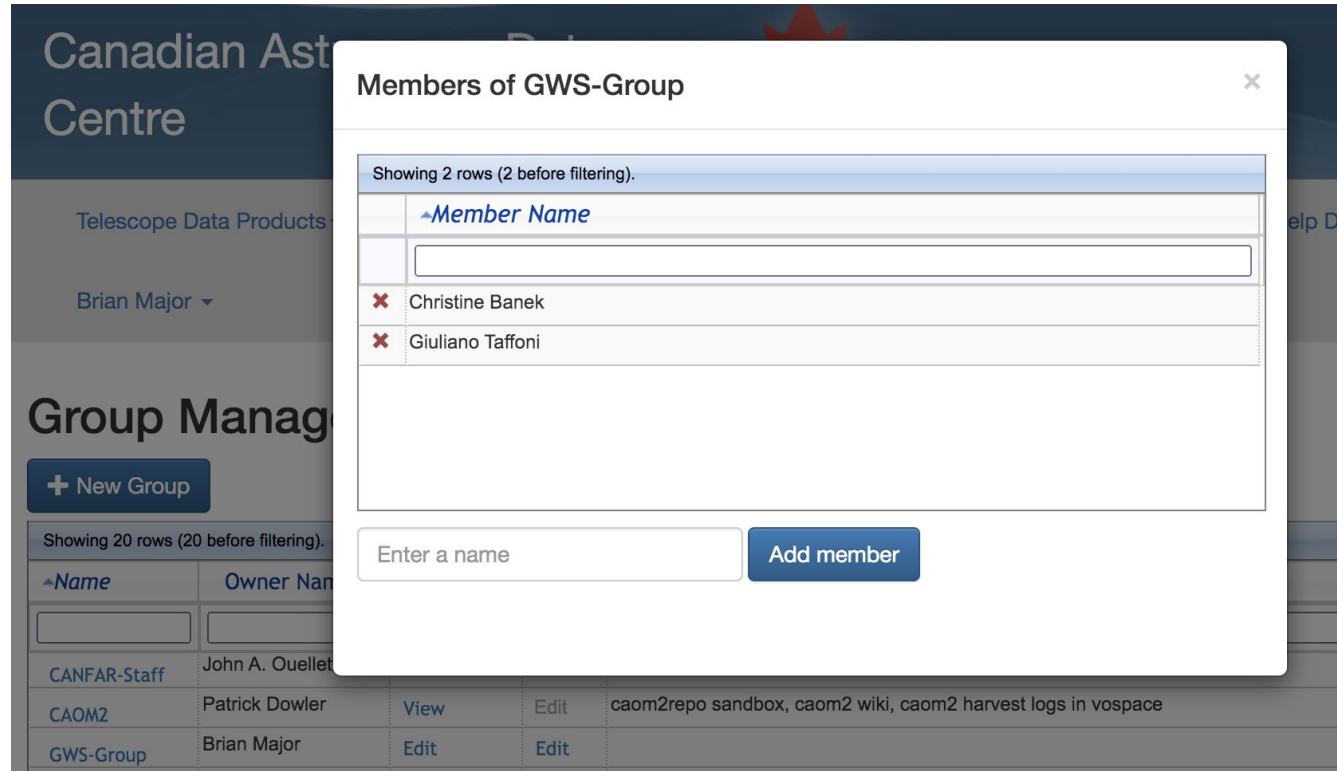
- Headless -- just run the container as is, no proxy necessary
 - API allows override of entrypoint and parameters
 - Step towards batch
- Desktop container
 - Entrypoint is always xterm
 - Window appears in Desktop session, but is it's own container running elsewhere.

Reproducibility:

- Avoid installing software in \$HOME (cavern)

Users manage their team

- Anonymous access to a science platform isn't likely.
- Projects and Users arrange for authorization and allocations with resources.

The screenshot shows the "Group Management" interface for the Canadian Astrophysics Centre. A modal window titled "Members of GWS-Group" is open, displaying a list of group members. The background interface includes a navigation menu with "Telescope Data Products" and "Brian Major" (with a dropdown arrow), a "+ New Group" button, and a table of groups. The modal window shows a search bar, a list of members with red 'X' icons for removal, and an "Add member" button.

Canadian Astrophysics Centre

Telescope Data Products

Brian Major ▾

Group Management

+ New Group

Showing 20 rows (20 before filtering).

Name	Owner Name			
<input type="text"/>	<input type="text"/>			
CANFAR-Staff	John A. Ouellet			
CAOM2	Patrick Dowler	View	Edit	caom2repo sandbox, caom2 wiki, caom2 harvest logs in vospace
GWS-Group	Brian Major	Edit	Edit	

Members of GWS-Group

Showing 2 rows (2 before filtering).

Member Name

✗ Christine Banek

✗ Giuliano Taffoni

Enter a name

Add member

Groups for allocations



- Allocations to CANFAR are assigned to **projects**
- Projects are represented by a **group**
- Membership checks in the groups are done through the Group Membership Service (**GMS**)
- CADC extension of GMS allows project admins to **control user membership**

- Resources assigned to projects/groups:
 - Storage space with quota in POSIX mounted VOspace (cavern)
 - Storage space with quota in object store VOspace (vault)
 - Storage space with quota in docker image registry
 - CPU core and ram allocations in kubernetes cluster
 - Project database schema in user-managed TAP (youcat)

Groups for discovering platform capabilities



- Storage space with quota in POSIX mounted VOspace (cavern)
- Storage space with quota in object store VOspace (vault)
- Storage space with quota in docker image registry
- CPU cores and ram allocations in kubernetes cluster
- Project database schema in user-managed TAP (youcat)

Question: What can this platform do?

More Pertinent Question: **What can I do** in this platform and how do I do it?

Answer: For all the groups in which you are a member: where are those groups assigned?

A reverse GMS/group lookup.

Groups for discovering platforms capabilities



Example: Group A is assigned to a project in an **image registry** with properties:

- Access URL
- Space used, space remaining
- Client protocols supported
- Authentication mechanism supported

Example: Group A is assigned to a project in **POSIX VOspace** with properties:

- Service ID/IVOID (leads to Access URL)
- Space used, space remaining
- Mount point used in containers (eg /cavern)
- Remaining details provided by VOspace itself

Pattern: Where groups are assigned says what the members can do:

- Registry/Discovery (what is it granting?)
- StandardID for the resource (how do I use it?)
- StandardID-specific details

Discovering all IVOA platform capabilities via Group Grants



Question: What can this platform do?

Revised: What can I do in this platform and how do I do it?

Bigger Question: What can I do with **all available IVOA platforms**?

Answer: Where have my groups been set across all platforms?

- Platform A Container Registry
- Platform A Storage
- Platform A Compute
- Platform B Compute

Thoughts:

- How to do find all the group grants? Nothing 'central' in IVOA except the Registry.
- Platforms open up access to other public/private container registries.
- GMS groups don't have institutional boundaries so are a good fit.

THANK YOU

brian.major@nrc-cnrc.gc.ca

