




**Towards a multi-messenger data science platform for the MMA era**



**Michael Coughlin (University  
of Minnesota), GRANDMA  
Collaboration Meeting,  
May 24, 2022**

# The History

O3 showed us how hard follow-up can be, and O4 is not going to get easier, as our interferometers remain at quite different sensitivities (plus the unfortunate situation with KAGRA's sensitivity).

However, because O3 was challenging, we built/adapted many, many quality tools both during the run and afterwards to analyze the results.

iCARE, gwemopt, simsurvey, POSSIS, gwemlightcurves/nmma, NIMBUS, various annotation pipelines, the GROWTH ToO-Marshal, the GROWTH Marshal, etc.

**Table 2.** Summary Statistics for O3, O4, and O5.

Run	BNS	NSBH	BBH
Median 90% credible area (deg <sup>2</sup> ) <sup>a</sup>			
O3	1672 <sup>+94</sup> <sub>-110</sub>	1970 <sup>+110</sup> <sub>-110</sub>	1069 <sup>+43</sup> <sub>-41</sub>
O4	1820 <sup>+190</sup> <sub>-170</sub>	1840 <sup>+150</sup> <sub>-150</sub>	335 <sup>+28</sup> <sub>-17</sub>
O5	1250 <sup>+120</sup> <sub>-120</sub>	1076 <sup>+65</sup> <sub>-75</sub>	230.3 <sup>+7.8</sup> <sub>-6.4</sub>
Median 90% credible co-moving volume (10 <sup>6</sup> Mpc <sup>3</sup> ) <sup>a</sup>			
O3	6.62 <sup>+0.97</sup> <sub>-0.97</sub>	44.1 <sup>+7.4</sup> <sub>-5.2</sub>	217 <sup>+23</sup> <sub>-16</sub>
O4	44.8 <sup>+6.4</sup> <sub>-6.5</sub>	191 <sup>+20</sup> <sub>-27</sub>	216 <sup>+16</sup> <sub>-20</sub>
O5	125 <sup>+21</sup> <sub>-12</sub>	448 <sup>+61</sup> <sub>-44</sub>	538 <sup>+23</sup> <sub>-24</sub>
Median luminosity distance (Mpc) <sup>a</sup>			
O3	176.1 <sup>+6.2</sup> <sub>-5.7</sub>	337.6 <sup>+10.9</sup> <sub>-9.6</sub>	871 <sup>+31</sup> <sub>-28</sub>
O4	352.8 <sup>+10.3</sup> <sub>-9.8</sub>	621 <sup>+16</sup> <sub>-14</sub>	1493 <sup>+25</sup> <sub>-33</sub>
O5	620 <sup>+16</sup> <sub>-17</sub>	1132 <sup>+19</sup> <sub>-23</sub>	2748 <sup>+30</sup> <sub>-34</sub>
Sensitive volume (10 <sup>6</sup> Mpc <sup>3</sup> ) <sup>ab</sup>			
O3	17.5 <sup>+1.4</sup> <sub>-1.3</sub>	101.1 <sup>+6.4</sup> <sub>-6.1</sub>	1047 <sup>+50</sup> <sub>-49</sub>
O4	109.0 <sup>+6.7</sup> <sub>-6.5</sub>	558 <sup>+26</sup> <sub>-26</sub>	4450 <sup>+130</sup> <sub>-130</sub>
O5	590 <sup>+29</sup> <sub>-28</sub>	2787 <sup>+89</sup> <sub>-87</sub>	19950 <sup>+310</sup> <sub>-310</sub>
Annual number of detections <sup>cd</sup>			
O3	5 <sup>+14</sup> <sub>-5</sub>	13 <sup>+15</sup> <sub>-9</sub>	24 <sup>+18</sup> <sub>-12</sub>
O4	34 <sup>+78</sup> <sub>-25</sub>	72 <sup>+75</sup> <sub>-38</sub>	106 <sup>+65</sup> <sub>-42</sub>
O5	190 <sup>+410</sup> <sub>-130</sub>	360 <sup>+360</sup> <sub>-180</sub>	480 <sup>+280</sup> <sub>-180</sub>

## The Vision

We have also seen our team evolve (both in terms of people coming and going, as well as career stages). Even during these 1+ year runs, these aspects change as well.

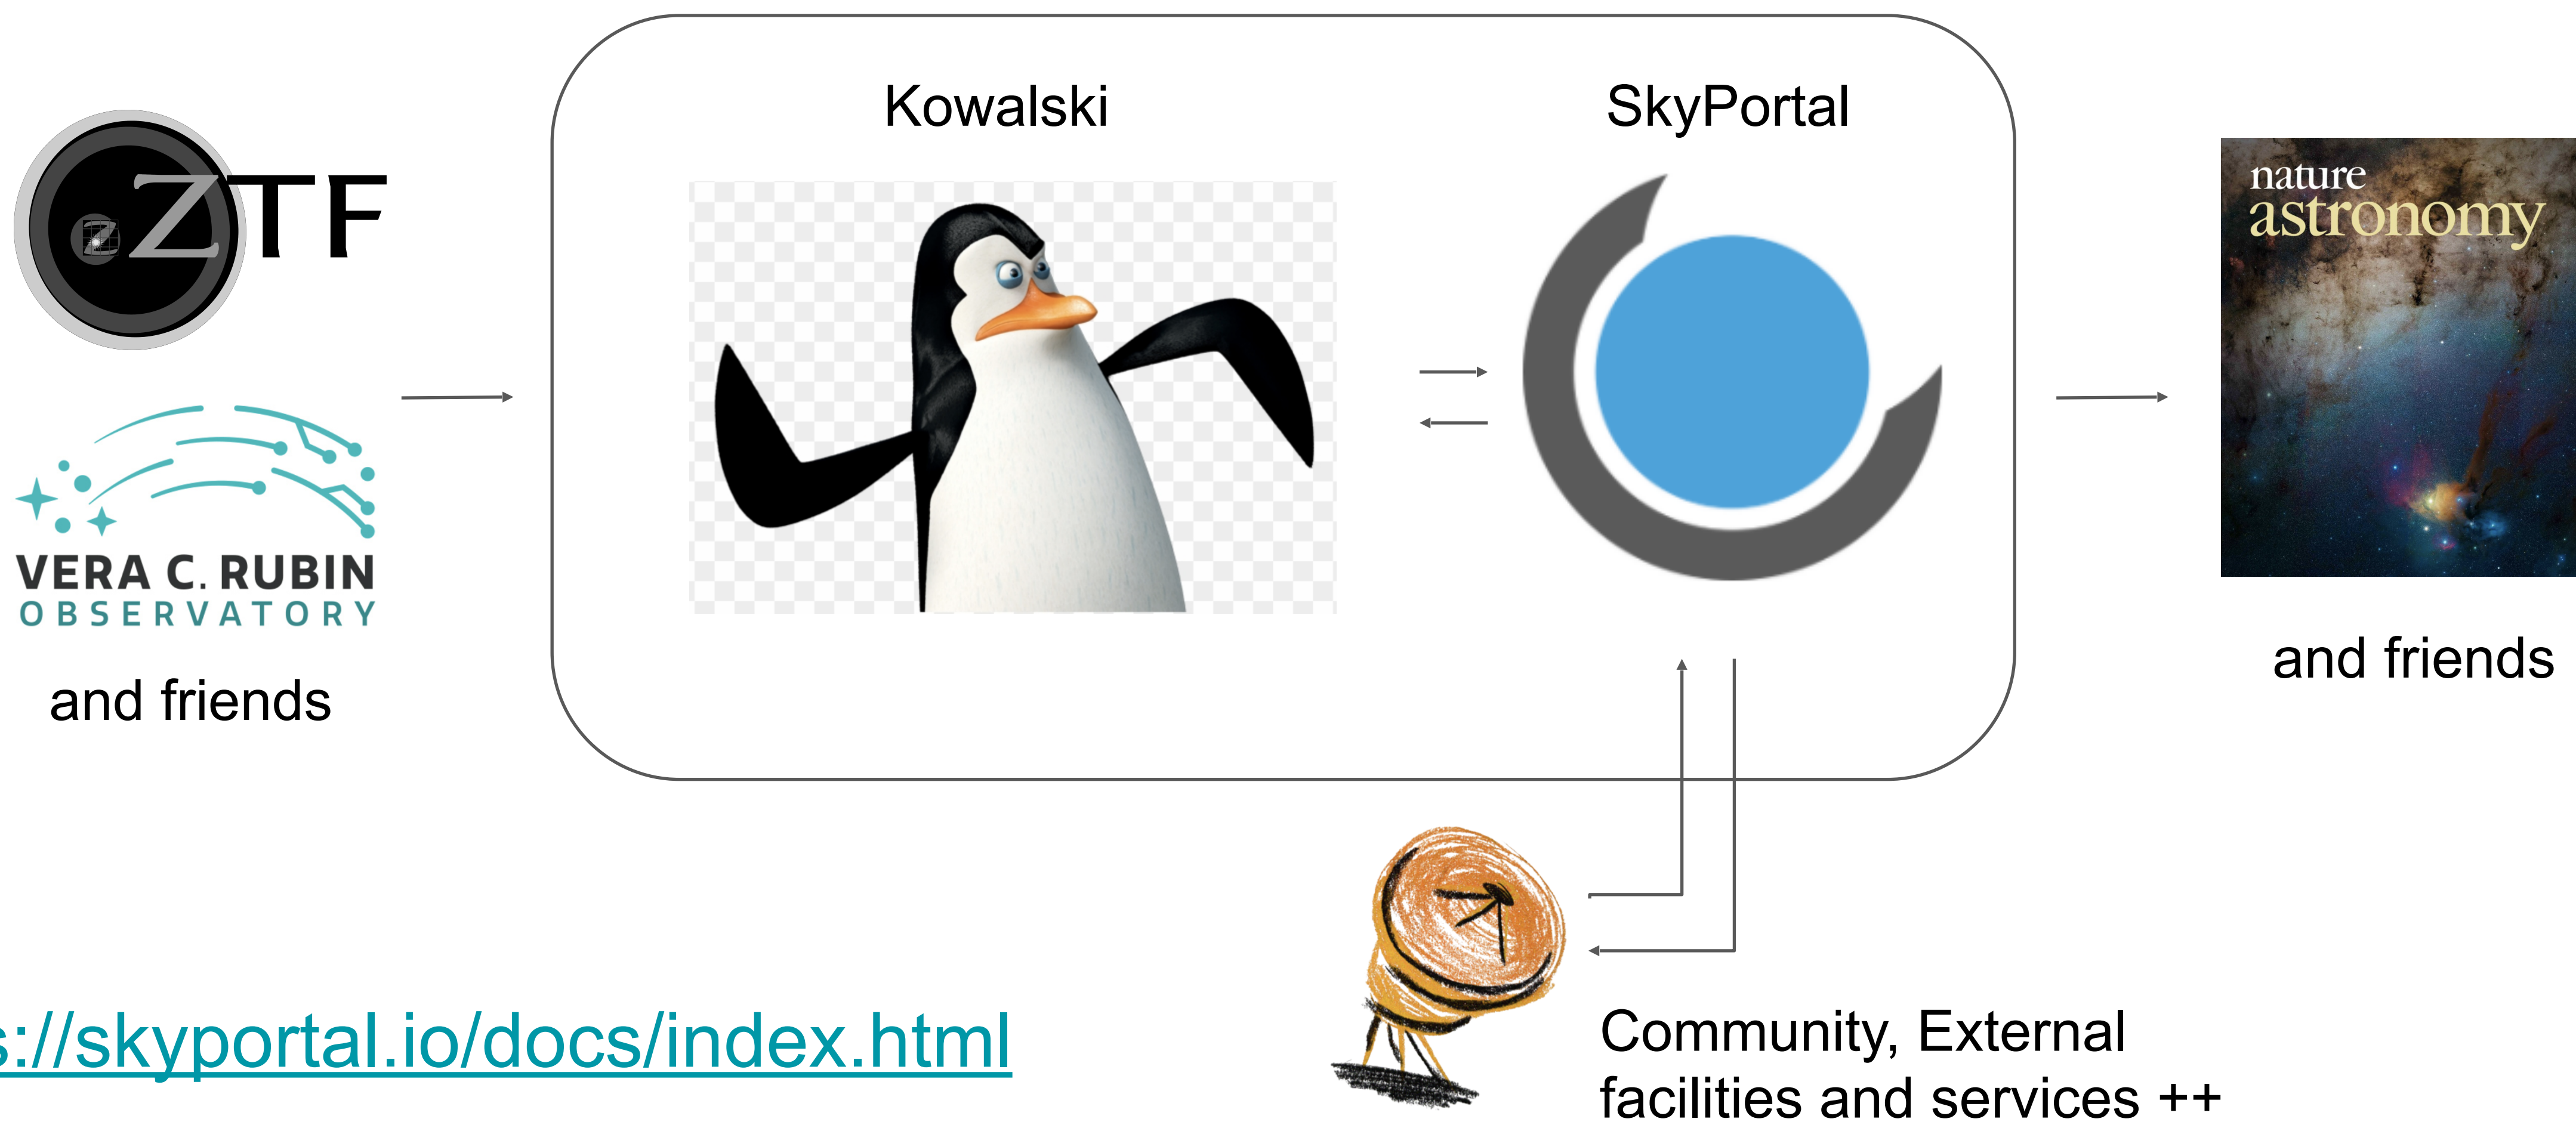
For this reason, we have set out to use SkyPortal as our one-stop-shop where many of these analyses should be available. This enables us to both save ourselves both in time and energy; the run is grueling (but obviously potentially incredibly rewarding) and limiting single points of success\* is important. Even if in principle more than one person knows how to do something, it can be often the case that basically one person knows how to do it, and that's a tough spot to be in.

\* A Patrick Brady special



# SkyPortal: Schematic overview

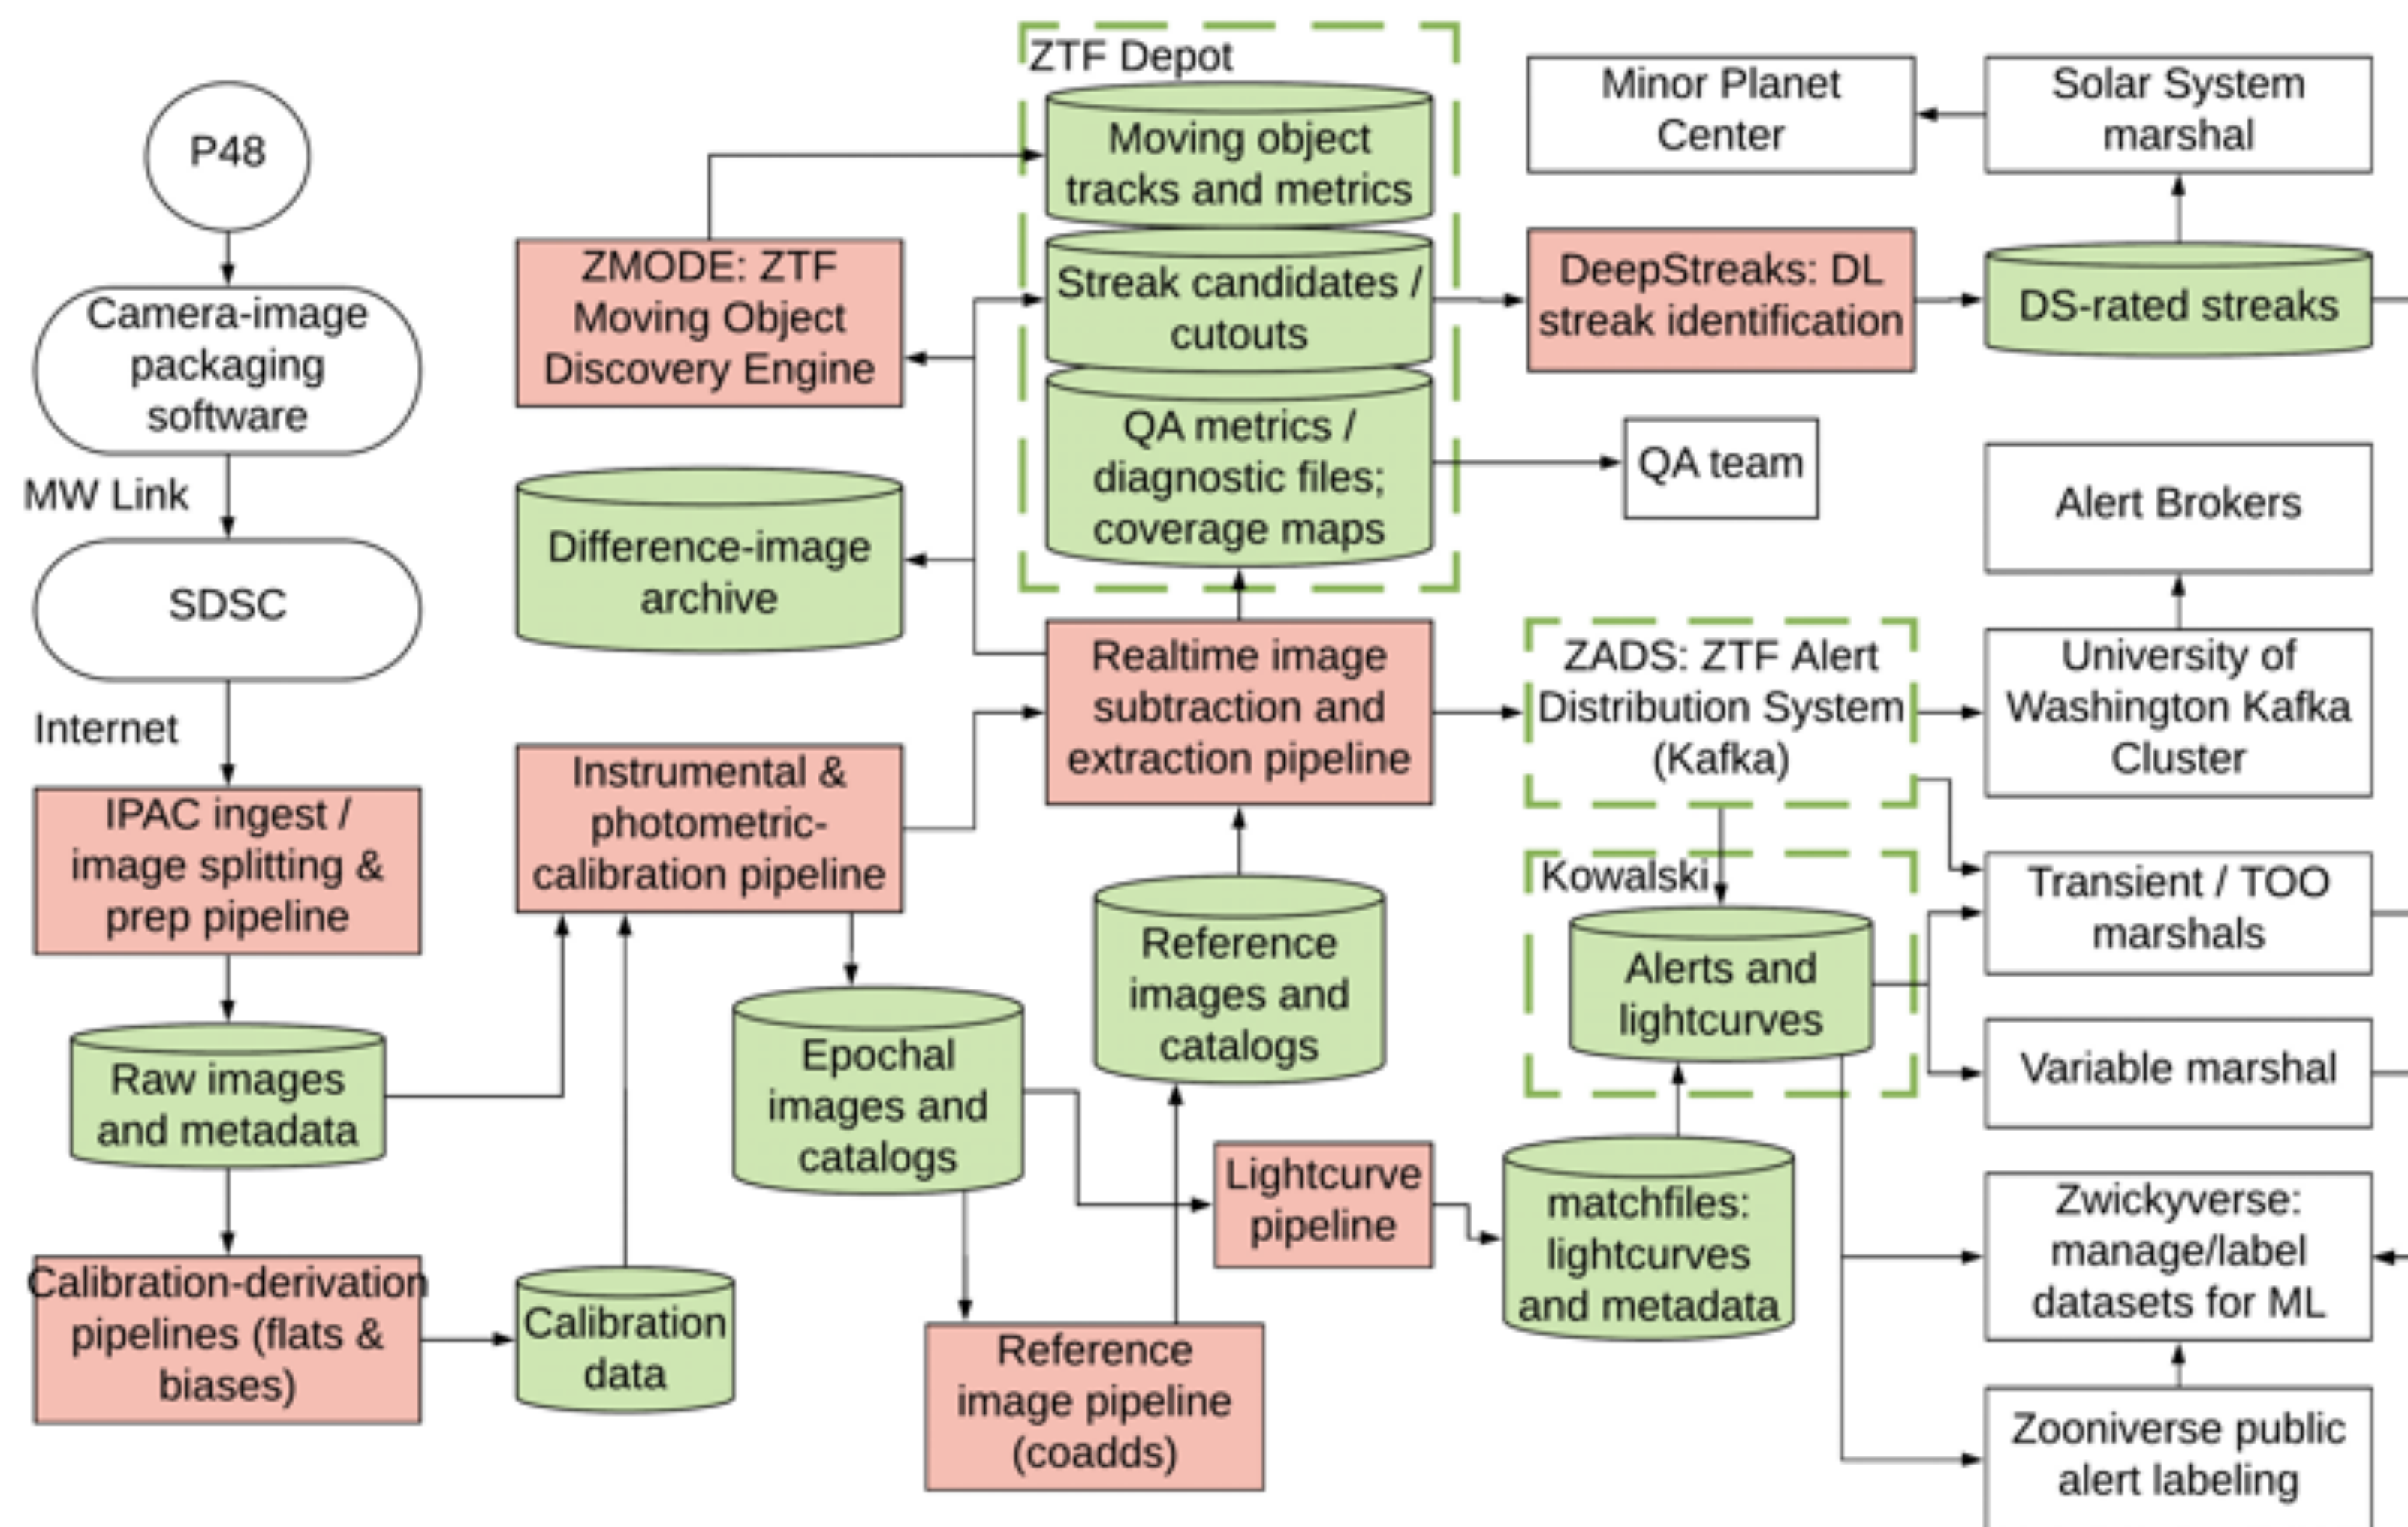
Observe → Mine/Discover → Study and Characterize → Profit!



<https://skyportal.io/docs/index.html>



# ZTF-I: data/processing flow



ZTF acts as a discovery engine. Discoveries are followed-up using a wide range of instruments (including ZTF itself)

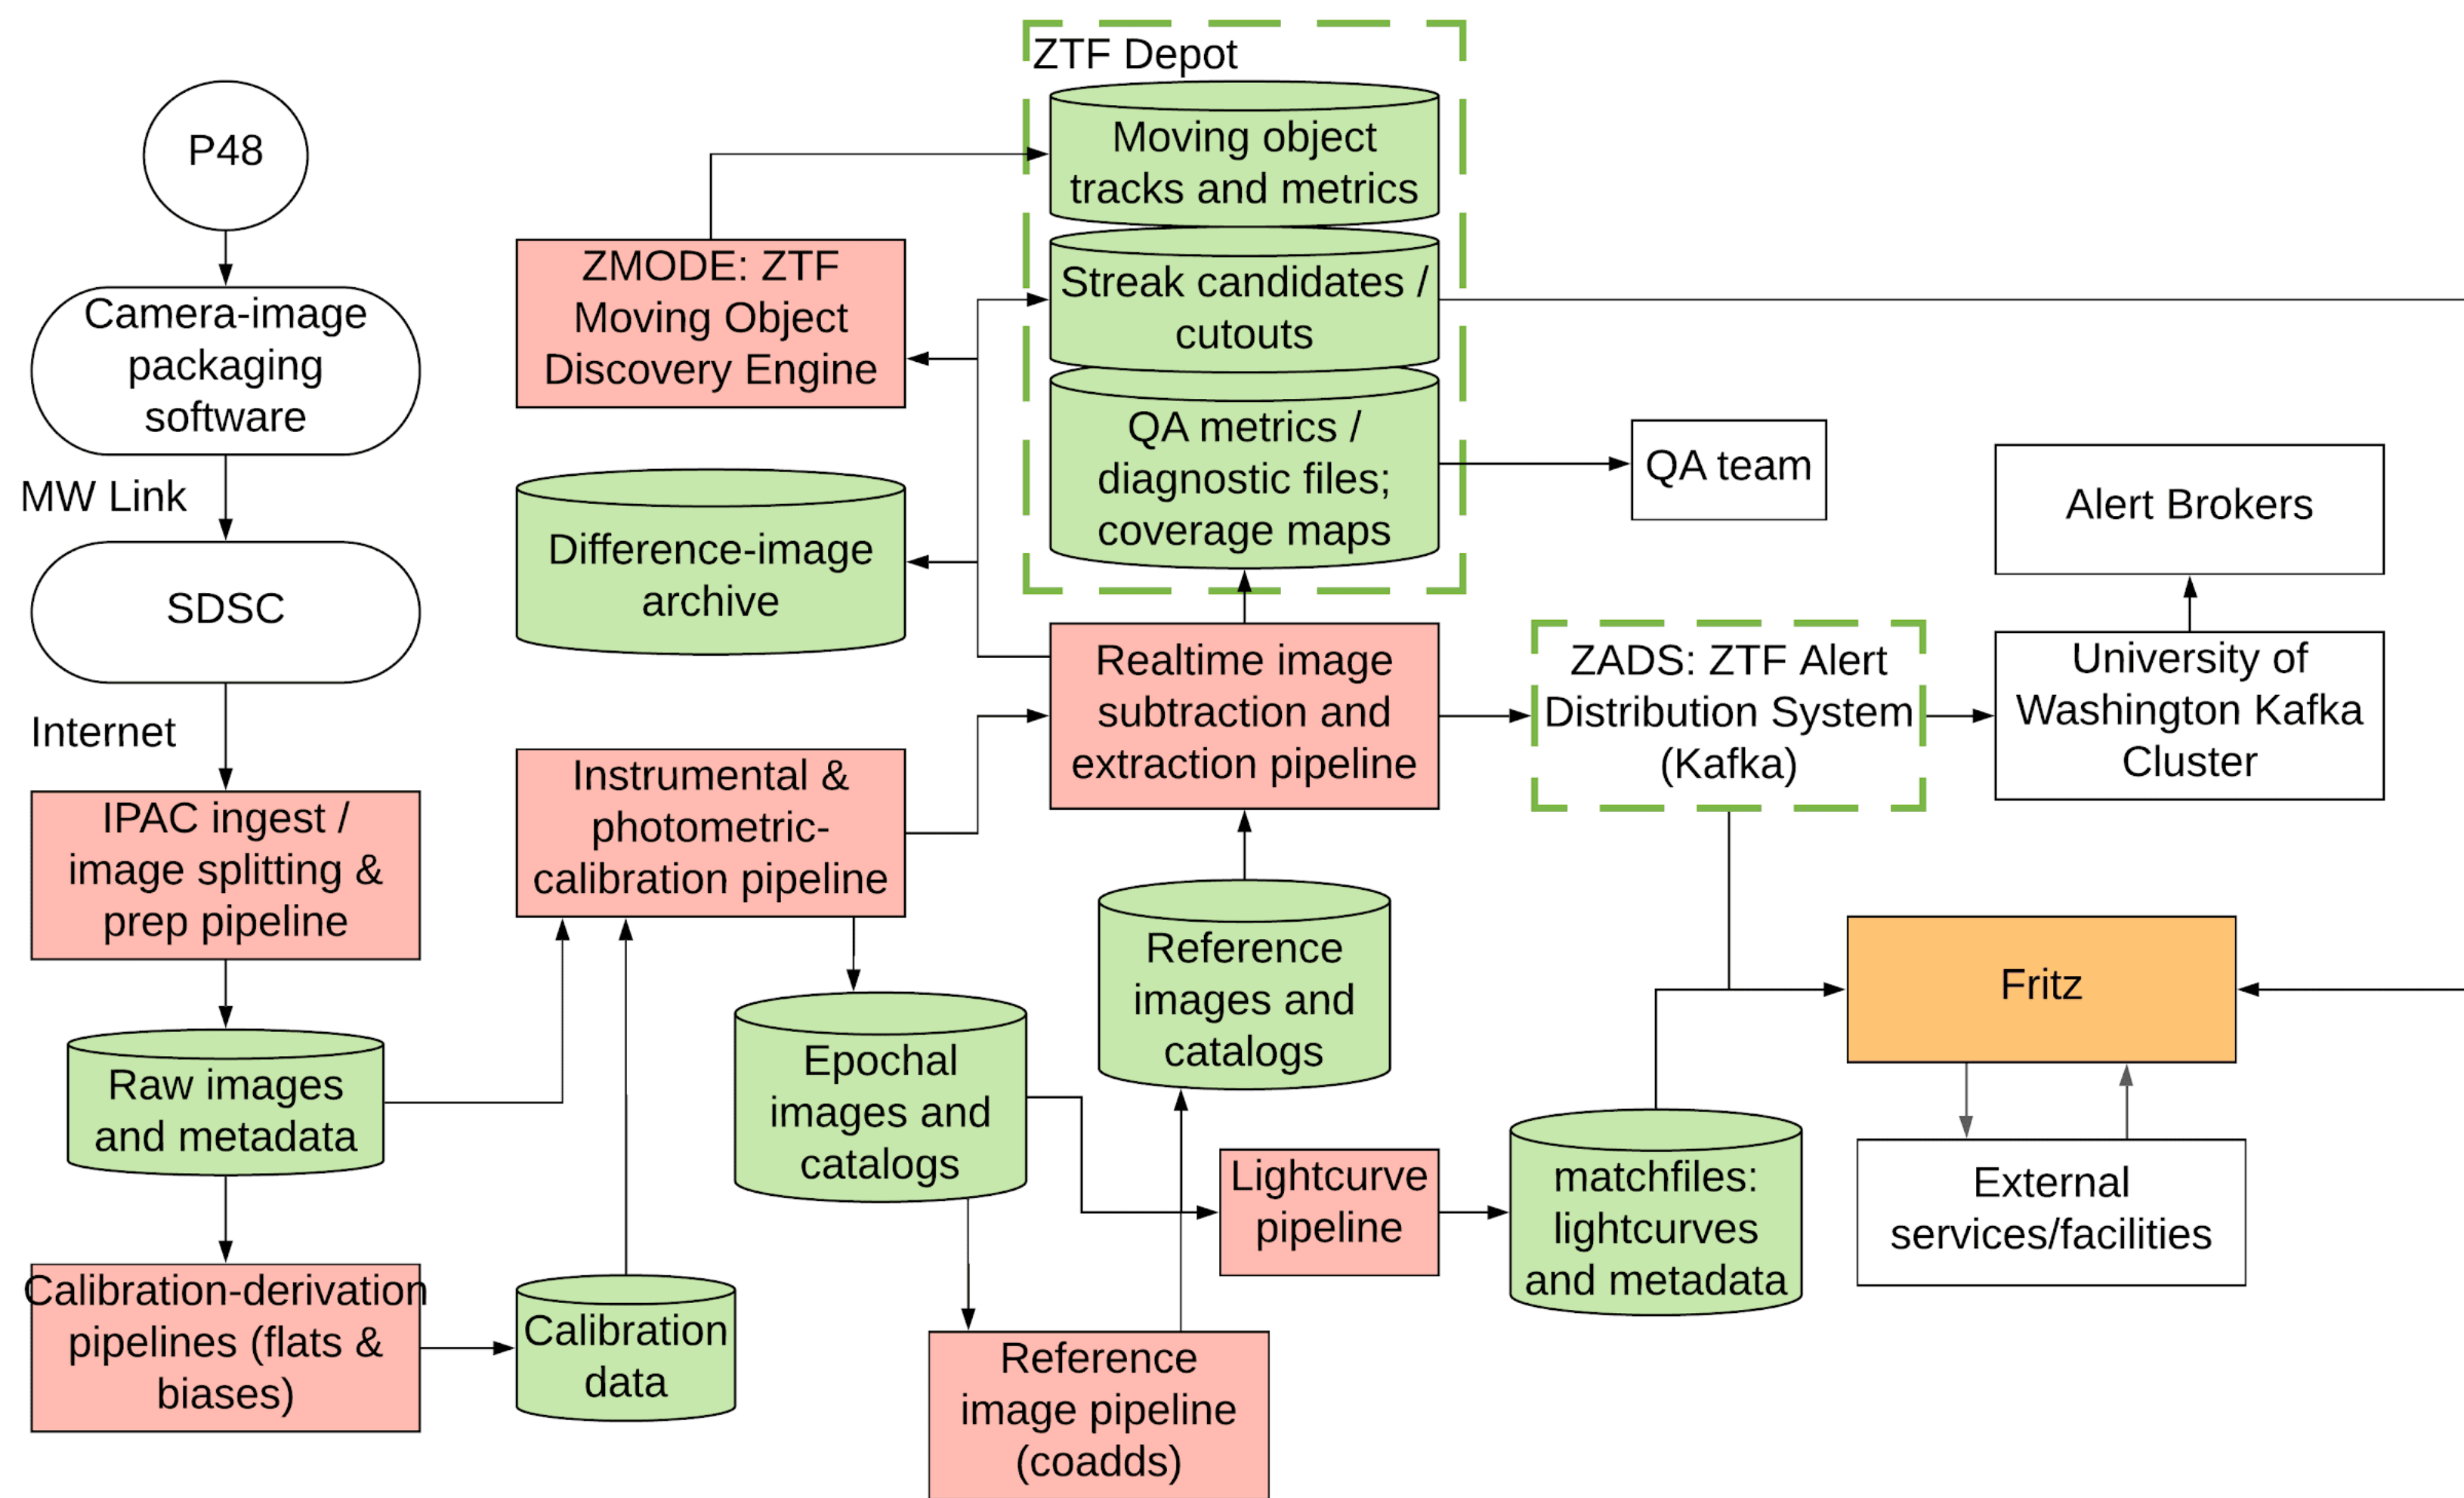
Single night	8h40m
# on-sky exposures	~700
Raw image data	~1 TB
Real-time data products	~4 TB
# unvetted $5\sigma$ alerts	$\sim 10^5 - 10^6$
# ML-vetted alerts	$\sim 10^3 - 10^5$
# unvetted streaks	$\sim 10^4 - 10^6$
# ML-vetted streaks	$\sim 10^2 - 10^3$

Nominal survey	3 x 260 n
Volume of data products	~3.2 PB
Volume of ref images	~60 TB
# CCD quad ref images	$\sim 2.8 \times 10^5$
Volume of matchfiles	~50 TB
# matchfiles	$\sim 2.8 \times 10^5$
# single-epoch PSF-fit source measurements	~800 B
# single-epoch aperture source measurements	~230 B

See Masci+ 2019



# ZTF-II: data/processing flow



ZTF acts as a discovery engine. Discoveries are followed-up using a wide range of instruments (including ZTF itself)

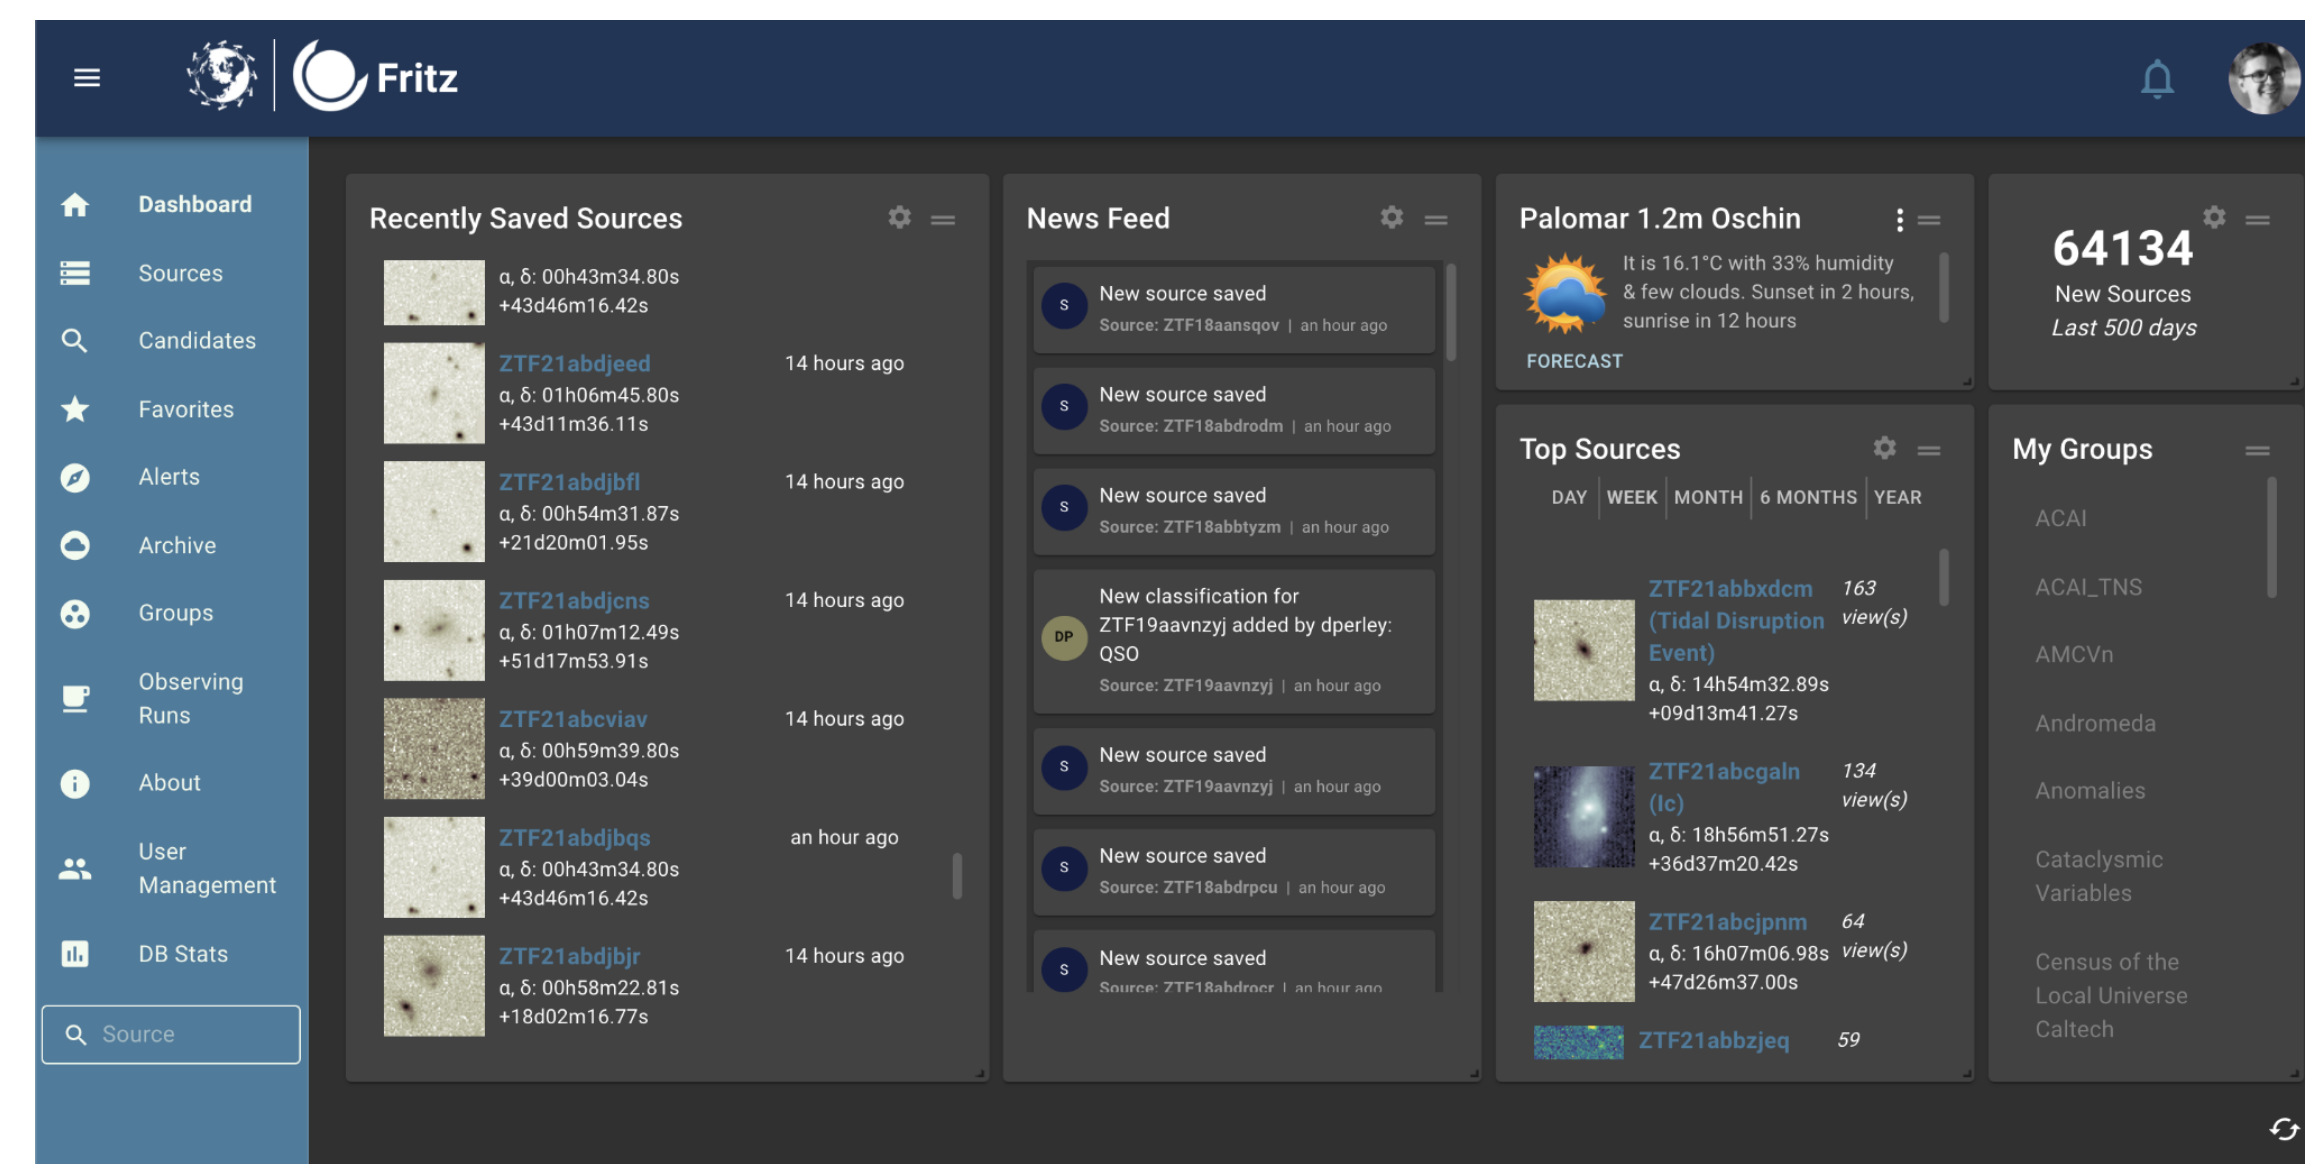
Single night	8h40m
# on-sky exposures	~700
Raw image data	~1 TB
Real-time data products	~4 TB
# unvetted $5\sigma$ alerts	$\sim 10^5 - 10^6$
# ML-vetted alerts	$\sim 10^3 - 10^5$
# unvetted streaks	$\sim 10^4 - 10^6$
# ML-vetted streaks	$\sim 10^2 - 10^3$

Nominal survey	3 x 260 n
Volume of data products	~3.2 PB
Volume of ref images	~60 TB
# CCD quad ref images	$\sim 2.8 \times 10^5$
Volume of matchfiles	~50 TB
# matchfiles	$\sim 2.8 \times 10^5$
# single-epoch PSF-fit source measurements	~800 B
# single-epoch aperture source measurements	~230 B

# SkyPortal

- Open source (free to use, modify, and distribute)
- Scalable, [API-first](#) system, with [fine-grained access control](#)
- Multi-survey data archive and alert broker
- Interactive, mobile-friendly collaborative platform for transient, variable, and Solar system science cases
- Workhorse for ML applications: classification and labeling at scale
- Follow-up observation management: robotic and classical facilities
- Well-tested, extensive docs, CI/CD

Initiated in Feb 2020  
Beta up in Sep 2020  
MVP live in Nov 2020





Stability/Scaling


# Roadmap

(In brief)

MMA (04)

3rd Party Analysis

# Stability / Scaling



# Twitter is over capacity.


Please wait a moment and try again. For more information, check out **Twitter Status**.

[Bahasa Indonesia](#) [Bahasa Melayu](#) [Deutsch](#) [English](#) [Español](#) [Filipino](#) [Français](#) [Italiano](#) [Nederlands](#) [Português](#) [Türkçe](#)

[Русский](#) [हिन्दी](#) [日本語](#) [简体中文](#) [繁體中文](#) [한국어](#)

© 2012 Twitter [About](#) [Help](#) [Status](#)





“Debugging performance issues is really hard. But it’s not hard due to a lack of data; in fact, the difficulty arises because there is too much data..It’s easy to confuse causes with symptoms”



Building the plane while we fly it...

...and (trying to) squash bugs big and small.



# MMA Infrastructure for O4

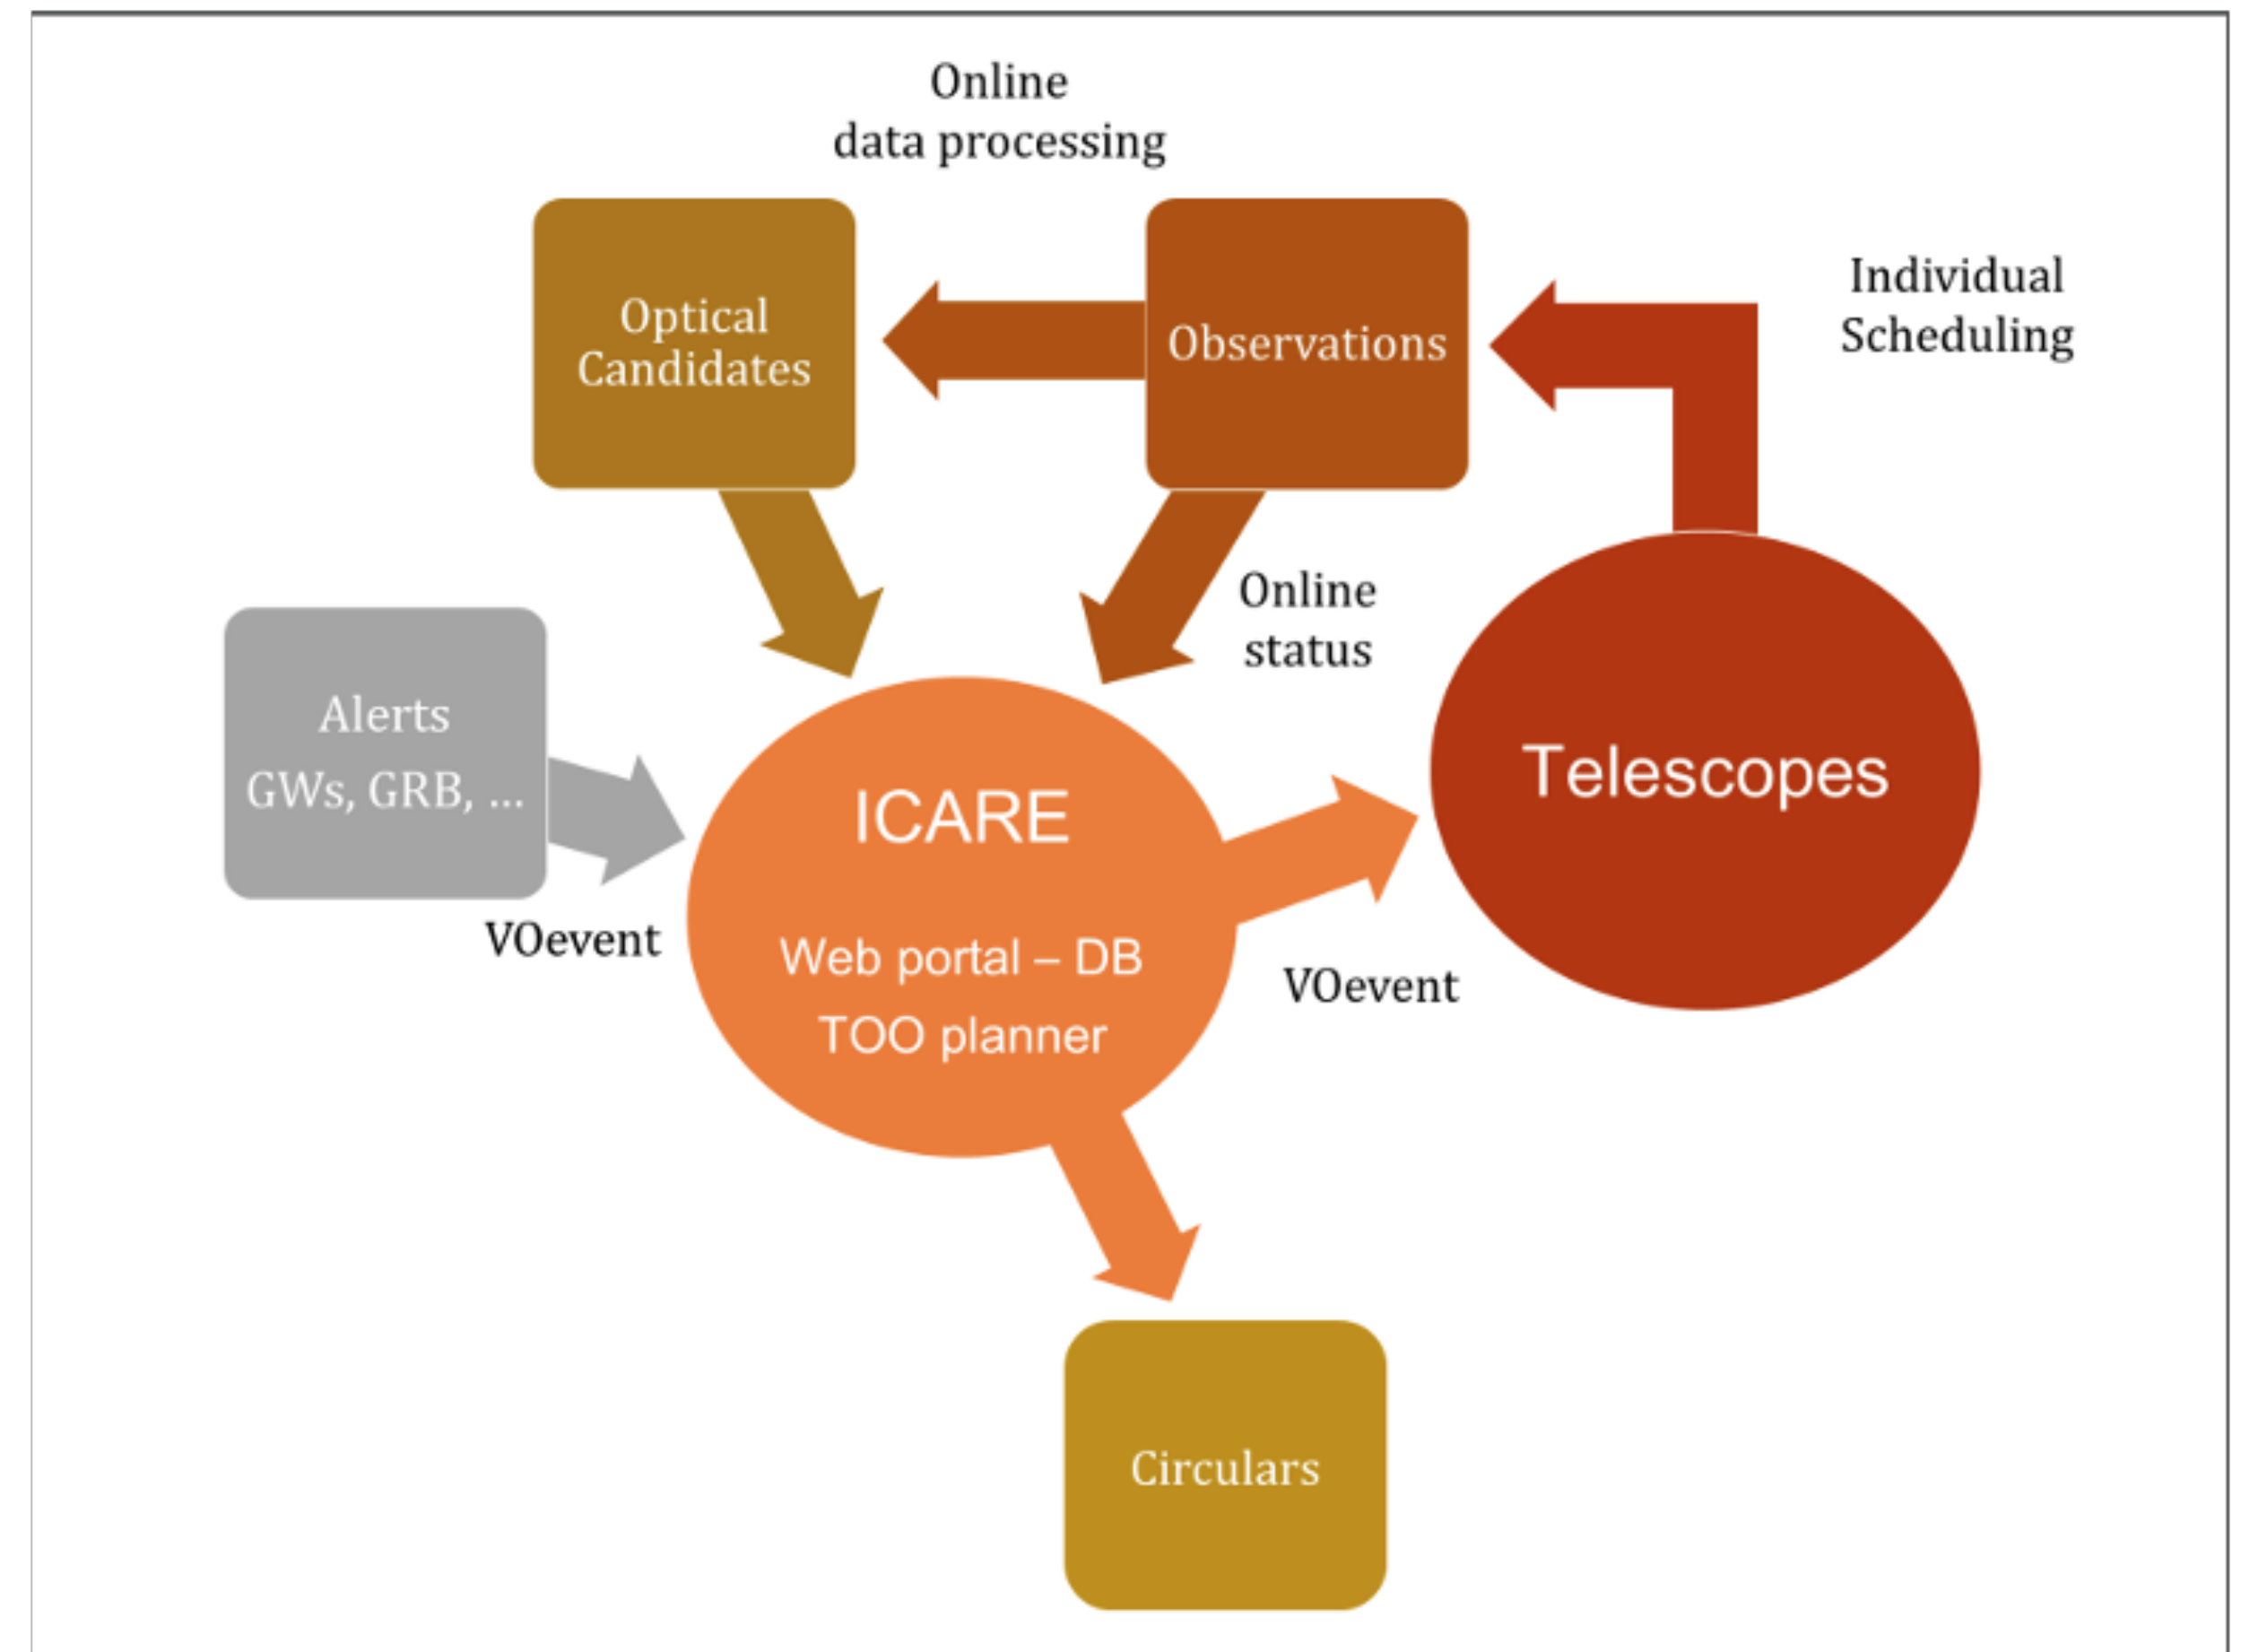


# Inherits from two very successful projects during O3 and beyond: GRANDMA's iCARE and GROWTH's ToO Marshal

## Events

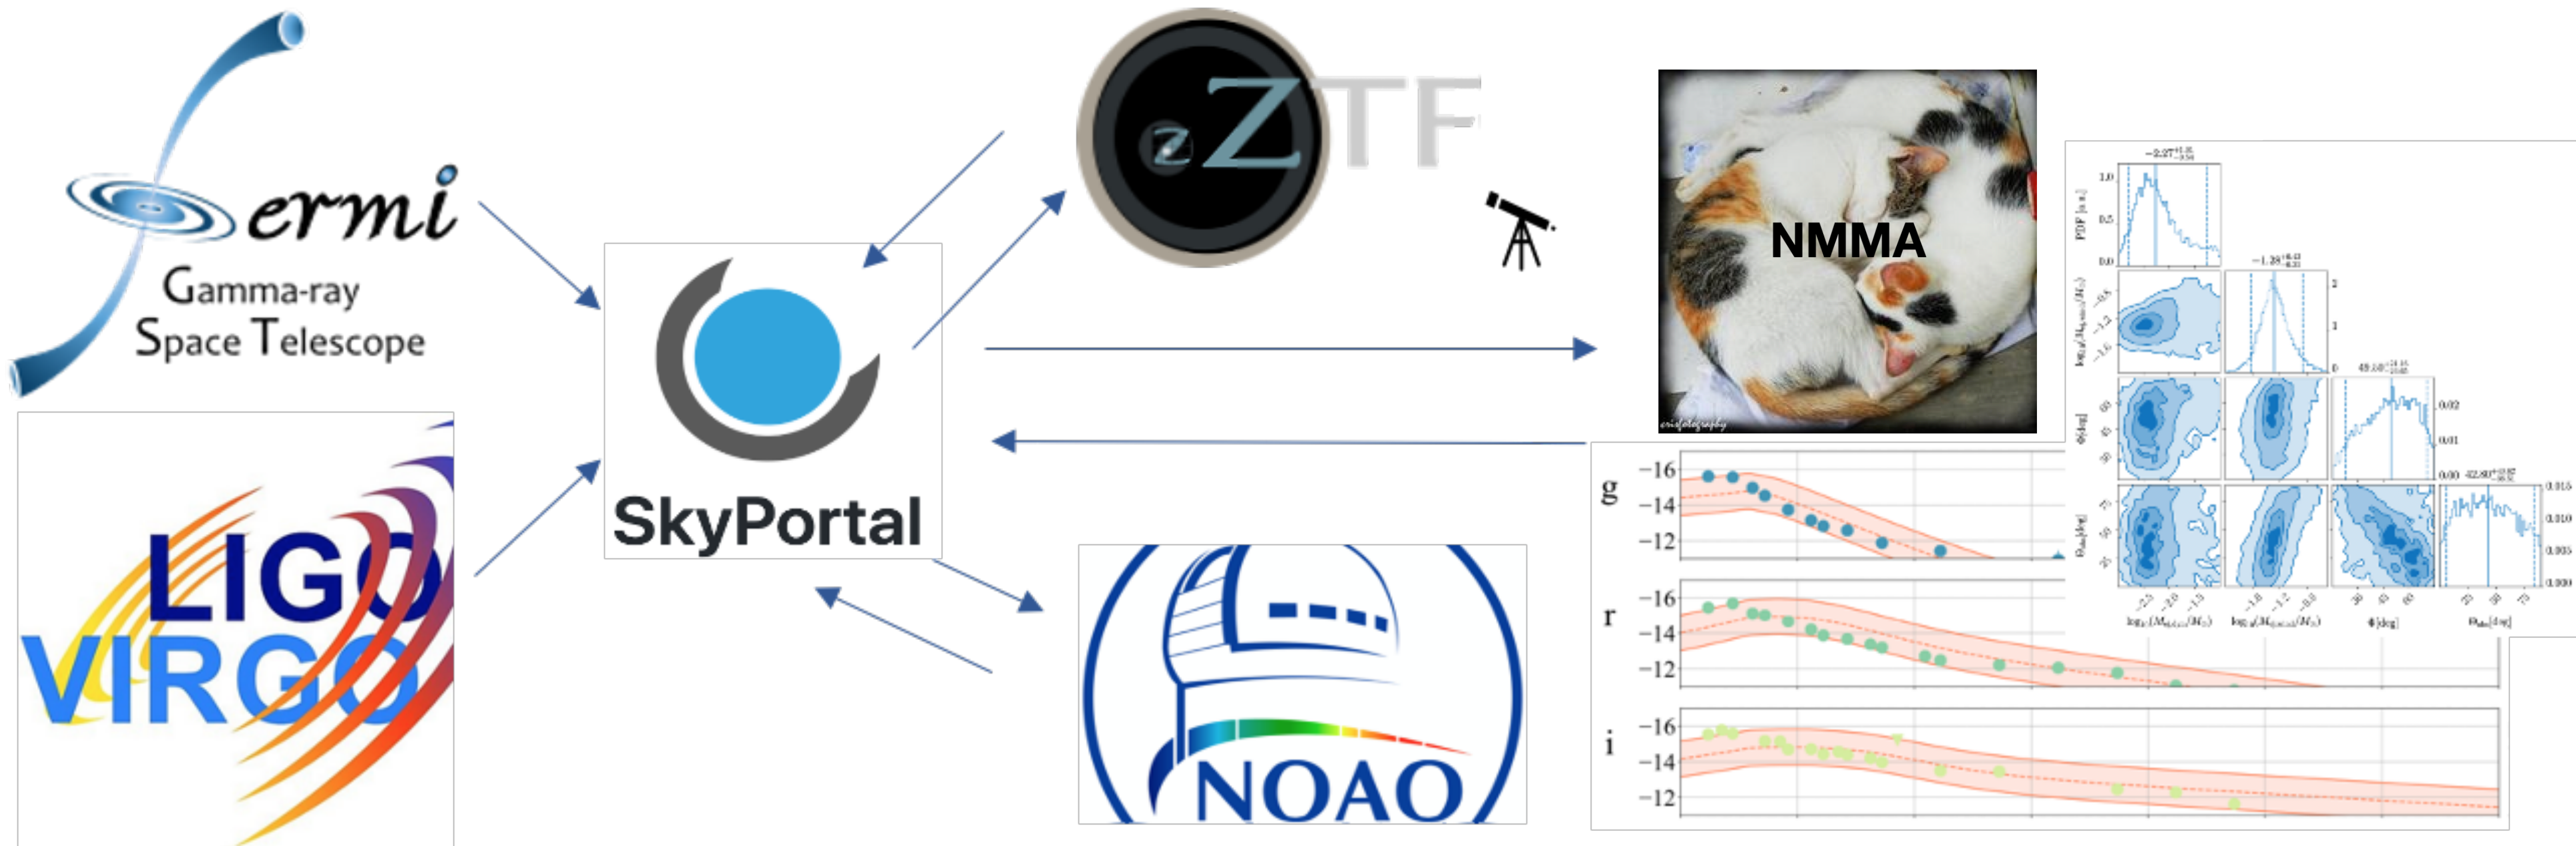
AMON  AllSky  BBH  BNS  Fermi  GRB  GW  HighMass  IMBH  
 LVC  MDC  MassGap  NSBH  Terrestrial  long  retracted  short  
 transient

210511 20:51:24 (18 hours ago)	Fermi GRB
<a href="#">210511 11:26:39 (a day ago)</a>	Fermi long GRB
210511 10:58:57 (a day ago)	Fermi long GRB
210510 19:20:01 (a day ago)	Fermi short GRB
210509 13:53:40 (3 days ago)	Fermi GRB
210509 12:15:04 (3 days ago)	Fermi short transient
210508 22:27:35 (3 days ago)	Fermi GRB
210508 21:37:25 (3 days ago)	Fermi short transient
210508 21:30:44 (3 days ago)	Fermi short transient
210507 19:01:51 (4 days ago)	Fermi GRB



Ahumada et al. 2105.05067, Anand and Coughlin et al. 2009.07210, Andreoni and Goldstein et al. 1910.13409, Antier et al: 1910.11261, 2004.04277, Coughlin et al.: 1907.12645, etc.

# A vision for the O4 workflow



Goal: Inform follow-up decisions

Goal: Remove difference between code we use to vet candidates in low latency and code we use to do science

## What do I hope is available by O4?

- Light curve fitting: fit to supernova, kilonova, GRB, etc. models with the click of a button (led by Berkeley, support from UMN / OCA for NMMA)
- Follow-up prioritization suggestion (led by Caltech/Ari with ORACLE code)
- Creation of space-time MOCs for ease of internal and external observation distribution (led by UWM)
- Improved scheduling with M4OPT (led by Goddard)
- Upperlimits/efficiency assessment: triggering of simsurvey/NIMBUS so we can quote sensible efficiencies in GCNs/papers (led by UWM, Stockholm, UMN and others). Are we tooled up for things like using our own survey limits as priors? GW170817 as a prior? etc?
- Shifts: fine-grained tooling to receive notifications of events / observations in time-zone appropriate ways (led by OCA/IJCLAB)
- Rapid/large-scale time-series interaction, i.e. data sets like from KPED or Chimera (led by Berkeley)



## What do I need from you?

- **All:** Missing features in the to-do list that you remember / expect you will find essential for O4
- **Those with existing observing programs (especially on those with APIs accessible through SkyPortal):** Try to trigger your programs through the website and work with us to fix any bugs / missing features. I know it will probably fail the first time or two, but this will allow us to be set by the time we need all hands on deck.
- **Those with existing code you use to interact in cool ways with SkyPortal but are otherwise poorly/tersely documented:** Take a few minutes to add some text and code to the SkyPortal Markdown docs (<https://github.com/skyportal/skyportal/tree/master/doc>)
- **Those with existing code you use that should probably be enabled within SkyPortal:** Post an issue / talk to me about what it does / check the big board
- **Those with some hiring power:** Send us your undergrads / summer students willing to learn some JS / React.js, we are sorely in need of some front-end help.

## What do I need from you?

- **All:** Missing features in the to-do list that you remember / expect you will find essential for O4
- **Those with existing observing programs (especially on those with APIs accessible and working)**
  - Note: None of this requires significant software input from your side!**
  - There is a pretty steep learning curve for joining the dev team, so we don't recommend that all the users necessarily become developers. That said, it is worthwhile to have at least one student or researcher devote the time to learn the inner workings by joining the team and contributing code. This helps the group get their own tools in faster and the way they need them, and it's also good for the developer that would learn a lot of useful technical skills along the way, such as industry best practices for open source development (code review, automated testing, etc).**
- **Those with existing code you use that should probably be enabled within SkyPortal:** Post an issue / talk to me about what it does / check the big board
- **Those with some hiring power:** Send us your undergrads / summer students willing to learn some JS / React.js, we are sorely in need of some front-end help.

the website  
fail the first  
on deck.  
portal but are  
the text and  
[portal/tree/](#)

Experience so far...




# Dev/Ops Personnel

Currently: 7 code contributors are mostly those donating their time



🙏 Michael, Jada (UMN),  
Guy, Josh, Stefan (Berkeley),  
Theo, Thomas (OCA/GRANDMA)

We appreciate bug 🐛 reports but we *really* need code contributors 🧑💻



# How to Engage

- Document Issues on the SkyPortal GitHub
- Contribute code via pull requests - we're happy to onboard you!
- Lots of (unique) learning opportunities in development and operations

# Lesson Learned: open source is

- Fritz/SkyPortal/Kowalski is open source
- A huge part of devs is volunteer labor of love
- While the core dev team is relatively small (in terms of effective person-hours), dozens of people have [contributed](#) meaningfully
- Leveraging what GitHub/OSS has to offer
  - Issues to track bugs and feature requests
  - PRs + thorough code review
  - GitHub Actions as the CI/CD platform
  - Don't argue about style, enforce pre-commit hook (black, flake8, eslint...)
- Project management tool: ZenHub
  - Should be as close as possible to GH



# Lesson Learned: team and community is

- Extraordinary individuals with a broad range of expertise
  - Respectful + supportive environment (that naturally extends to mutual admiration) allows the team to go far (check out [this article](#) for the spirit)
  - Critical but fair and open-minded review of ideas and code allows to iterate fast, converging on better solutions
  - Staying in sync: Slack + Daily 15-min stand-ups + weekly 1h meetings
- Constructive community feedback is essential for success
  - Beta-testers!
  - Enormously useful for finding/fixing bugs and implementing new features
  - Need to be clear about the communication channels: a dedicated Slack channel for smaller issues + GH issue templates for larger stuff
  - Critically important: prioritization + clear big picture for the project



**François Chollet**  @fchollet · Apr 12

Doing everything your users request is the shortest path to ruining your product -- especially when these users come from a legacy product and just want the new thing to be like the old thing

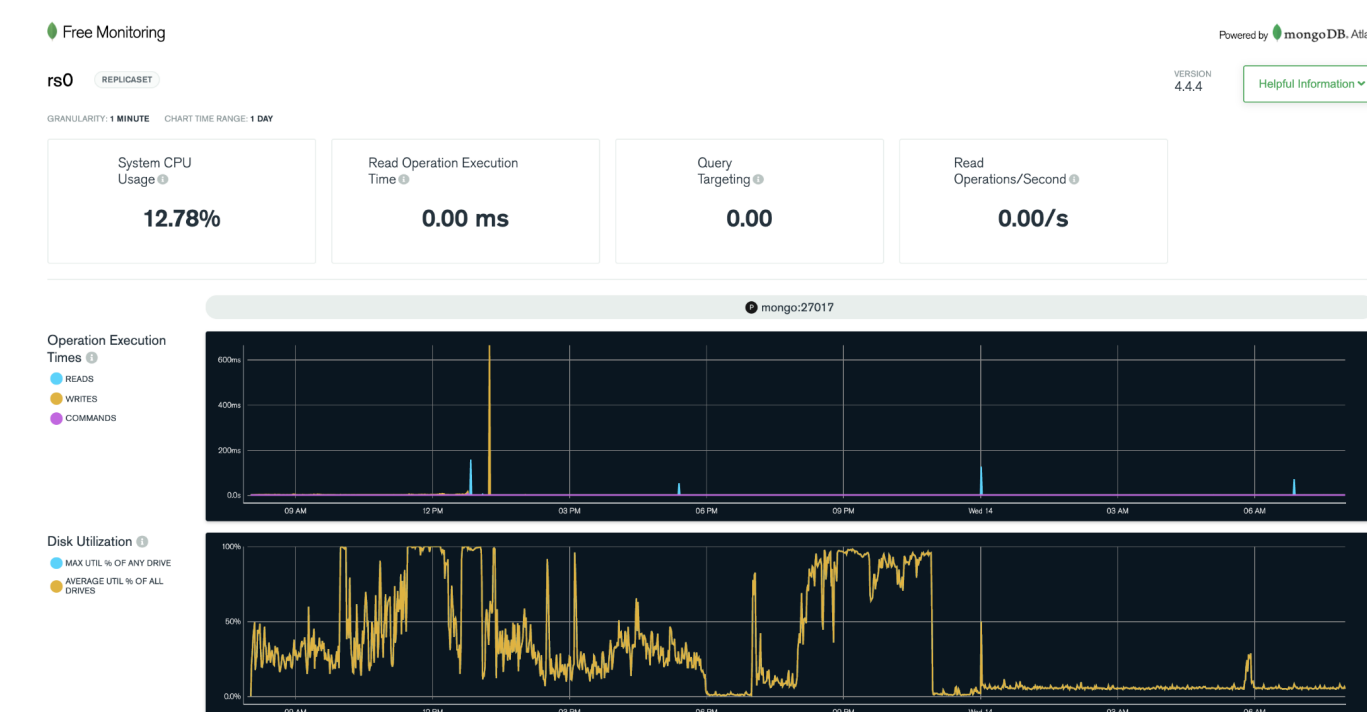
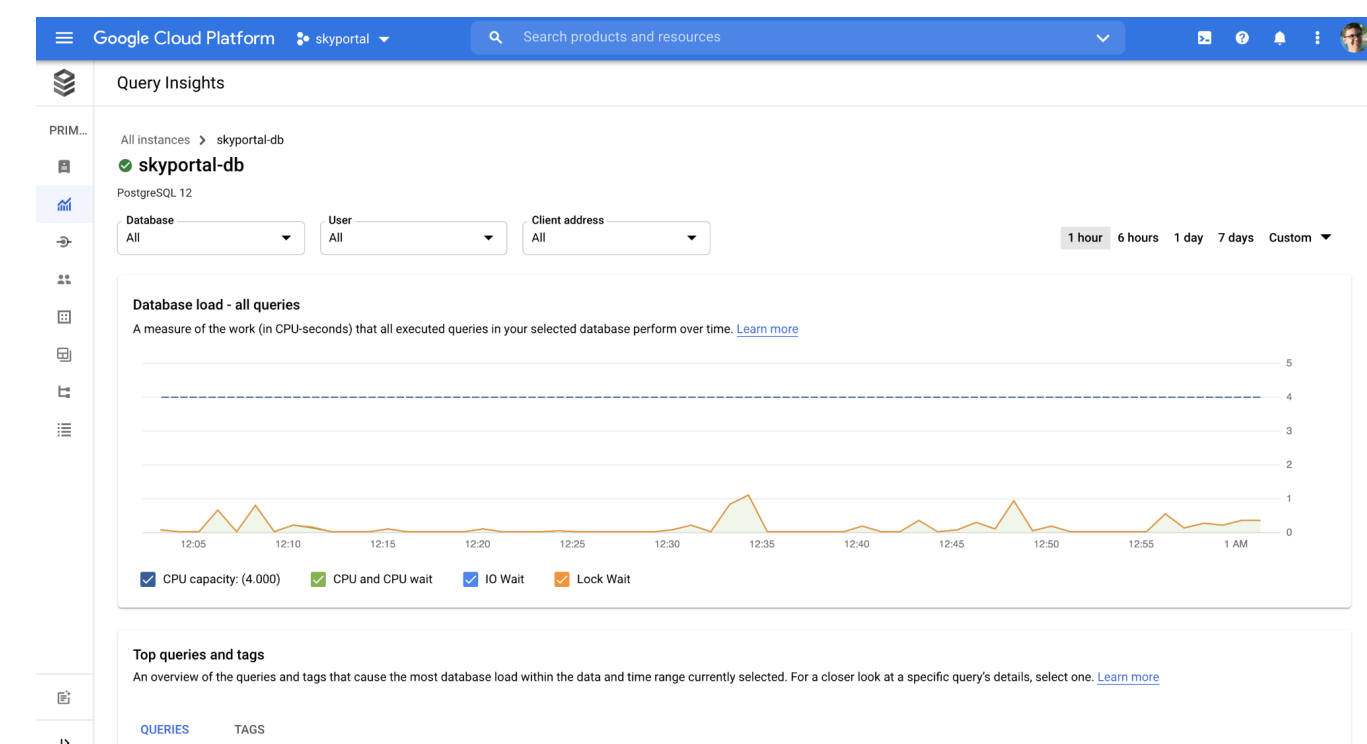


# Lesson Learned: testing

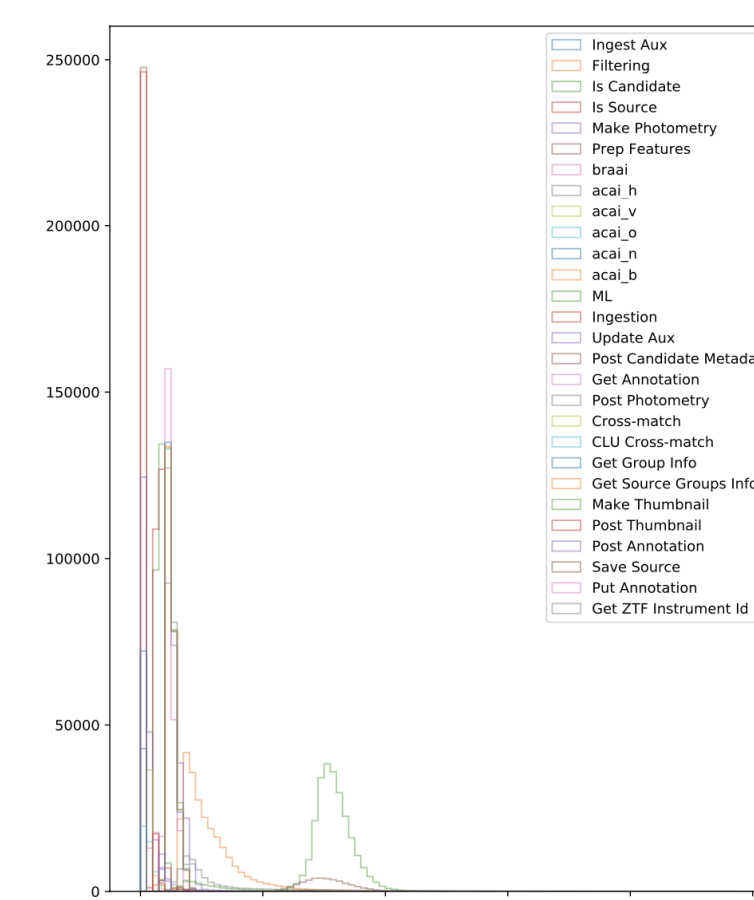
- You are not testing your code enough
  - No, you're not!
  - From unit testing to integration testing through API and frontend, every bit is essential
  - Is that docker image still building from scratch? Note the word "continuous" in CI
  - Never underestimate the scale of a disaster that six innocently-looking lines of code can bring
  - Database migrations should be tested both ways - roll-backs are more common than we'd wish
  - Understand (and embrace!) flakiness
- Staging environment
  - Helps catch a lot of bugs before they have a chance to reach production, e.g. innocently-looking migrations that can take forever

# Lesson Learned: production is really hard

- Even harder is to deliver updates/new features to prod
- Testing *is* your friend, but it won't catch everything that can happen
- Running a subset of the test suite on a read-only replica of the prod db
- Weekly (at least) deployment to prod
- Resilient infrastructure for deployment
- Monitoring the performance of the different components
- Query Insights on the GCP
- API endpoint response times, temporal evolution



Kowalski Performance Report  
Start Date: 04 / 12 / 2021  
End Date: 04 / 13 / 2021





Thank you!