ZTF Operations Cost Review
Science Data System

Frank Masci, David Imel, Steve Groom, Ben Rusholme

& the IPAC-Caltech ZTF Team

December 20, 2018

HEISING-SIMONS FOUNDATION

Outline

- Timeline
- Objectives of operations phase for Data System
- Overview of Data System operations model
- Automated and manual tasks (daily and long-term)
- Staffing plan
- Budget
- Continued development tied to baseline design
- Risks

Data System is currently in "steady-state" survey operations

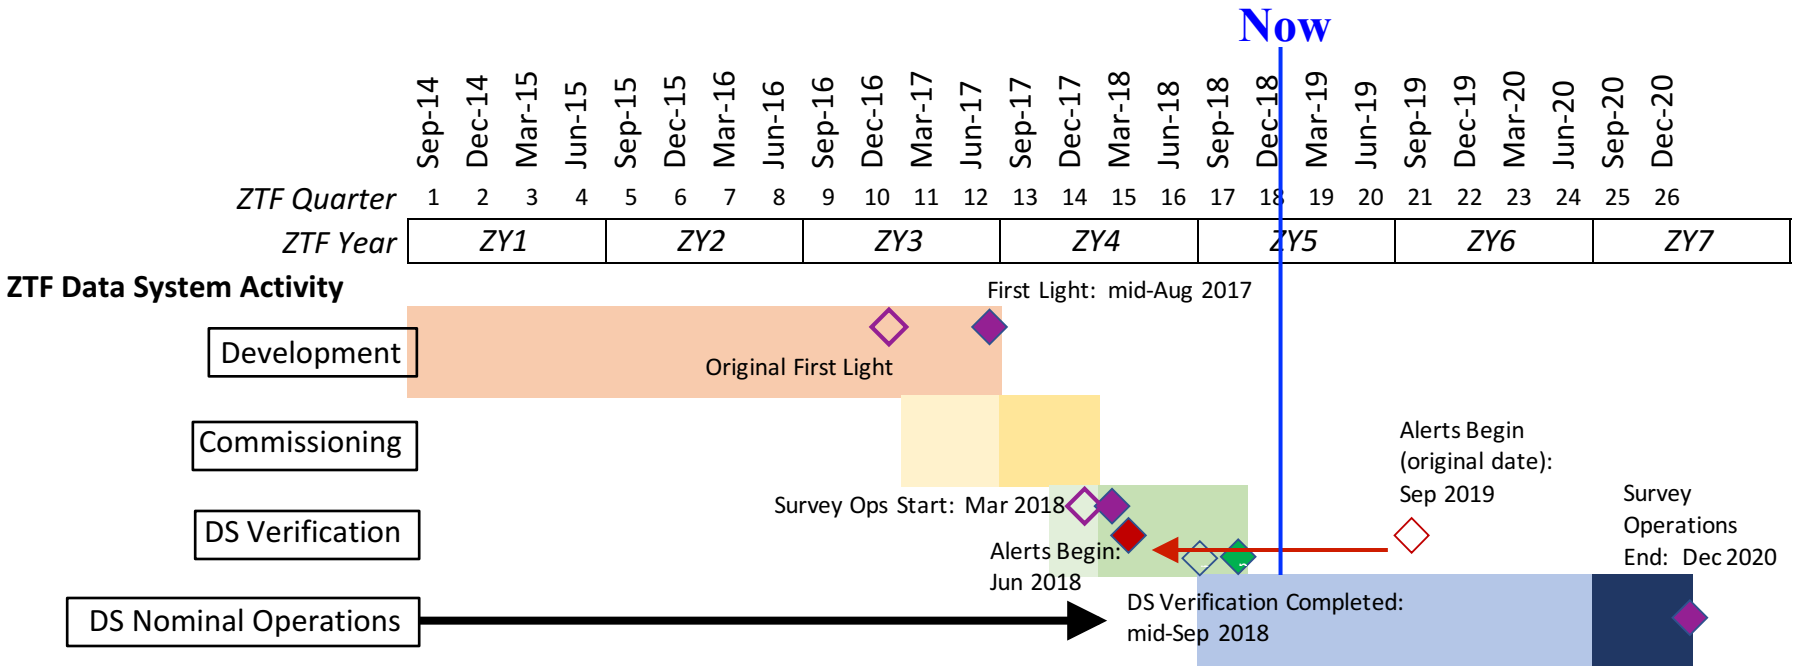

- Science operations commenced on March 17, 2018
- Public Alerts officially commenced on June 4, 2018
- Data system verification period ended in mid September 2018
- We are now in nominal operations

Objectives of Operations Phase

- Support and sustain data processing and archiving until end of 2020.
- Support public data releases.
- Continue to refine calibrations and automated classifiers.
- Monitor baseline production system for inefficiencies and optimize where necessary.
- Continue to develop additional services **that have been approved and funded**.

Automated daily operations timeline

The ZSDS is a highly automated system.

The **Virtual Pipeline Operator (VPO)** executive orchestrates most functions lights-out:

- checks for incoming data
- executes pipelines according to data type
- archives products in real-time
- distributes alert packets in real-time through Kafka
- generates QA metrics in real-time for ZTF-Depot
- morning cleanup of operations disks
- dispatches nightly summary report via e-mail
- updates epoch-based all-sky coverage maps
- updates reference image library with new references
- updates reference image all-sky coverage maps
- updates Solar System Object ephemeris files
- executes ZMODE: moving object tracklet generation
- prepares for next night: candidate DB tables; setup paths
- database vacuuming

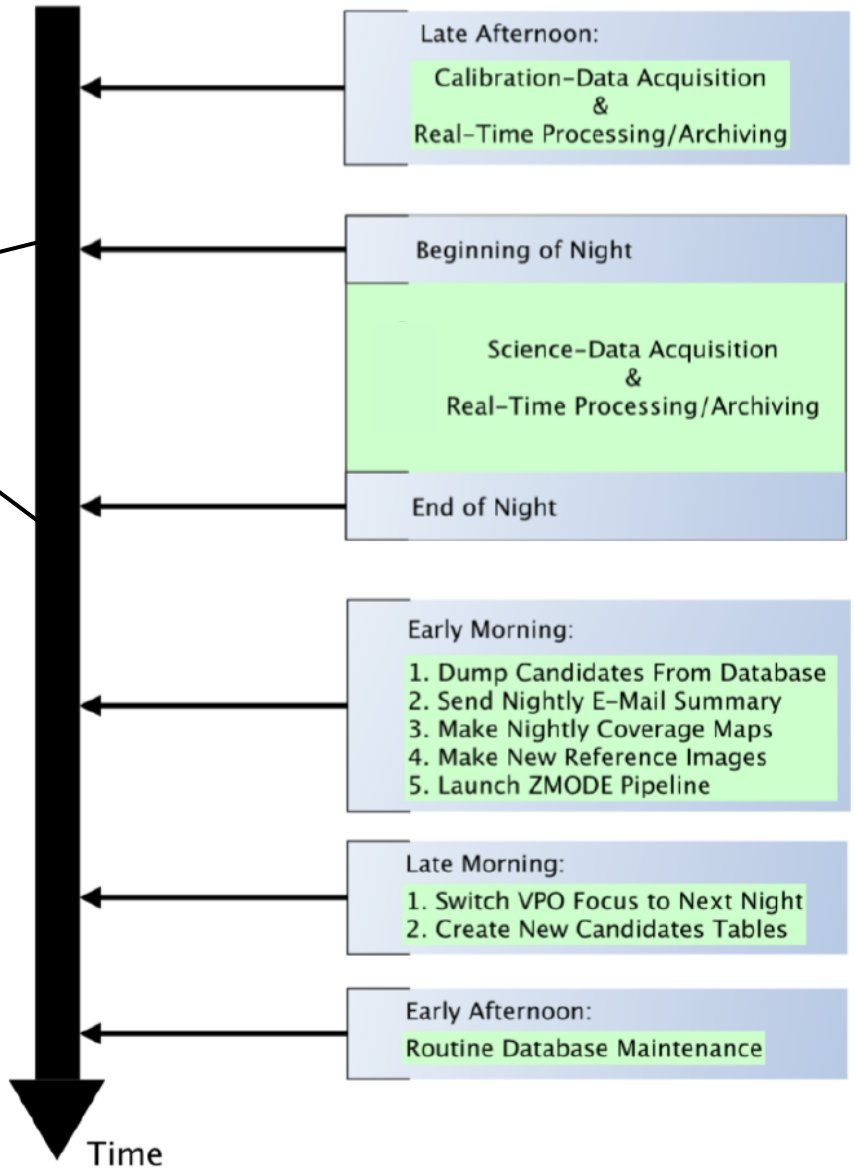

Routine *manual* daily operational tasks

Performed by pipeline operator, trained backup personnel, and Data System Lead.

Prior to forthcoming night (regardless of expectations on weather):

- Check with Observing System and engineering staff on any special activities; non-robotic tests; downtime at P48
- Ensure Virtual Pipeline Operator (**VPO**) daemon is ready to ingest and process data
- Monitor VPO functions, starting with ingestion of calibration data late afternoon
- Monitor reference image generation/updates and other longer-term processing queues **(slide 8)**
- Monitor any concurrent reprocessing tasks / requests; notify project communications group of completion
- Notify project of any scheduled maintenance-related downtime.

End of night (morning) pipeline operator tasks:

- Reprocess and re-archive any data with archive failures in real-time: e.g., due to inadvertent network/DB issues
- Check cluster node performance (throughput), fileservers, DB-servers, network; notify system admin group of issues
- Check disk space on pipeline operations file system, data-ingest staging area, *ZTF-Depot*, and then clean up

0.50 FTE

End of night (morning) activities by Data System Lead:

- Synopsis and analysis of pipeline errors (if any) and cross-check with statistics in nightly summary report
- Quality assure a subset of the previous night's data
- Report any fatal errors to project and QA monitoring team, particularly unexpected glitches from OS

0.20 FTE

Archive (IRSA) related activities

- Daily monitoring:
  - Monitor archive storage consumption; create new disk volumes as old ones fill up.
  - Monitor archive-ingest performance and user-loads.
- Prepare and manage archive systems consisting of
  - File servers and storage.
  - Database servers and storage.
  - Application servers for data transfer, ingestion, query & retrieval.
- Maintenance for:
  - Automatic data ingestion: nightly products in realtime; reference images daily.
  - Data access API's and GUI's for image & catalog files, lightcurve viewing, time series tool.
  - Loading of Object Database Table following matchfile updates.
- Access control database.
- Perform and monitor status of regular backups, on and off site.
- Assist with performance analysis and tuning.
- Additional activities expected in FY19:
  - Lightcurve retrieval API
  - Preparation for public release(s)
  - Addition of ZTF reference images to IRSA Viewer tool (TBD)
  - Monitor performance of services following public release(s)

0.89 FTE

Longer term *manual* operations tasks

- **Lightcurve (source-matching) pipeline:** every month or longer (TBD), contingent on data volume accumulated, trigger source-matching pipeline by matching to new epochal catalog data acquired since last update.
- **Reference image synopses:** typically every week; check data quality, coverage statistics; possible regeneration.
- **Reprocessing requests:** ad-hoc bulk reprocessing/re-archiving of data if anomalies/quality-issues identified.
  - Reprocessing is more likely prior to each Public Data Release if warranted.
- **Reconfiguring pipeline operations system:** e.g., generating new partitioned static calibrator catalogs (for astrometry & photometry) if any new fields not on the predefined science grids are scheduled. This is rare.
- **Financial reporting to project:** once a month.
- **General maintenance/upgrades** to operations and archive DBs and file servers: performance tuning and hardware.
- **Configuration Management (CM) tasks:** software rebuilds, operating system updates/patches, RTBs.

0.40 FTE

- **System Administration tasks** related to IPAC-wide computing infrastructure during scheduled downtime.
- **Troubleshooting system-generated alerts** from hardware/performance issues when they occur.

0.25 FTE

Ongoing pipeline maintenance tasks

- Deploying and validating refinements to automated classifiers for point-source alerts and streaks.
- Importing and testing of improved Star/Galaxy classification scores for PS1 to associate with alert streams.
- Checking / re-deriving bad-pixel masks used in production.
- Refinements to calibration image-products, calibrator catalogs, and pipeline parameters if improvements identified through feedback received from science programs.
  - Goal is to avoid parameter retuning and recalibration as much as possible.
- Pipeline parameter updates or cleanup in response to unforeseen behavior to OS / TCS configuration or scheduling.
- Updates to Explanatory Supplement, in particular Cautionary Notes sections as “features” become known.
- Help-desk support / queries.

0.82 FTE

Summary:
FTE breakdown by task and project phase

Data System Task	≤ Aug 2017 Dev	Sep 2017 → Sep 2018 Commissioning, SV & ZSDS Verification	Oct 2018 → Dec 2020 ZSDS Nominal Ops
Task management and reporting to project; respond to help-desk; budgeting; costing; documentation	0.50	0.30	0.32
Pipeline upgrades, optimization tweaks, tuning	2.70	0.35	0.10
Archive development, user-interfaces, and services	1.50		
Simulation, QA, on-sky performance trending with feedback to pipeline developers	0.15	0.10	
Database administration (archive and pipeline DBs)	0.20	0.25	0.10
Ongoing PTF / iPTF reprocessing	0.20		
Pipeline operational maintenance & executive		0.40**	0.90
Pipeline operator tasks: monitoring and reprocessing		0.50	0.50
Archive ingest and IRSA-related operations: manage archive volumes, tools, services, docs, help-desk, monitor system performance & user loads; data releases		1.00	0.89
System admin: maintenance, monitoring, install/ patching of hardware & system software; backups	0.50	0.50	0.25
TOTAL FTEs:	**5.75 (peak)**	**3.40 (month average)**	**3.06 (plan)**

** was 1.0 FTE prior to Dec 31, 2017

Workforce: ZTF Year 4 (FY18)

September 2017 – August 2018

ZTF Year 4 Workforce Plan versus Actual

8/10/2018 NP

FTE

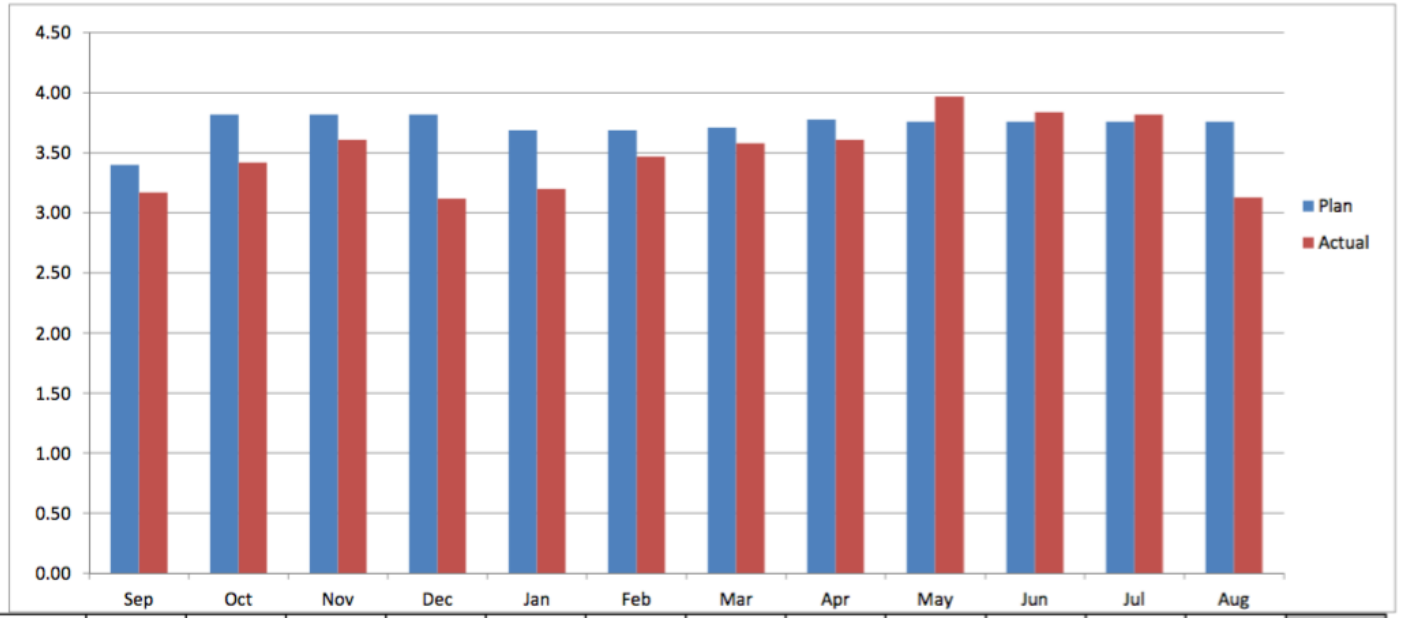

	Sep	Oct	Nov	Dec	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	WM Total
Plan	3.40	3.82	3.82	3.82	3.69	3.69	3.71	3.78	3.76	3.76	3.76	3.76	44.77
Actuals	3.17	3.42	3.61	3.12	3.20	3.47	3.58	3.61	3.97	3.84	3.82	3.13	
Variance	-0.23	-0.40	-0.21	-0.70	-0.49	-0.22	-0.13	-0.17	0.21	0.08	0.06	-0.63	

Plan does not include realization factor (vacation, sick leave...), ~ 15%

Workforce: ZTF Year 5 (FY19)

September 2018 – August 2019

ZTF Year 5 Workforce Plan versus Actual

12/10/18 NP

FTE

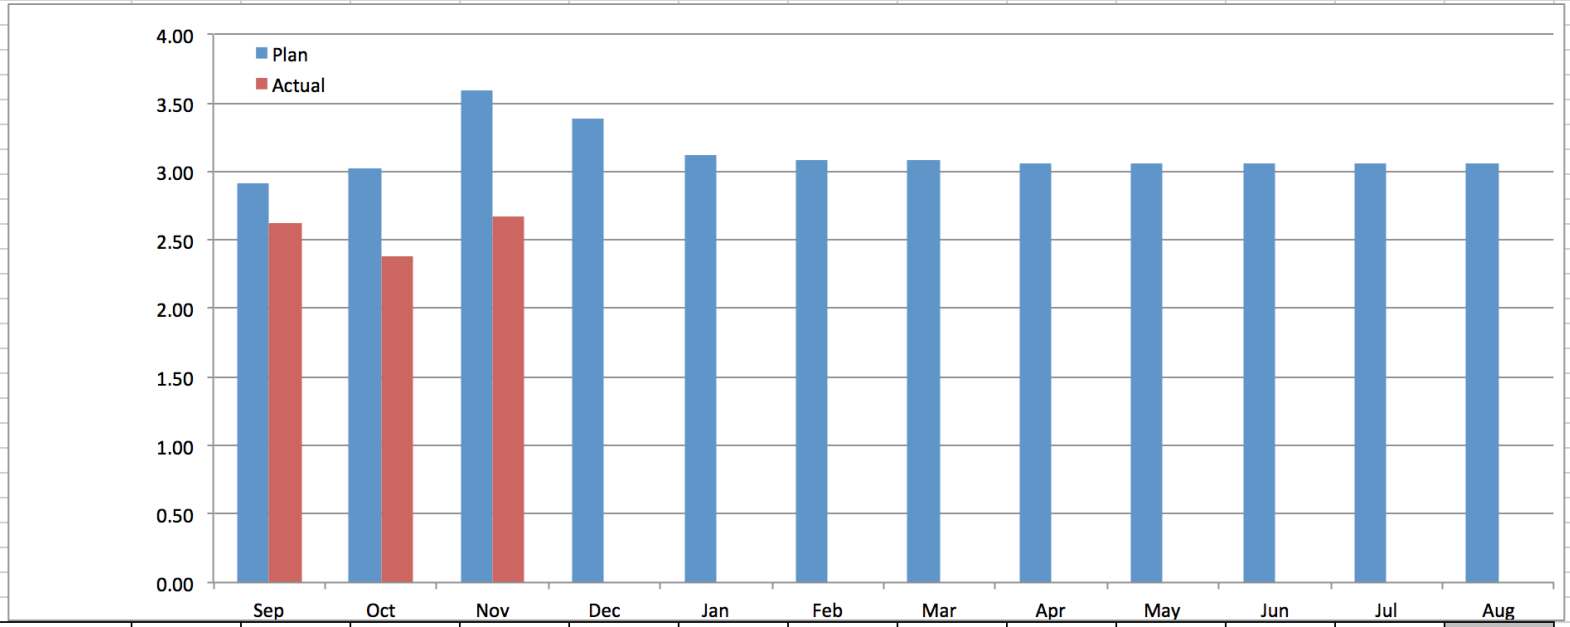

	Sep	Oct	Nov	Dec	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	WM Total
Plan	2.91	3.02	3.59	3.38	3.12	3.08	3.08	3.06	3.06	3.06	3.06	3.06	37.46
Actuals	2.62	2.38	2.67										
Variance	-0.29	-0.64	-0.92										

Hardware procurements until end of 2020

- Compute cluster hardware is in place: 66 nodes (1192 cores); sufficient to support baseline survey.
- Budgeted $337k for FY19 and $36k for FY20.
  - we can defer most of FY19 cost to later since consumption rate has been lower than expected.
- The above includes:
  - RAM and SSD for database servers: ~ $25k.
  - Offsite tape storage: ~ $8k (FY19), ~ $16k (FY20/21).
  - 6 Nexsan E48 w/ 8TB drives: ~$50k each (spread into FY20/21; subject to consumption rate).
  - 600 LTO6 tapes: ~ $15k (spread into FY20/21; subject to consumption rate).

ZTF Years 5, 6, 7 Budgets (FY19, 20, 21)

	ZY5	ZY6	ZY7
Period	**Sep'18 – Aug'19**	**Sep'19 – Aug'20**	**Sep'20 – Dec'20**
FTE	3.12	3.06	3.06
Total labor	$404,785	$436,534	$149,877
Procurement (H/W)	$337,673	$36,102	-
Travel & other	$7,000	$7,000	$5,250
IASC (archive)	$166,241	$176,028	$59,078
Total overhead	$324,413	$412,009	$142,446
Total budget	$1,240,112	$1,067,673	$356,651

Continued Development (R&D) tied to baseline production system

In progress:

- Forced photometry service **already funded by Heising-Simons grant**; delivery: end of February 2019.

- Optimizing DB queries to better handle associations with historical events for alert packets.
  - Significant loads on Operations DB were recently detected; symptom of a growing number of records?

- Lightcurve (*matchfile*) generation updates: filter "bogus" transients from image edges: end of January 2019.

Near future:

- Update from Gaia DR1 to DR2: for both astrometric calibration and alert association: mid January 2019:
  - Will require analysis, regression testing, regenerating static field-based catalogs, reformatting queries & code.

- Update pointing / WCS offset file (*wcs.cfg*) to provide better CCD-quadrant priors: end of February 2019.

- Prepare for First Public Data Release (Anticipated in the Spring of 2019; TBD):
  - Refine content/scope, reprocessing subset of public data, analysis and documentation.
  - Finalize API tools to query lightcurves from IRSA DB (retrieval via GUI already exists).
  - DS staff require a four month lead time prior to actual release date.

Refinements contingent on support from partnership

- These are outside the Data System baseline costing for operations.
- The following are optimizations to primarily support specific partner science programs.
- Correct dome flats for edge / scattering / CCD-etching effects prior to stacking.
  - Includes optimal (re)weighting when combining LED-sets of exposures per filter.
- Star-flat assessment and application (DESY group input).
- Exposure-time correction map (flat augmentation, ~ 0.2% max, at focal-plane edges).
- *i*-filter fringe correction (DESY group input).
- We expect the partnership to provide standalone drop-in software modules if approved by the project.
  - However, DS resources are still needed to integrate and test new software modules, and reassess runtime.
  - These refinements will be scheduled according to priorities.
  - Maintaining the baseline production system and responding to "hiccups" has priority.

Risks

- Catastrophic issues with the Observing System, e.g., Camera or Telescope Control System requiring a significant redesign of the processing / calibration methodology, retuning, or DS architecture.

- Failures in critical hardware components (file or DB servers) requiring immediate replacement.
  - IPAC has spare hardware available, however, ZTF project would need to replenish this.
  - Depending on failure, may take some time to rebuild data content from backups.

- The project may want to set aside contingency funds.

Back up slides

Future ideas (contingent on resources / demand)

- Currently, alerts are distributed as *avro* packets; consumers ingest these into their databases to enable positional association and retrieval of metadata.

- We are currently storing all alert packets in tar-files *per CCD-quadrant* in the archive.
  - The files are only searchable on an image-basis using the standard API and GUI, but no search capabilities exist at the alert (source) level.
  - It would be extremely versatile to search *for individual* alerts, their photometric histories, metadata & cutouts, all of which are archived; this information already resides in a DB at IPAC.

- Make alert packets more compact (subject to project approval and partnership feedback).
  - strip out image cutouts and schema-header; serve cutouts through archive instead.

- Sandbox (work space) environment for users to perform analyses close to where the data resides.

Database Management / Preservation