

ZTF Data System: Operations Plan

Frank Masci & the IPAC-Caltech ZTF Team

ZTF Working Group Meeting, March 2018



Outline

- Overview of **manual** “day-in-the-life” tasks
- Overview of **automated** tasks (Virtual Pipeline Operator)
- Longer term manual operations tasks
- Reference image (re)generation plan
- Staffing plan, activities, and schedule
- Pipeline software and parameters update/approval plan
- Support schedule (hours of operation)
- Contact information and communication channels

Typical “day-in-the-life”: *human pipeline operator tasks*

Manual tasks: pipeline operator or trained backup personnel

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
- Check with OS scheduling lead on any new science fields that are *not* on the predefined grid; if so, need to prepare DS with new static source-catalog calibrator catalogs and other configuration files
- Ensure Virtual Pipeline Operator (**VPO**) daemon is ready to ingest and process data (**see slide 4**)
- Monitor VPO functions, starting with ingestion of calibration data late afternoon
- Trigger/monitor reference image generation and any reprocessing queues (**see slide 6**)
- Monitor any concurrent reprocessing tasks/requests; notify project/communications group of completion
- Notify project in advance of any scheduled downtime to DS operations

End of night (morning) tasks -- signaled by reception of manifest summaries from P48:

- Synopsis of pipeline errors (if any) and cross-check with night summary report and statistics
- Report any fatal errors to project and QA monitoring team (Richard W), particularly issues from Observing System
- Review cluster node performance (throughput), file servers and DB-servers; notify system admin group of issues
- Check disk space on operations sandbox, ingest staging area, *ztf-depot*, and then clean up
- Database “vacuuming” and maintenance (in coordination with DBA; **see slide 5**)
- Continued monitoring of daily **VPO** tasks, including ref-image generation, source-matching and other reprocessing

Typical “day-in-the-life”: *Virtual Pipeline Operator (VPO)*

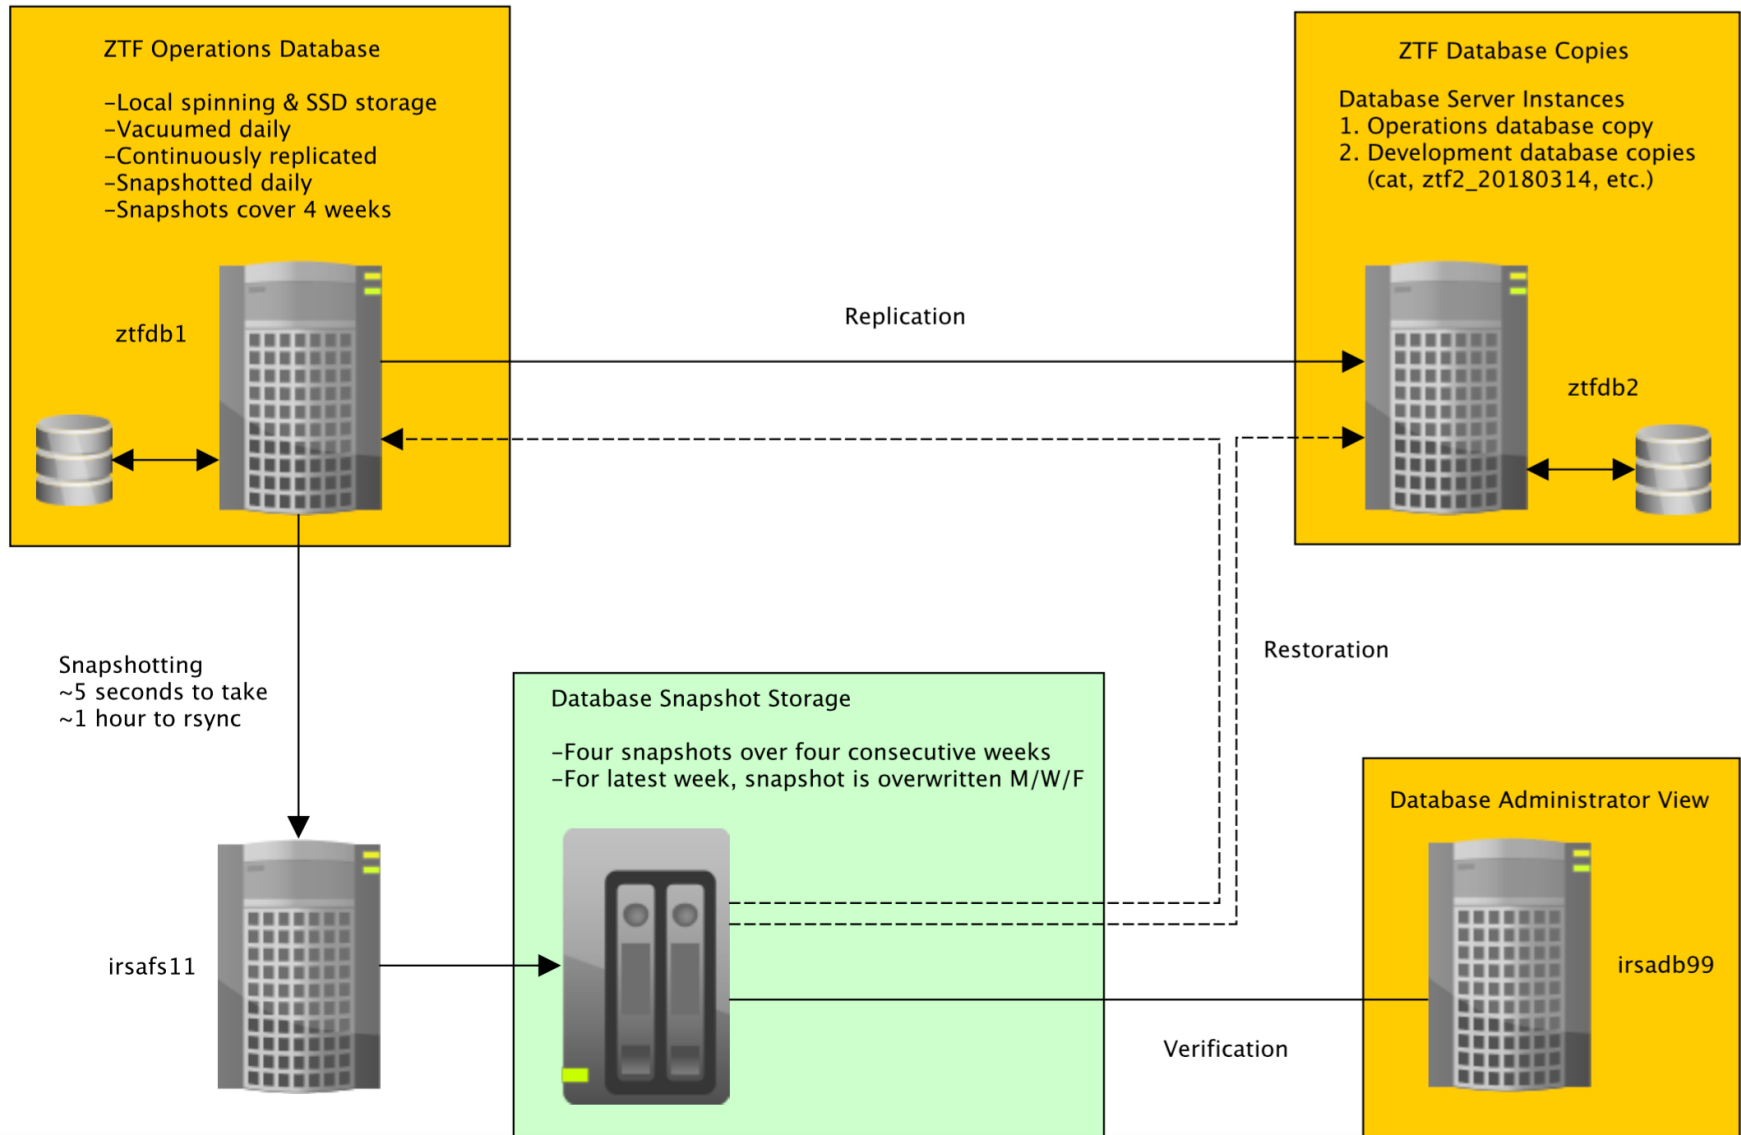
VPO runs lights out, does the heavy lifting; automated functions:

- Checks for incoming data (24/7), executes raw-data ingest and relevant science pipelines according to image types
- Pipeline orchestration and prioritization rules in accordance with pipeline executive (SLURM)
- Creates date-dependent paths in operations and ztf-depot file systems
- Archives processed products and performs verification checks using “checksums”
- Tracks statistics on number of products and number of sources extracted throughout night
- Posts alert packets to UW using Kafka mirroring
- Looks for end-of-night signal, performs accountability check on raw data files received versus sent from mountain
- Dispatch nightly-summary email report with statistics
- Triggers moving-object (ZMODE) pipeline following completion of nightly real-time processing (tracklet pipeline)
- Generate depth-of-coverage maps: both incremental (for recent night) and cumulative
- Copy recent night transient-candidate metadata from DB to filesystem for fast history association in next night
- Monitor disk space in raw-data inbox (xfer server) and local scratch on nodes, then clean up

Automated on archive (IRSA) side:

- Create new disk partitions and volumes as old ones fill up
- Periodic backups to tape and offsite

Database Management / Preservation



Longer term (manual) operations tasks

- **Weekly or longer:** update Solar-System ephemeris files in operations with latest from Minor Planet Center
- **Lightcurve (source-matching) pipeline:** every month or longer (TBD), contingent on data volume accumulated, trigger source-matching (lightcurve) pipeline by matching to new catalog data acquired since last update
- **Reference image synopsis and (possible) regeneration:** contingent on data quality; see **slide 7** for plan.
- **Reprocessing requests (both bulk and ad-hoc):** desirably, want to keep this at a minimum following stabilization of all pipeline modules and parameters; want to avoid heterogeneity in archived products.

Other periodic tasks:

- Maintenance / upgrades to operations and archive databases
- General hardware maintenance/upgrades, including archive tools and services
- Configuration Management (CM) tasks: software rebuilds, operating system updates/patches, RTBs
- Anticipated downtimes will be broadcast in advance to project through communications group
- Throughput monitoring of entire system: cluster nodes, file-servers, DB-servers, network, user access loads
- Tending to *xymon* alerts from all servers

Reference Image (re)generation plan

- Reference images are used in the image-differencing and source-matching pipelines downstream.
- Filters to select “good” quality science images as determined from analyses are now defined (see PL document):
 - FWHM; astrometric & photometric calibration quality (ZP and color terms); noise and background levels; magnitude limit; only data with DATEOBS \geq 2018-02-05 (following reconfigured camera);
 - *minimum* number of images satisfying above criteria = **6**: to make references quickly for the mini-surveys
 - *maximum* number of input images following a sort into ascending FWHM; currently = **15**.
- Currently, following a night, an automated process checks which fields, quadrants, filters observed satisfy the above criteria. If satisfied and if **do not** already have an archived reference image, reference image pipeline is triggered.
 - Each reference image has quality metrics that are internally thresholded to set a usability *status* flag in the database; failures will be placed in a regeneration queue (see below)
 - This process should (rapidly) converge to a complete reference image library following repeated coverage of the science grid fields
- We have the ability to regenerate all (or a subset) of references at any time following a global synopsis of all science grid fields and filters. Beware of downstream dependencies(!); *RealBogus* may need to be retrained.
- If a reference image is determined to be bad after use (i.e., persistently gives a bad quality difference image), its usability flag will be unset in DB and it will be placed in the regeneration queue
 - may warrant a revision of the input selection criteria so that all fields/quadrants/filters can attain a reference.
- Any updates to the input selection criteria for reference image generation will be subject to CCB purview (slide 9).

Staffing Plan, Tasks and FTE breakdown

- We are currently in the Data System Verification phase; this will continue until **July 31, 2018**
- This phase will include possible low-level pipeline development, bug fixes, and tuning in response to science analyses, pending CCB approval (see slide 9)

Sep 2017 → Jul 2018

Data System Task	Dev	Commissioning, SV & DS Verification	Nominal Ops
Task management and reporting to project; respond to help-desk; budgeting; costing; documentation;	0.50	0.40	0.30
Pipeline upgrades, optimization tweaks, tuning	2.70	0.35	
Archive development, user-interfaces, and services	1.50		
Simulation, QA, on-sky performance trending with feedback to pipeline developers	0.15	0.15	
Database administration (archive and pipeline DBs)	0.20	0.30	0.10
Ongoing PTF / iPTF reprocessing	0.20		
Pipeline maintenance and operations: specifically pipeline operator tasks, reprocessing, monitoring		1.0**	1.50
Archive ingest and IRSA-related operations: manage archive volumes, tools, services, docs, help-desk		1.00	1.00
System admin: maintenance, monitoring, install/patching of hardware & system software; backups	0.50	0.50	0.25
TOTAL	5.75	3.70	3.15

** was 1.5 FTE prior to Dec 31, 2017

Updates to production pipeline software or parameters

- All updates to pipeline software now go through a Change Control Board (CCB) before deployment to production
 - Proposed update is communicated to Project Scientist and identified stakeholders where relevant
 - Potential impact and risk is assessed through discussions with developer (and scientists if applicable)
 - Given approval, update is implemented and tested in a production-like environment
 - Depending on update, outputs from testing are validated and communicated to stakeholders
 - Software and/or parameter updates are deployed to production
- A history of all updates is currently documented in GitHub for everyone to see
- Includes:
 - pipeline parameter updates or any system-related hardware / operating system upgrades
 - improvements to ZTF-specific archive user-tools and services (on IRSA side)
- Excludes:
 - development of new *un-scoped* pipeline functionality and/or archive-centric tools throughout survey ops
 - will require a separate cost analysis before proceeding
- Issues are currently tracked in JIRA. Periodic snapshots can be made and forwarded to project management

Support and Troubleshooting

- The current cost estimates (slide 8) assume the *PTF* operations model
- Follow regular Business Hours model of operation (Pacific Time)
- No after-hours on-site support, such as troubleshooting during night-time operations (e.g., hardware failures)
- Communication will be available until late evening using the *ZTF general* Slack channel to report status / glitches
- We will diagnose and fix any issues at sunrise (Pacific Time):
 - replace hardware with backup hardware on standby
 - reprocess previous night or portions thereof if needed, with re-distribution of alerts
 - schedule downtime if needed

Contacts and communication channels

- We request that all project members, including partners, continue to send questions and queries to the ZTF communications group: ztf.communication.coordinators@gmail.com
 - This account is moderated by members of the ZTF project, selected by the Project Scientist (not the DS team)
 - Depending on the query, we advise the moderator to then forward this to the specific individual(s) below
- **No data in archive, or depot, or no alerts generated, or for any ad-hoc reprocessing request:**
 - **Primary:** Frank Masci: fmasci@ipac.caltech.edu (626) 395 1962
 - **Backup 1:** Russ Laher: laher@ipac.caltech.edu (626) 395 2596
 - **Backup 2:** Ben Rusholme: rusholme@caltech.edu (626) 395 1883
- **Pipeline operations status, updates on (re)processing queues, reference image (re)generation status:**
 - **Primary pipeline operator:** Ron Beck, beck@ipac.caltech.edu (626) 395 1940
 - **Backup pipeline operator 1:** Ben Rusholme: rusholme@caltech.edu (626) 395 1883
 - **Backup pipeline operator 2:** Russ Laher: laher@ipac.caltech.edu (626) 395 2596
- **Archive data-access issues:**
 - **User accounts and authentication:** Stefanie Wachter: wachter@ipac.caltech.edu
 - **Problem with a GUI or API service:** email IRSA help-desk: irsasupport@ipac.caltech.edu
 - **Backup 1:** Steve Groom: sgroom@ipac.caltech.edu (626) 395 1878
 - **Backup 2:** Serge Monkewitz: smm@ipac.caltech.edu (626) 395 1863
- **Data-quality centric, formats, content, or algorithmic related:**
 - Frank Masci: fmasci@ipac.caltech.edu (626) 395 1962
- We will continue to monitor and post updates to the ZTF *general* Slack channel throughout operations