Minor Planet Center

Newsletter - February 2025

2025 FEBRUARY 28

In this month's issue:

The long road to Vera Rubin Observatory/LSST and NEO Surveyor

State of the Art

The Vera Rubin Observatory/LSST is expected to see first light in August 2025, while the NEO Surveyor mission is preparing for a launch currently scheduled for September 2027. These milestones mark significant moments for the MPC and the broader Solar System and Planetary Defense communities.

Despite their differences in design and objectives, both telescopes are poised to transform our understanding of the Solar System's small body population. The MPC has spent years preparing to ingest and process the vast influx of data these missions will generate, all while maintaining current operations. This has been a challenging task.

The MPC is a complex system with many moving parts, serving a diverse user base with varying levels of expertise in software development and data analysis. From users needing access to individual object files to those replicating our entire database and keeping their systems in sync, we strive to support them all.

We also recognize that entire operational systems, from telescope observations to orbit-computation pipelines, rely on our existing products and that any change on our side requires careful planning, time, and dedicated personnel on the community side. That's why we've focused on improving our current systems while gradually transitioning to new processes, ensuring continuity for our users and the broader scientific community.

With this newsletter, we aim to provide an overview of the work accomplished in recent years, the ongoing efforts behind the scenes, and our plans for the coming months.

The work behind the scenes

Current status of the PostgreSQL tables

The majority of MPC data is stored in PostgreSQL tables, which are replicated to our community via the Small Body Node. Full details on this process can be found in our May 2023 and April 2024 newsletters. Additionally, our Documentation page includes a dedicated section on the Replicated tables, with an introduction on their usage, detailed schema descriptions (including known issues), and examples of useful queries.

Current replicated tables

NEOCP Table

A full set of tables containing observations and orbits for objects currently on the NEOCP. These tables are highly stable, with no planned modifications, and are internally used for all NEOCP-related processing.

Identification tables

- ***Current identification table***: contains all primary and secondary packed and unpacked provisional designations for every object ever designated by the MPC.
- ***Numbered identification table:*** includes all numbered objects.

Both tables contain minor planets, comets and natural satellite designations. They are stable and used internally across all MPC processes.

Primary object table

A comprehensive record of all primary designations for designated minor planets, comets and natural satellites.

Observations table

This table includes all the observations published by the MPC or currently in the Isolated Tracklet File (ITF). It is largely stable, but some considerations apply based on the small-body population:

- *Minor planets*: while minor inconsistencies affect fewer than 0.01% of entries, they are a lower-priority issue. All of our processing pipelines now rely solely on this table.

- *Natural satellites*: the table is clean and consistent with the flat files. We will soon transition natural satellite pipelines to use it as the primary data source.
- *Comets*: we are in the process of reconciling the flat files with this table. This is particularly complex for comet fragment observations, so for now, comet-processing pipelines continue to use the flat files.

Orbits table (Work in Progress)

The orbits table, currently under development, was released early to allow the community to familiarize themselves with the data and provide feedback. At present, it only includes minor planet orbits: comets and natural satellite orbits will be added in the future. Work is ongoing to align the table with existing flat files, and we anticipate a stable version in the coming months. Once finalized, it will become the primary internal data source for minor planet orbits. Further updates will be shared in future newsletters.

Observation Alterations Tables (Work in Progress).

Designed to record corrections, remeasurements, and deletions, these tables are still under active development as we refine the supporting code.

Upcoming Discovery Information Table.

Discovery information, both provisional and permanent (for numbered objects), has not yet been incorporated into any PostgreSQL tables. We are actively working on a new table to address this and expect it to be finalized in the next month.

One of the MPC's top priorities is ensuring that all data required for processing and publication is accurately stored in PostgreSQL tables and that these tables serve as the primary source of information for all processing pipelines.

Virtualization and migration of the astrometry processing system database to virtual machines

In our January 2025 newsletter we explained how over the past year, the MPC has been working on virtualizing its systems in preparation for the significant data influx expected from Vera C. Rubin Observatory (VRO/LSST) and NEO Surveyor.

Virtualization allows multiple virtual machines to run on a single physical server, creating a more flexible, scalable, and reliable computing environment. This improves data processing, system performance, and maintenance efficiency, while reducing downtime.

On February 12, 2025 we migrated the astrometry processing system database to a new virtual machine. The migrated database includes all the operational tables as well as all the tables that are replicated via SBN.

Aside from the three-hour maintenance windows and the catch-up time required for the replica after migration, our end users should not have experienced any significant downtime.

We are currently monitoring the performance of the migrated database.

Our next steps include:

Migrating the system that ingests the observations to Amazon Web Services.
The new ingesting system has been operational for some time, running in parallel with our current system to allow thorough testing and performance improvements. This system is also being used for observation ingestion as part of testing plans with the VRO/LSST and NEO Surveyor teams.

Given the significance of this transition, we will conduct multiple test runs to ensure a seamless process. As always, we will send out notifications in advance, but aside from the required maintenance window during the switch, we do not anticipate any major issues or disruptions for our end users.

The transition to the new AWS ingestion system will enable us to address several outstanding challenges, including the ability to support large-scale remeasurement submissions from users.

Migrating our major processing pipelines to the new state of the art.
We have been developing a new orbit-fitting processing pipeline built using modern software development practices. This pipeline leverages RabbitMQ, a messaging and streaming broker, to efficiently queue each ingested tracklet into the appropriate processing path. Designed as a Python package, the new pipelines can run via Docker, allowing flexibility in deployment. Thanks to its scalability and portability, the pipeline is optimized for our new **cloud-based infrastructure**, where processing cores can be dynamically allocated as needed.

The new pipeline's output will be **identical** to current products, ensuring a seamless transition for users with no visible differences after the switch.

This upgrade marks a significant improvement for the MPC. While stress tests have confirmed that our system could withstand VRO/LSST going live today, they have also highlighted potential bottlenecks,

particularly in orbit fitting, due to the large data influx. This new pipeline addresses those challenges, ensuring a more efficient and resilient system.

Internal migration towards a database-centric system - MPC publications

We are transitioning all product generation to retrieve data directly from our PostgreSQL database tables.

The MPC's official publications comprise five main products.

1. Minor Planet Electronic Circulars (MPECs).
Ad-hoc publications used to share:

- New NEO designations (e.g. https://minorplanetcenter.net/mpec/K25/K25DO4.html),
- New comet designations (e.g. https://minorplanetcenter.net/mpec/K25/K25DN5.html),
- New TNO designations (e.g. https://minorplanetcenter.net/mpec/K25/K25D19.html),
- New natural satellite designations (e.g. https://minorplanetcenter.net/mpec/K24/K24DB4.html),
- Comet recoveries and identifications (e.g. https://minorplanetcenter.net/mpec/K25/K25DN8.html),
- Identifications of NEOs (https://minorplanetcenter.net/mpec/K25/K25DN8.html),
- Identification of TNOs, when the two single oppositions TNO form a multi-opposition orbit (https://minorplanetcenter.net/mpec/K25/K25DK9.html),
- Recovery of NEOs, TNOs (e.g. https://minorplanetcenter.net/mpec/K25/K25DM8.html, https://minorplanetcenter.net/mpec/K25/K25DF0.html).

Given the current status of our PostgreSQL tables, MPECs for *new NEOs* are now entirely generated using data from our PostgreSQL database. *Comets*, however, still primarily rely on flat files. *Identifications and recoveries* use a combination of database tables (e.g. observations) and flat files (e.g. orbits).

Designations of NEOs are partially automated. For a description of the NEOCP Automated Processing Pipeline, please refer to the January 2024 newsletter.

Recovery and identifications of NEOs are fully automated, see our documentation pages on identifications and our recovery page.

We are actively working on fully automating all the other processes.

2. Daily Orbit Update (DOU)
Daily publication to provide revised orbital elements for minor planets, incorporating the latest

observations for NEOs. The observations are currently generated from our observations PostgreSQL table, while orbits are still sourced from flat files.

The DOU is fully automated.

3. Mid-month Minor Planet Supplements (MPS)

Circulars, which are used to publish minor-planet observations sorted by designation and date between monthly circulars, are now generated entirely from our observations PostgreSQL table. The mid-month batches are released approximately weekly, depending on the volume of received observations, and include a complete list of observations from the previous week.

The mid-month circulars are fully automated.

4. Monthly Minor Planet Circulars (MPCs)

These Circulars are official Minor Planet Center (MPC) publications that provide essential updates on minor planets, comets, and natural satellites. Each edition offers a comprehensive summary of MPC operations during the period between two circulars. They include new designations, updated orbital elements, recent observations, newly assigned and deleted identifications, newly numbered objects, and their discovery information.

The process of generating the monthly circulars is one of the most complex systems currently in place at the MPC. Modernizing both the process and its underlying code is now a high priority to improve efficiency and ensure long-term sustainability. This includes the derivation of all the main products from our database table. The observations are currently generated from our table, while everything else is still sourced from flat files.

The Monthly Circulars are not yet fully automated, and improving their automation remains one of our key priorities. Additionally, we are exploring alternative methods for officially publishing the same information in more efficient formats. Recreating files containing 500 million observations is already a time-consuming process, and with data volumes expected to grow significantly in the coming months and years, the current approach will soon become unsustainable. Our goal is to develop a more scalable and manageable solution for the long term.

This is a work in progress.

5. Observations and orbits of comets

A weekly product is released to provide new observations and updated orbits of comets. Significant effort

has been dedicated to automating comet processing; however, given the current state of our database tables, this process still relies entirely on flat files for data sourcing.

This is a work in progress.

Additional lower priority projects

We have already outlined several high-priority projects currently in progress. Many of these initiatives link to community requests submitted through multiple Jira tickets. In addition to these major projects, we continuously manage daily operations and address smaller issues or requests as they arise.

We also maintain a list of medium- and lower-priority tasks, many of which are dependent on the completion of higher-priority projects before they can move forward.

Ongoing projects include, but are not limited to:

A new Ephemeris Service and MPChecker
We are actively upgrading the code for a new version of the ephemeris service and MPChecker, both of which depend heavily on the structure of the orbit table. Recognizing the importance of these tools to the community, we are working on a prototype for an early release, allowing users to familiarize themselves with the new system in advance.

A new monitoring system
Based on user feedback, we have been working to improve communication through the MPC-ml mailing list and the Status page. We recognize that our updates are not always as timely or effective as we would like, and we are taking steps to address this.

To enhance monitoring and transparency, we are implementing Grafana and Prometheus, open-source systems that provide a dimensional data model, flexible query language, efficient time-series database, and modern alerting capabilities. While these dashboards are currently for internal use, our goal is to develop metrics that can be displayed on the MPC website, giving users insight into system performance in real time.

Heliolinc @ the MPC
Heliolinc is a heliocentric linking algorithm developed by a team including Matthew Holman, the former MPC Director, and Matthew Payne, the current MPC Director. It is designed to efficiently identify and link detections of moving objects across multiple nights of observations.

In preparation for NEO Surveyor, we have been developing a new pipeline that fully integrates the Heliolinc algorithm to improve our current linking system. The pipeline is currently undergoing testing.

Our new available services

As part of our transition to a database-centric system, we have also been refining how we share data with our users, incorporating feedback from the community.

Documentation

The first step in this effort was reorganizing and improving our documentation using Django.

New APIs

In parallel, we launched the data.minorplanetcenter.net subdomain, built on a Docker infrastructure, to serve as a dedicated server for hosting web content (see our April 2023 newsletter). This subdomain has been instrumental in the development of the new APIs, which are described in several newsletters: October 2023 (WAMO), February 2024 (designation_identifier), March 2024 (Observations), July 2024 (NEOCP observations), November 2024 (Orbits).

MPC Explorer

The MPC Explorer is designed as the future replacement for the db_search functionality. Our goal is for it to serve as a comprehensive tool, providing detailed information on observations, orbits, statistics, discoveries, and ephemerides for every object in our database.

The MPC Explorer is fully built on the new APIs and PostgreSQL tables, making it more stable and reliable than the current db_search, for all the available data.

We are actively working on expanding its capabilities by including discovery information and generating lists of notable objects (e.g. past impactors, dual status objects, etc.).

Additionally, we plan to expand the NEOCP section to include the possibility of displaying variant orbits, providing our users with a more complete and detailed view of these objects.

It is very important to highlight that all the data available on the MPC Explorer can be accessed at the same time using the available APIs.

The best way for our community to help us prepare for VRO/LSST and NEO Surveyor is by integrating our new systems into their daily operations, transitioning away from old data products and adopting the newly generated ones. We will provide you at least 60 days' notice before officially deprecating legacy services, ensuring ample time for adjustment. However, it is important to recognize that VRO/LSST and NEO Surveyor mark a significant shift in how the Solar System community interacts with and utilizes data. Your early adoption and feedback will be invaluable in making this transition as smooth as possible.

WAMO Tool Updates

As mentioned in the January 2025 newsletter, the WAMO API has been updated to use a new format. We remind everyone that **requests using the old format will be deprecated tomorrow, March 1, 2025**.

In addition to using your *submission_block_id* (e.g. 2024-05-02T21:03:35.001_0000FzZw_01) to request information about your observations batch, you can now also use your ***submission_id*** (e.g. 2024-05-02T21:03:35.001_0000FzZw).