

PS1 FAQ - Frequently asked questions

This page includes answers to some of the more common questions about PS1 data. If you have suggestions for additional topics you would like to see addressed, please send email to archive@stsci.edu.

- The Pan-STARRS project
 - What is Pan-STARRS?
 - What types of data were obtained?
 - How does PS1 compare to other surveys such as the Sloan Digital Sky Survey (SDSS)?
 - Which data is available in which data release?
 - DR1
 - DR2
- PS1 data access
 - What should I know about the current data release?
 - How do I get started? Are there examples of data queries?
 - How do I get PS1 data for a list of sources I have?
 - What tools are available to extract and examine PS1 data?
 - How do I access PS1 image data?
- PS1 catalog data
 - How do I make sense of the entries in the PS1 data tables?
 - How can I distinguish stars from galaxies in the catalog?
 - Some stack objects have multiple entries. How can I select the correct entry?
 - How can I retrieve catalog data with python through the MAST API?
 - How can I search for moving solar system targets in the PS1 images and catalogs?
- PS1 astrometry
 - How good is PS1 astrometry?
- PS1 photometry
 - What filters did PS1 use?
 - How good is PS1 photometry?
 - Which magnitudes should I use?
 - What are the brightest and faintest stars for which the data are reliable?
 - How reliable is PS1 photometry of nearby, bright galaxies?
 - Why is my photometry from images so different from the catalog photometry?
- PS1 Images
 - How can I convert pixel values on the stack images back to the original counts?
 - Why do some catalog objects apparently not have images available?
 - How do I convert pixel values into magnitudes?
 - Why are the coordinates incorrect when I display a skycell image in DS9?
- PS1 Timing
 - What is the time scale used for the PS1 observation times? Why do the observation times in the image header and catalog disagree?

The Pan-STARRS project

What is Pan-STARRS?



The Panoramic Survey Telescope and Rapid Response System ([Pan-STARRS](#)) is an wide-field astronomical imaging and data processing facility developed at the University of Hawaii's Institute for Astronomy. Pan-STARRS1 (PS1) is the first telescope of Pan-STARRS to be completed and is the basis for DR1 and DR2. The PS1 survey used a 1.8 meter telescope and its 1.4 Gigapixel camera (GPC1; see [PS1 GPC1 camera](#)) to image the sky in five broadband [filters](#) (g, r, i, z, y). The PS1 Science Consortium funded the operation of the Pan-STARRS1 telescope, situated at [Haleakala Observatories](#) near the summit of Haleakala in Hawaii, from 2010 to 2014, and the data released in DR1 and DR2 are based on these observations.


What types of data were obtained?

Pan-STARRS1 has carried out a set of distinct synoptic imaging sky surveys including the 3pi Steradian Survey and the Medium Deep Survey in 5 bands (grizy). On average, about 12 individual epochs (called warps) were taken in each filter at a given position, which are then stacked to obtain deep static images of the sky. The mean 5-sigma point source limiting sensitivities in the individual grizy epochs are (22.0, 21.8, 21.5, 20.9, 19.7) respectively. The mean 5-sigma point source limiting sensitivities in the stacked 3pi Steradian Survey in grizy are (23.3, 23.2, 23.1, 22.3, 21.4) respectively. The upper bound on the systematic uncertainty in the photometric calibration across the sky is 7-12 millimag depending on the bandpass.

How does PS1 compare to other surveys such as the Sloan Digital Sky Survey (SDSS)?

Metcalf et al. MNRAS 435, 1825 (2013) - has a detailed comparison of the PS1 Small Area Survey 2 data (a test 3pi region taken under optimal conditions) with SDSS. Here are plots comparing 5-sigma counts from SDSS DR8, PS1-SAS2 and Stripe 82. This is a best-case scenario. Generally speaking, averaged over the whole sky, PS1 is similar in depth to SDSS in g and r, ~0.5 deeper in i and ~1 mag deeper in z (for point sources).





Which data is available in which data release?

DR1

- Stacked images
- Mean and stack photometry

DR2

- Warp images (=individual epochs)
- Photometry of warped images

PS1 data access

What should I know about the current data release?

Users of the current PS1 DR2 data should be aware of a few issues and inconsistencies in the data.

- Approximately 4.9 square degrees (0.016% of the survey area) has sources that are missing due to various issues in the database population procedures. These are expected to be filled in eventually.
- The stack positions (raStack, decStack) have considerably larger systematic astrometric errors than the mean epoch positions (raMean, decMean).
- The astrometric and photometric keywords in the image headers are slightly inconsistent with the astrometry and photometry for catalog objects. The catalog quantities are more accurate due to post-image-processing improvements.
- Some images near the celestial pole have registration issues that lead to poor photometry and astrometry.
- There are some astrometry issues at the $\sim 0.5''$ level near the pole (mostly above 85 deg declination) which have caused problems with the stacking. As a result, both warp and stack images may have incorrect WCS, and in extreme cases double objects may be visible in the stacks. There will also be errors in the stack positions (but not mean epoch positions) in the catalogues.

- FITS image headers use the obsolete PC001001, PC001002, etc., keywords for the WCS. The fluxes have been non-linearly scaled using an asinh transformation. The files also use the tile-compressed image format. The full skycell FITS images do not have a RADESYS keyword, which leads some software (including DS9 v8) to incorrectly interpret the coordinates as being equinox 1950 rather than equinox 2000. The fix is to insert the keyword RADESYS = 'FK5' in the header. The headers also lack the TIMESYS = 'TAI' keyword, which is important if you care about the image observation dates at an accuracy of 30 seconds.

See the [PS1 DR2 caveats](#) page for more information.

Users of the current PS1 DR1 data should be aware of a few issues and inconsistencies in the data.

- The stack positions (raStack, decStack) have considerably larger systematic astrometric errors than the mean epoch positions (raMean, decMean).
- The astrometric and photometric keywords in the image headers are slightly inconsistent with the astrometry and photometry for catalog objects. The catalog quantities are more accurate due to post-image-processing improvements.
- There are some astrometry issues at the ~0.5" level near the pole (mostly above 85 deg declination) which have caused problems with the stacking. As a result, both warp and stack images may have incorrect WCS, and in extreme cases double objects may be visible in the stacks. There will also be errors in the stack positions (but not mean epoch positions) in the catalogues.
- FITS image headers use the obsolete PC001001, PC001002, etc., keywords for the WCS. The fluxes have been non-linearly scaled using an asinh transformation. The files also use the tile-compressed image format.
- The non-PSF magnitudes from the stacks can have zero-point errors up to 0.3 mags. This is most prevalent in the i-band.
- The stack and forced warp magnitudes have an incorrect flatfield pattern imprinted upon them at the level of a few hundredths of a magnitude (and on the scale of an exposure, so about 3 degrees).

See the [PS1 DR1 caveats](#) page for more information.

How do I get started? Are there examples of data queries?

See the [PS1 Sample queries](#) page.

How do I get PS1 data for a list of sources I have?

The [How to retrieve and use PS1 data](#) page describes the three current interfaces to access the PS1 catalogs. All of them can be used to cross match a list of sources.

1. Follow the **File Upload Form** link in the upper right corner of the [simple web form](#) to cross-match a list of positions with the catalog.
2. Use the **VO Cone Search** interface in a script to query many positions. This is easily accomplished, e.g., in Python. Note that it is possible to get results back from this interface in a variety of formats, not just in VOTable format. See the [documentation](#) for more information.
3. Large cross-matches can be carried out in [Casjobs](#). These queries can be customized to extract only the columns of interest and to apply constraints based on additional catalog information (e.g., to require that the object be detected in a particular filter). One of the SQL query examples in Casjobs shows how to do a cross-match. Login to Casjobs, go to the **Query** page, select the **PanSTARRS_DR1 context**, and click the **Samples** button (just below the context menu). The **PS1 crossmatch** query does a cross-match with one of your MyDB tables:

```
/* Create a table of positions within a file in your local directory that looks like this
```

```
ID, RA, DEC
1, 1.58851, -0.07845
2, 5.53825, -1.51730
3, 6.33772, -12.34491
```

```
Upload this file into MyDB as a table named MySources using the Import tool in CasJobs.
Run the following query to obtain the PS1 catalog information
about matches that lie within 20 arcseconds of the search positions.
Reports input values, PS1 positions and PSF magnitudes. */
```

```
select s.ID, s.ra, s.dec, o.objID, o.raMean, o.decMean,
       o.nDetections, o.ng, o.nr, o.ni, o.nz, o.ny,
       m.gMeanPSFMag, m.rMeanPSFMag, m.iMeanPSFMag, m.zMeanPSFMag, m.yMeanPSFMag
from MyDB.MySources s
cross apply fGetNearbyObjEq(s.ra,s.dec,20.0/60.0) nb
inner join ObjectThin o on o.objid=nb.objid and o.nDetections>1
inner join MeanObject m on o.objid=m.objid and o.uniquePspOBid=m.uniquePspOBid
```

What tools are available to extract and examine PS1 data?

How do I access PS1 image data?

See the documentation on the [PS1 Image Cutout Service](#) for information on how to access the PS1 image data. There is a simple web interface that provides access to JPEG previews, FITS cutout images, and full-frame FITS skycell images. The web interface also creates color images by combining the 5 PS1 filters. The interface provides services that can be used for scripted downloads of large numbers of cutouts.

There is also a simple [Jupyter notebook script](#) that shows how to access the `ps1filenames.py` and `fitscut.cgi` scripts from Python.

PS1 catalog data

How do I make sense of the entries in the PS1 data tables?

The PS1 databases are complex, with information stored in many different tables. The [PS1 Sample queries](#) page gives some examples that can be helpful in getting an overview of the use of the tables. The [PS1 Database object and detection tables](#) page has detailed descriptions of the tables that have data on individual objects. You can search within this page for particular columns or quantities. There are other pages that describe tables with [observational metadata](#) (e.g., exposure dates and times) and [catalog metadata](#) (e.g., information on flags).

How can I distinguish stars from galaxies in the catalog?

There are several possible approaches to separate extended objects from point sources. See the page describing [how to separate stars and galaxies](#) for suggested approaches. Note that a recently added MAST High Level Science Product, [PS1-PSC](#), provides point-source probabilities for 1.5 billion objects in the PS1 DR2 catalog. It is accessible via the [MAST Casjobs](#) interface and may be the current best choice method to separate stars from non-stars.

Some stack objects have multiple entries. How can I select the correct entry?

For various reasons, there can be several entries for a given stack object. Currently, the best way to select a unique entry is to use the PrimaryDetection column in [StackObjectThin](#), a geometric flag which assigns a unique area on the sky to each skycell. However, in about 0.5% of the stack objects, there are multiple entries with PrimaryDetection=1 for a given object. This happens when there is a very close split stack detection on the primary skycell. The PrimaryDetection flag is then set for both entries because both detections are in the primary skycell (for that location on the sky) and both are assigned to the same object. In severe cases, there may be more than 2 entries with PrimaryDetection=1. It is also possible for an object to exist without a PrimaryDetection=1 detection if it lies in an area overlapping an adjacent skycell but is not detected on the primary skycell, only on the non-primary one(s) (e.g. this might happen for an object very near the detection limit, or if there was some masking only on the primary skycell).

There is also a BestDetection flag in [StackObjectThin](#), which was intended to solve this issue. However, there has been a conversion issue which corrupted the flags, and the BestDetection flag cannot be currently used. We are working on fixing this issue for DR 2.1.

How can I retrieve catalog data with python through the MAST API?

The [new MAST interface](#) to the Pan-STARRS catalog supports queries to both the DR1 and DR2 PS1 catalogs. It also has an associated [API](#), which is used in the script described [here](#).

How can I search for moving solar system targets in the PS1 images and catalogs?

See the [How to search for moving targets in PS1 images and catalogs](#) page for a Python Jupyter notebook that demonstrates one approach to searching the PS1 archive for either images or catalogs that have observations of solar system targets. A MAST interface is currently under development to make this kind of search easier.

PS1 astrometry

How good is PS1 astrometry?


The positions of DR1 sources were originally determined from the astrometric calibration of PanSTARRS using sources in common with the 2MASS catalog. The systematic uncertainty in this calibration is not precisely known, but is likely to be close to 0.1".

However, the positions of the **mean** objects in the DR1 release catalog were recalibrated using Gaia DR1 positions as additional constraints in the astrometric solution. The Gaia measurements were given very high weight, as detailed in [Magnier et al. \(2016\)](#). A comparison of Gaia and PanSTARRS positions for the objects in common yields a typical residual of 5 mas (1-sigma, 2-d), with a slightly higher component in the direction of right ascension. This comparison may not fully reflect PanSTARRS position errors, because of the correlation in the quantities that were compared (Gaia measurements are included in the determination of individual PanSTARRS positions).

A side effect of the inclusion of Gaia DR1 positional data is that the **epochMean** column in the [ObjectThin](#) table, which gives the mean epoch for the RA and Dec measurements, is also affected for objects that included Gaia DR1 data in their position calculations. The result is that the **epochMean** date is often later than any of the PS1 measurements given in the [Detection](#) table. Typically the value for objects that include Gaia DR1 data is close to the Gaia DR1 epoch of 2015.0 = MJD 15023. Note that in these cases, the **epochMean** value is completely unrelated to the mean date for the PS1 photometric measurements.

An independent test of sources with Gaia data that were **not** used in the recalibration was carried out by Makarov and collaborators (V. Makarov, C. T. Berghea, & J. H. Frouard, Technical Memorandum, AADD USNO, 2017). They consider the over 19 million stars for which the *duplicated_source* flag is set in the Gaia DR1 catalog for which a Gaia-corrected position is available in PanSTARRS DR1. Their results, summarized in the first plot below, suggests that the residual systematic uncertainties for recalibrated PanSTARRS positions is closer to 20 mas (1-sigma, 2-d) for sources with Gaia magnitudes between G=15 and G=18, increasing for fainter magnitudes to about 35 mas at the Gaia magnitude limit (G=21.7). Errors are larger towards brighter magnitudes as well, possibly because saturation and proper motion effects become significant in the PanSTARRS measurements.

Comparison with GAIA stars *not* used during the PS1 calibration




PanSTARRS-Gaia astrometric comparison for sources not used in the PanSTARRS calibration. Courtesy Makarov et al (2017).

The majority of the sources in this category have a value of the *astrometric_excess_noise* parameter in Gaia DR1 below 1 mas (median value 0.61 mas). The *astrometric_excess_noise* parameter estimates the magnitude of residuals per observation beyond the nominal measurement errors after the single-object astrometric fit (see [Lindgren et al 2016](#) for more details).

Makarov et al (2017) obtain similar results from a comparison with quasars with high-quality radio positions but without a Gaia counterpart, although the number of sources in this case is much smaller.

Comparison with OCARS quasars without GAIA counterparts

- OCARS: catalog of radio-sources observed in the optical



PanSTARRS-Quasar astrometric comparison for radio sources without a Gaia counterpart. Courtesy Makarov et al (2017).

It is thus likely that systematic uncertainties of 20 mas (1-sigma, 2-d) apply to all mean objects in PanSTARRS DR1, with possibly larger uncertainties in regions with a scarcity of sources in common between Gaia and PanSTARRS.

In addition to this systematic uncertainty, each mean source has a statistical measurement uncertainty reported in the following fields in the ObjectThin table:

raMeanErr	arcsec	REAL	4	-999	Right ascension standard deviation from single epoch detections.
decMeanErr	arcsec	REAL	4	-999	Declination standard deviation from single epoch detections.

Note that the **stack** object astrometry has not been recalibrated with Gaia constraints, and therefore it still includes the larger systematic uncertainty from the 2MASS calibration.

As an additional caveat, stars with high proper motion ($> 1''/\text{yr}$) are likely to be split into multiple mean objects; proper astrometric solutions for such sources require use of individual detections (which will be available in DR2) and epoch-aware co-processing. Similarly, objects with high proper motion might appear elongated in stacked data and may not be properly identified as point sources.

PS1 photometry

What filters did PS1 use?

The PS1 camera surveyed the sky using 5 filters: **g**, **r**, **i**, **z**, and **y**. The effective wavelengths (and spectral resolutions) of these 5 filters are 481 nm ($R = 3.5$), 617 nm ($R = 4.4$), 752 nm ($R = 5.8$), 866 nm ($R = 8.3$), and 962 nm ($R = 11.6$), respectively. Please refer to Table 4 in [Tonry et al. \(2012\)](#) for bandpass details. [Schlafly et al. \(2012\)](#) provides updated zeropoints in Table 1.

How good is PS1 photometry?

The PS1 photometric system is shown by [Schlafly et al. \(2012\)](#) to have reliability across the survey region at the level of (8.0, 7.0, 9.0, 10.7, 12.4) millimags in (**g**, **r**, **i**, **z**, **y**). The Haleakala site is good enough to enable $< 1\%$ photometry over much of the sky. The PS1 photometric calibration pipeline process is described in [Magnier et al. \(2016\)](#).

Which magnitudes should I use?

There are several different kinds of magnitudes in the PS1 catalog (aperture, PSF-fitting, Kron, etc.), and there are also several different sources for those magnitudes (means from multi-epoch measurements, measurements from deep stack images, "forced mean" measurements where a stack image is used to identify objects but the photometry is determined by fits to the single-epoch warp images). Which magnitude should you use for your science?

A [comparison between different photometric measures](#) provides some guidance. The answer is complex, but here is a high-level summary:

- **For point sources use PSF magnitudes.**
 - **Mean PSF magnitudes** have the lowest noise (because the PSF model is most accurate in single-epoch images). They are good for brighter objects, but for objects near the single-epoch detection limit they will be biased (due to the absence of sub-threshold detections), and objects too faint to detect in a single epoch are missing.
 - **Stack PSF magnitudes** are noisier because the PSF model is less accurate. But the stack detections are more than a magnitude deeper and so have many more faint objects than mean detections.
 - **Forced mean PSF magnitudes** use PSF-fitting photometry on the single-epoch images at positions of stack detections. They are a reasonable compromise: they have slightly lower noise than the stack PSF magnitudes, and they are deep and unbiased (because they use data from all warps). Their noise is higher than the mean PSF magnitudes, however.
- **For extended objects use Kron magnitudes.**
 - **Stack Kron magnitudes** are usually the first choice as a general-purpose, deep magnitude.
 - **de Vaucouleurs** and **exponential** model fits could be better in some cases, and the **mean measurements** can be useful for objects that are barely resolved (where the PSF is important).
 - Extended object photometry using the PS1 catalog **will require research and analysis by the user** to determine the best approach.

What are the brightest and faintest stars for which the data are reliable?

The answer is it depends on band and FWHM. A very conservative estimate for the bright limit is as follows:

g	r	i	z	y
14.5	15	15	14	13


You may find individual fields where you can do up to a magnitude brighter than this.

The faint limits vary across the sky - see the [section on the photometric depth](#).

How reliable is PS1 photometry of nearby, bright galaxies?

There are two main problems you should be aware of which affect bright galaxies:

(1) over-subtraction of the sky background - this affects most Messier galaxies and some (most?) of the larger NGC objects. As the background is subtracted independently from each filter, this can induce colour changes across the galaxy. In principle the polynomial used to store the background is stored and can be de-applied. However, as the subtraction is done at the warp stage, if you want stacked data the stacking procedure would have to be re-run. This facility is not available in DR1.




Messier 31 from the 3pi stacks

showing an extreme example of

the over-subtraction of the background.

(2) issues with the [row-by-row bias and continuity corrections](#) - these are applied and therefore seen on individual ccd scales (~600 pixels). As this is done at the detrending stage this is currently very difficult to correct. There is an aspiration to re-analyse all bright galaxies but this has not yet been done.



Messier 106 - this is an extreme stretch of the r-band stack showing problems on the scale of individual ccids.

Why is my photometry from images so different from the catalog photometry?

The pixel values in PS1 images must be [transformed using a non-linear equation](#) to convert them to fluxes before doing photometry on them. If you do not apply that transform, the photometry results will not be remotely close to the correct photometry. (Note that this transformation is required for full sky cell images; smaller cutout images have already been converted to linear fluxes.)

Even after applying this non-linear transformation, photometry from the images will not match the catalog perfectly. The [ubercal](#) procedure was used to modify the photometric calibration using sources observed in multiple images. That allowed correcting small errors in the photometry that depended on the exposure conditions, the epoch of the observation, the position in the detector, etc. Since this calibration was completed after the creation of the images, the image FITS headers and flux scaling do not reflect the improvements. Measurements made directly from the images will consequently not be as accurate photometrically (or astrometrically) as measurements in the catalog.

Flux differences that originate from post-image calibration should be small, however. Gross differences in photometry are much more likely to result from not applying the non-linear flux transformation.

PS1 Images

How can I convert pixel values on the stack images back to the original counts?

The short answer is you can't really get from the stack pixel values back to measured counts.

The long answer is it depends on why you want these values. The stacks are a weighted combination of the single-epoch warp images, with some scaling to get the zero-point to 25, so although you can get back to the numbers on the stack frames, it is not clear how these will relate to the warp counts, although they probably agree within a factor of 2 or better.

It is not clear that you could ever reverse engineer the stacking to recover the actual sum of the input warp counts (although there is some information in the image headers which might help).

It is actually worse than this in some cases, because the exposure time recorded is *not* the true exposure time (warps can be dropped from the stacking process after the total exposure time has been calculated). In this case the numbers in the stack could be quite different from warp counts.

Also, sky has been removed, so if you wanted the number of total counts (object + background) to do photon counting S/N stats that would be tricky.

With these caveats, to get back to a version of the numbers on the stack frames, you can use this equation:

$$10^{((\text{zeropt} + 2.5 \cdot \log_{10}(\text{exposure_time}) - \text{catalogue_mag}) / 2.5)}$$

where the zeropt is 25 for the stacks.

Why do some catalog objects apparently not have images available?

The PS1 catalog includes some sources south of the nominal -30° declination limit of the PS1 3PI survey. If the positions of those sources are entered in the [PS1 Image Cutout Service](#), it reports that there is no data available because they are south of the declination limit. This naturally leads to the question of why an object in the catalog does not have an associated image.

The PS1 project did take a small amount of data south of -30 degrees declination. We are not sure about the history of these fields – they could have been either accidents or experiments. When there were exposures available, the catalog pipeline software did process the images and extracted sources from them.

However, the image processing software did not produce any stacked images (combining the exposures) for these southern fields. The data only got processed to the single-exposure level, so there presumably are no stack detections or other data products derived from the stack images. There are some individual exposure images in these regions, but those will not be available for download until the DR2 data release (which we are hoping will be completed in January 2019.)

That is why there can be sources in the catalog that do not have associated stack images available. These sources should only occur in the far southern sky.

How do I convert pixel values into magnitudes?

This is explained in [PS1 Stack images](#), the "Photometric Calibration" section.

Why are the coordinates incorrect when I display a skycell image in DS9?

If the coordinates in your image appear to be far off (by arcminutes or more) from what you expect, you probably have encountered the RADESYS problem using recent versions of DS9. The full skycell FITS images do not have a RADESYS keyword. That leads some software to incorrectly interpret the PS1 image coordinates as being equinox 1950 rather than equinox 2000. At the moment the only known software with this issue is DS9 v8, but it could happen with other software as well. (Older versions of DS9 correctly interpret the PS1 FITS headers as J2000 coordinates.)

The fix is to insert the keyword RADESYS = 'FK5' in the header. PS1 FITS cutout images have a correct RADESYS keyword (as of 2019 March 13), but full skycell FITS images do not.

We hope eventually to fix this issue (and some other problems described in the "FITS image format quirks" section of the [PS1 DR2 caveats](#) page), but since the PS1 archive includes 1.5 petabytes of PS1 images, the task of updating all the image files is not simple.

PS1 Timing

What is the time scale used for the PS1 observation times? Why do the observation times in the image header and catalog disagree?

The times in the warp image headers and in the catalog (e.g., the obsTime column in the [Detection table](#)) are defined using international atomic time (TAI) rather than UTC time. Those times differ by the addition of leap seconds, which leads to header times that differ by 34 or 35 seconds from the UTC times. (See [Rots et al. 2015](#) for more details.) If you are concerned with timing at this level, you may need to convert the times to UTC. For the warp images, the fix for this is to insert the keyword TIMESYS = 'TAI' in the header. FITS cutout images have a correct TIMESYS keyword (as of 2022 January 20), but full skycell FITS images do not have a TIMESYS keyword. For the database times, the fix is to subtract 34 or 35 seconds (depending on the date) from the TAI time to get the UTC time. The `astropy.Time` module includes convenient functions for converting between TAI and UTC times.

Also note that the epoch given in the FITS header for the warp images is the start time of the observation, while the epoch in the database for `Detection` entries is the mean time of the observation, which is the start time plus 15 seconds (since the exposure time for PS1 images is 30 seconds). So you should add 15 seconds to the warp header's MJD-OBS keyword value to get the equivalent value from the `Detection.obsTime` column in the PS1 database.

Thanks to Peter Van Wylen for discovering the TIMESYS issue and helping identify the fix. And thanks to Jules Halpern for pointing out an error in the original description of the fix.