

NAGIOS Services for ISS

Overview of NAGIOS services for ISS

Most data managers will have several service checks in common, and there are explanations for some of those services below. The section after that describes checks more specific to PERDIGAO.

HTTP

Verifies that the web server on the data manager is responding. This is not critical for data flow, but if it goes down then the nagios status page will not be available.

Mirror Raid

For systems with a mirror RAID filesystem for the system and data, this checks the health of the RAID. If one disk in the RAID should ever fail, this check should indicate critical.

NTP

The data managers use the Network Time Protocol to synchronize the system clock with known good times from time servers on the Internet. This is not always stable, since the delay over the satellite link is significant. If this stays critical for more than a day, then it needs to be investigated.

Surface Broadcast

When a data manager is ingesting data from a surface weather station, it broadcasts the latest values on the local network for GAUS. This check verifies that the broadcast is being received, that it has a recent timestamp, and that the values look reasonable.

USB Backup

Twice a day the system and all the data are backed up onto the external USB drive, and the backup script generates a log file in `/iss/log/usb-backup-USB_BACKUP.log`. If the log file has not been updated, appears to contain errors, or shows that the USB disk is filling up, then this check becomes critical.

If this fails, it does not affect the data flow, but it does need to be fixed. One thing to try is to unmount the disk (if mounted), unplug it and make sure it powers down, then plug it back in again. It is safe to see if the disk can be mounted through the desktop USB icon, but it should be left unmounted when not in use. The `usb-backup` script will mount it automatically when it runs the backup.

Data file checks

Each instrument on the ISS network typically creates multiple kinds of data files. These data files may be created by a program or script on the data manager, or they may be copied from another host using `rsync`. Usually all of the data files accumulate under the data store directory, `/iss/ds`.

A python script periodically scans the data store directory for new files. It recognizes each kind of file and stores metadata about the files into a database, `/iss/ds/iss_catalog.db`. In particular, the script reads and decodes each kind of file to determine the time coverage for that file. It also knows how often the files should appear (sample density) and how quickly (latency). Given these parameters for all the different kinds of data, the script can determine if any category of data appears to have stopped flowing.

Since the data files get copied from various sites back to EOL, the nagios checks have a lot of context information in their name. This is how they break down:

- `site_check_category`

Here is a check running on the `iss1` data manager at the profiler site:

- `metcraxii/profiler_latency_lidar/ppi_cnr/image/png`

The site is known as 'metcraxii/profiler', the check is for 'latency', and the category is 'lidar/ppi_cnr/image/png'. All of the LIDAR checks have a prefix of 'lidar' in the category name, while all of the profiler checks have a prefix of 'profiler' or 'mapr'.

Aggregate Checks

Rather than generate alerts when any one of the categories fails a check, NAGIOS only generates alerts for special aggregate checks called 'anyfail' and 'allfail'. For example, it is impossible to predict which kinds of files the LIDAR will generate and when, so as long as some of the lidar categories have recent data, the 'metcraxii/profiler_allfail_lidar' check will pass. Conversely, all of the profiler categories should be very steady, so 'metcraxii/profiler_anyfail_profiler' fails as soon as a single category fails. This means that while there is a problem with synchronizing the 'profiler/449/rim/mat' files, the aggregate profiler check fails also.

NAGIOS monitoring at METCRAX-II

Here are specific notes for each of the METCRAX-II sites involved in data flow.

Profiler site

This instance monitors data from the profiler, lidar, Campbell data logger, and WXT. The Campbell and WXT data are ingested directly on the data manager, while the profiler and lidar data are downloaded periodically through rsync. There are nagios services for all of the individual categories of data, even when there are multiple kinds of data files for one instrument. Thus there are a dozen services for LIDAR data, and several for the profiler. These are grouped into three aggregate checks:

- metcraxii/profiler_anyfail_profiler
- metcraxii/profiler_allfail_lidar
- metcraxii/profiler_anyfail_surface

WXT and Campbell Surface Data

The aggregate surface check includes the WXT and the Campbell, so this check fails if either one is out. There is only one category for the Campbell data, surface/campbell/nc, corresponding to the netcdf files which are generated directly by the cam_ingest program. The WXT data however, are first recorded in NIDAS raw format with sample ID 1.10 (raw /nidas_1_10), then parsed into actual variables with sample ID 1.11 (raw/nidas_1_11), and finally written to netcdf files by the data_nc program (surace/wxt/standard/nc). So if the WXT itself cannot be read, then all three downstream categories of data will stop.

Of the three WXT checks, the most critical is the nidas_1_10, since that represents the raw samples read directly from the WXT. If that check is failing, then something is wrong in the connection between the data manager and the WXT. This has been a problem over several days. It seems that cycling the power or re-plugging the serial connection on the WXT up at the LIDAR restores the data flow, but then it gives out several hours later.

Profiler

As mentioned above in the section about aggregate checks, all of the profiler categories should be very steady, so 'metcraxii/profiler_anyfail_profiler' fails as soon as a single category fails. This means that while there is a problem with synchronizing the 'profiler/449/rim/mat' files, the aggregate profiler check fails also. A failure in one or a few categories probably means the processing is behind or a data rsync is failing for some reason, but there is probably nothing that the current operator can do about it. Gary will get an email and he can work on it if it's a problem.

The profiler process running on the profiler computer is **not** checked directly by NAGIOS. If all of the profiler categories are failing, then certainly a hung profiler process or similar error could be the cause. Either way, the profiler computer still needs to be checked separately.

LIDAR

As mentioned above in the section about aggregate checks, it is impossible to predict which kinds of files the LIDAR will generate and when, so as long as some of the lidar categories have recent data, the 'metcraxii/profiler_allfail_lidar' check will pass. If all the checks fail, then either the LIDAR has stopped operating, there has been a network interruption between the data manager and LIDAR, or the rsync to the LIDAR is failing. A network interruption would also show up as a failure of the ping check to the LIDAR host.

GAUS site

EOL

This is where all the data converge and things really get complicated. There are essentially three data contexts at EOL: the subsets of data copied from each of two metcraxii sites, plus the set of images and plots which are generated from the data and copied onto the web. If data transfers from the sites start to fail, or the web processing starts to fail, then the corresponding NAGIOS service checks running on snoopy will start to fail.

- [Overview of NAGIOS services for ISS](#)

- HTTP
 - Mirror Raid
 - NTP
 - Surface Broadcast
 - USB Backup
 - Data file checks
 - Aggregate Checks
- NAGIOS monitoring at METCRAX-II
 - Profiler site
 - WX
T
and
Ca
mp
bell
Sur
fac
e
Data
 - Pro
filer
 - LID
AR
 - GAUS site
 - EOL