

Australia News

NOVEMBER, 2023

Contents

Contents.....	1
From the GM's Desk.....	3
KISTERS Company News.....	3
Dedicated downloads website	3
KISTERS Instruments.....	3
Australia User Group meeting 2023	3
Australia User Group meeting 2024	5
WISKI Product News	6
KISTERS expands the Hydrological Data Management System in the Philippines	6
WISKI Service Desk and release planning.....	8
WISKI HTTP/TLS Certificate: Expiration and renewal	8
Fade-out of the Kiscript IDE.....	9
The Murray-Darling KiWQM Plan	9
KISTERS Product News	13
KISTERS Raven	13
Hydstra Product News	13
Hydstra 14 - call for feature requests.....	13
Hydstra Releases.....	14
Hydstra Systems and Patches (v12, v13) are available.....	14
We're Going to Need a Bigger Boat!	14
HYZIPDIRS support SVRIMP1 though SVRIMP4.....	14
V14.00 - heads up: Hydstra will no longer override the TEMP and TMP environment variables	14
Distributing new HYACCESS files	15
V14.00 - heads up: changes to how ARI axes are labelled on probability plots.....	15
Experts As A Service.....	15
Clearer user attribution for SVRRUN jobs	15
Time Series Images	16
Tailor HYPERUPD to Your Needs.....	17
Checking Your HYDLOG Files for Errors and Warnings.....	17
Ignoring Warnings about Data Discontinuities.....	17

Organising Backups in Notepad++	17
Importing Documents into DOCPATH	18
HYXPLORE Handy Shortcuts.....	19
HYRATED – Global Mouse Mode	19
HYAUDIT Default Datasource	19
V14.00 – changes to HYAUDIT Period Data Within Range test	19
Hydstra V14 - What do you want in it?	20
Some General Observations on Password Policies	20
WEBSERVICE.INI feature to control default values for optional parameters to webservice calls	20
V14.00 - HYAUDITOR replaces HYAUDIT.SUMREP	21
V14.00 - Style Upgrade for HTML Reports.....	21
V14.00 – Changes to default Windows font for HYXPLORE, parameter screens, HYMANAGE etc	23
V14.00 - HYMANAGE Delete Transaction Logs.....	23
V14 - User Constraints in PASSWD	23
Hydstra Administration Service	23
Proposed DEFAULT User Expiry in V14.....	24
Marking Missing Data in the Data Manager’s Workbench.....	24
Some new keyboard shortcuts in the Data Managers’ Workbench.....	25
Hydstra V11 Unsupported	26
Who Deleted My Database Records?.....	26
HYLATEST enhancements	26
HYMANAGE Date Format	26
Reminder to Open Files in Offline Mode When Possible	26
V14 - Storing the Analysis Date and Time of WQ Results.....	27
Keeping SVRIMP Folders Under Control.....	27
Modernising your Hydstra Fonts	27
AUTOJOB.BAT routine to submit tasks to SVRRUN	28
WEBSERVICE.PL withdrawn in V14.....	28
Avoid Locking Telemetry Files in HYDMWB.....	28
Cleaning up IIS Log Files with HYZIPDAT	29
Backup up Hydstra using HYCLONE and ROBOCOPY	29
Chromium browser support in HYXPLORE coming in V14.....	29
SVRRUN has a high-performance task ingestion method based on SVRSUB.PRM files.....	29
Dealing with rating table exceptions	29
SVRIMP attempt to press on with other variables for the affected site if one of them has an invalid/missing data trans	30
HYCSV lets you configure a single header row, or no header row at all	30
Staff News	31
Information	32

From the GM's Desk

Dear colleagues,

The countdown to Christmas has almost begun! Our recent re-launch of the KISTERS group brand including instruments has also involved a new web site at www.kisters.com.au and a new customer resources portal at <https://resources.kisters.com.au>. Access to the resources portal is via username and password, details within this newsletter.

We recently had the first post covid US KUG in September at San Diego, it was great to see so many users and user presentations. Prior to KUG Dylan and I took the opportunity to visit several US clients, we hope to do the same next year. We are planning a KUG in Canberra Australia on 19-20 June 2024, followed with an African KUG in August.

I really like the article of photos being available in the Hydstra data managers work bench, this feature is a great resource for validating data management decisions and providing imagery in the context of events.

We love data.

By Paul Sheahan, General Manager, KISTERS Pty Ltd

Paul Sheahan
General Manager
KISTERS Pty Ltd



KISTERS Company News

Dedicated downloads website

We have launched a dedicated website for KISTERS downloads & resources at resources.kisters.com.au.

You can use this website to download WISKI and Hydstra software releases, patches, documentation, KISTERS Australia newsletters, recorded sessions from User Group meetings, training videos, and other resources as they are published.

If your organisation hasn't yet received login details for the Resources website please contact us at conflks.atlassian.net/servicedesk or hydstra.support@kisters.com.au.

KISTERS Instruments

Effective from the 1st of June 2023 HyQuest Solutions Pty Ltd are now known as KISTERS Instruments Pty Ltd.


Australia User Group meeting 2023

Yet another successful KISTERS User Group meeting was held in Sydney on 22 May 2023, the day before the AHA conference. Thanks to everyone involved. We recorded most of the sessions and we have published them on the new KISTERS Resource website.





A fun highlight of the User Group was the special prize awards of iFly experiences, which went to Joe Gendall, Georgia Spankie, and Chris Michl.



Australia User Group meeting 2024

Hold the date, the 2024 Australian User Group will be held in Canberra on **19-20 June 2024**.

WISKI Product News

KISTERS expands the Hydrological Data Management System in the Philippines

KISTERS is enhancing flood forecasting in the Philippines through the "Hydrological Data Management System (HDMS)" project with the Philippine Atmospheric, Geophysical, and Astronomical Services Administration (PAGASA).

This initiative aims to modernize the essential infrastructure for accurate hydrological forecasting, vital for safeguarding flood-prone communities nationwide. The software component of the project is composed of various KISTERS solutions, with WISKI at its core. The Array Storage is used to store and provide the raster data sets for hydrological forecasting, and WISKI Web serves as a web-based tool for data visualisation and publication.

The project has been developed in two phases, both of which have been successfully delivered and implemented, resulting in the integration of 8 out of the 18 river basins into the system.

In the first phase of the project, executed between May and October 2022, the implementation of HDMS was carried out in five locations, including the Hydro-Meteorology Division in Manila and four River Flood Forecasting and Warning Centers situated in Pampanga, Agno, Bicol, and Cagayan de Oro.

Phase II expanded the system to encompass three additional locations: Cagayan, Davao, and Pasig Marikina. This phase unfolded between March 2023 and October 2023, and included on-site training sessions at each of the 3 locations. For this purpose, KISTERS' representatives, Chris Michl and Melisa Isgró, were on-site in August conducting intensive training sessions at each location (see pictures 1 and 2).



Training in Tuguegarao (Northern Luzon) and Davao City (Mindanao)

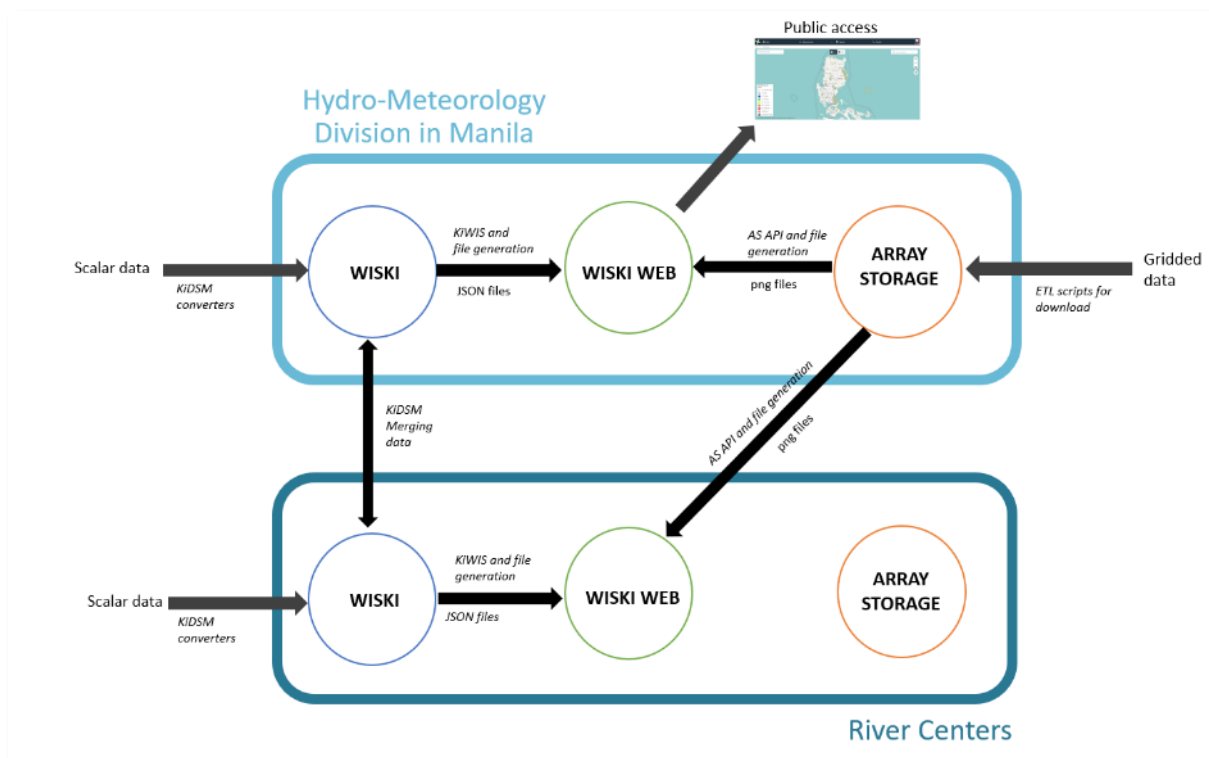


Training groups in Tuguegarao (left) and Davao City (right)



Training group at Pasik Marikina (Manila, headquarter HMD). Each course was delivered at each location over three days with a mentimeter test at the end of each day (great fun)

Regarding the system's architecture, the HDMS is designed to support the independence of each river center in the Philippines, including the Hydro-Meteorology Division in Manila. This approach is crucial in a flood-prone nation. It ensures that each river center can operate independently even during technical disruptions. In this setup, the Hydro-Meteorology Division in Manila manages the main WISKI system, serving as a central hub for receiving data from all river centers. Conversely, each river center has exclusive access to its locally generated data. This system facilitates continuous data exchange between Manila and the river centers, enhancing PAGASA's operational resilience. As it is shown in the following diagram each river center has its own internal WISKI web and Array Storage. However, only the Hydro-Meteorology division in Manila will be publishing the data for the public on their website.



WISKI Service Desk and release planning

We are excited to present the WISKI product release plan concept, offering two distinct branches for customer releases. These branches are designed to cater to different aspects of your usage and provide you with a range of options for staying up to date with our product offerings:

WISKI LTS (formerly "WISKI Vintage"), recommended for existing customers, with focus on critical bug fixes and security updates, and long-term support for at least two years after initial released as LTS version.

WISKI Product, our main branch, perfect for ongoing projects, will see regular release every three months, with additional service packs every four weeks, also offering paid enhancements. Please see below our general release plan with the latest version for each main branch for the next month:

Version and SR (Service Release)	Date	Remark
7.4.13 SR7/SR8/SR9/SR10/SR11	since May 2022	WISKI LTS - most on-going versions
7.4.13 SR12	since Nov 2023	WISKI LTS - latest LTS version
7.4.15 SR4	since Sep 2023	WISKI product - current main version
7.4.15 SR5	available in Dec	WISKI product

The 7.4.13 version has now been moved into the WISKI LTS version, with the latest version, 7.4.15, becoming the new WISKI main branch. Version 7.4.13 is known for its established stability, while 7.4.15 is set to bring more enhancement and new features. Currently, our focus is on deploying 7.4.13 SR11 and SR12, along with 7.4.15 SR4, to systems that have yet to be updated. We look forward to collaborating with you to determine which version aligns better with your requirements in our upcoming planned update.

In short, if you are not in the process of planning with us new functionalities in WISKI/KiWQM or KiECO, we suggest to stick to the 7.4.13 version throughout 2023. In all other cases, we will discuss with you individually the update path.

WISKI HTTP/TLS Certificate: Expiration and renewal

The WISKI Server encrypts all the inter-process and inter-system communication such as Client-Server communication based on the TLS "Transport Layer Security" encryption protocol as it is used in HTTPS communication.

The authentication, validation and encryption are performed with the help of the so called "digital certificates" (private keys for the server and public key for the client) that are validated and signed by an official trusted certification authority such as VeriSign or Microsoft or in some cases via the generated WISKI Server CA (Certification authority) self-signed certificate.

All certificates have limited life time. In the WISKI default, root CA certificates are created with a life time of 10 years and self-signed certificates with a life time of 3 years. Before the end of life of a certificate, the certificate must be renewed. If the certificates are not renewed, the communication processes cannot establish a secured communication leading to a system failure.

Which WISKI releases are affected?

All the WISKI systems starting from 7.4.11 SR7 need to be checked as described in this notification.

How to check the expiration date of certificates

Currently, the best way to check the validity of the server certificate chain is to use the system browser by calling e.g. "https:// [WISKI-Server-Hostname]:7479" - replace "WISKI-Server-Hostname" by your FQDN hostname to avoid TLS warnings in the browser.

The certificates including the expiring date can be inspected by using buttons on the left of the browser URL - the exact implementation differs and depends on the browser.

Please refer to the following page: <https://www.globalsign.com/en/blog/how-to-view-ssl-certificate-details> for more precise instructions for different browsers.

In the near future, our development teams will provide a tool to monitor, check and assist the renewal of certificates.

Until then and in case of an approaching certificate expiration, please get in touch with our support team to help with the renewal.

Fade-out of the KiScript IDE

The KiScript IDE is an Eclipse-based GUI to execute KiScript commands on the WISKI Server. It was used in the past by a small number of customers to run customized scripts for multiple purposes such as data analysis, data collection or reporting.

The current IDE application is based on the Eclipse 3.x framework which is outdated and not supported anymore by the Eclipse Foundation.

We will not upgrade to the current 4.x version of Eclipse due to high technical efforts and breaking changes in combination with the small user base of the IDE.

Since the KiScript IDE will not be part of any future WISKI release, we like to inform you about the related impact on your system based on the following cases:

1. You are not using the KiScript IDE and it is not installed on your system:

No action needed.

2. You are not using the KiScript IDE but it's installed on your system:

Please uninstall the application.

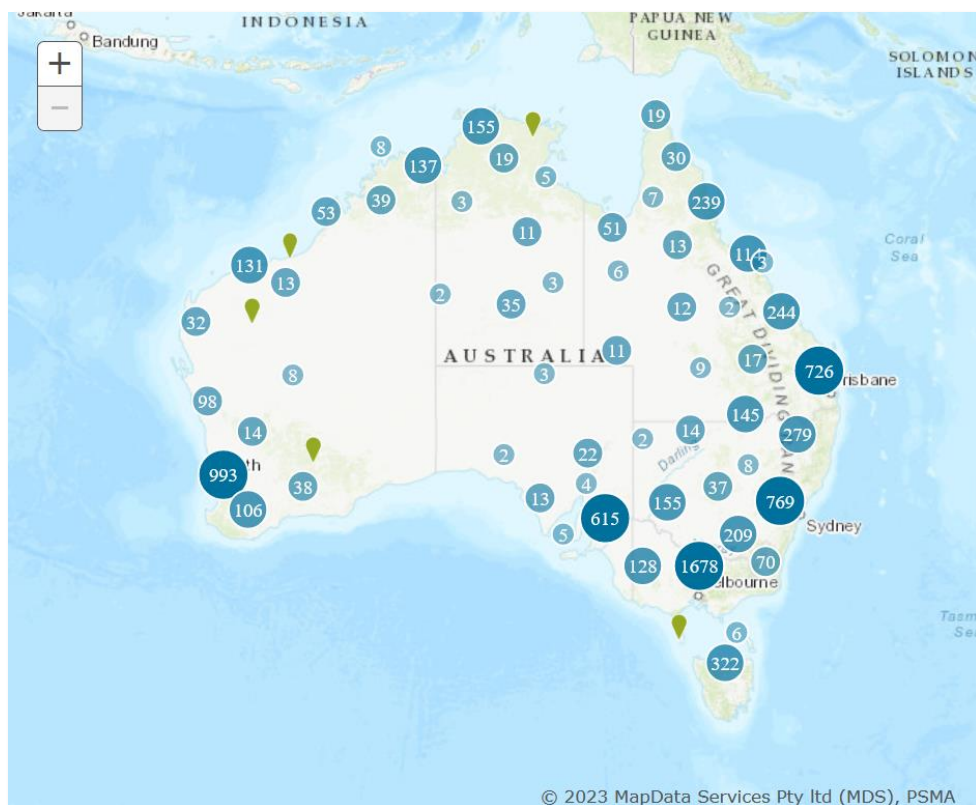
3. You are actively using the KiScript IDE:

Please contact our support team for the upcoming update on your system and send us your current use-cases and requirements. We would like to involve you closer in the decision, design and planning of alternative solutions.

The Murray-Darling KiWQM Plan

Through the Water Act of 2007, the Bureau of Meteorology was required to, “collect, hold, manage, interpret and disseminate Australia's water information” – no small task.

Organisations across the states and territories are required to submit certain data to the Bureau in the xml-based Water Data Transfer Format (WDTF), specifically designed to provide ample information to the BOM in a standard format for the meaningful storage and publication of data from all across the country.



Water Data Online - <http://www.bom.gov.au/waterdata/>

The end-result is a national system underpinned by data being delivered in a series of standard formats from many different agencies, chiefly WDTF. To the dismay of water quality scientists, no such legislated and enforced standardisation exists for discrete water quality data.


The Murray-Darling Basin Authority (MDBA) – while smaller in scale to the Bureau – are also curators of massive amounts of data from a range of organisations, including discrete water quality data.

As existing WISKI customers, a migration project was recently delivered using the **KISTERS Water Quality Management (KiWQM)** module that had a series of key objectives:

- Standardise historic datasets from legacy systems in a data cleansing project
- Migrate historic data from legacy systems to KiWQM
- Automate the regular transformation and import of data received from different organisations and source systems
- Enable access to up-to-date water quality data through web products (Portal and API access)

The migration itself was completed in close consultation with data managers at MDBA to clean up historical datasets and start afresh in KiWQM. To aid in dealing with large amounts of metadata-rich source data, an intermediary SQL Server database was used to aid in the standardisation of the process, enabling each transformation to be scripted and reproduced if needed in a timely manner.

Why SQL Server? Well, if you've ever tried to run vlookups and nested IFs against tables in Excel with millions of rows and hundreds of columns, you'd know why. Unfortunately, MS Access didn't cut the mustard either:



After raw data was collated and indexed in SQL Server, strict rules were developed to standardise, including (but not limited to):

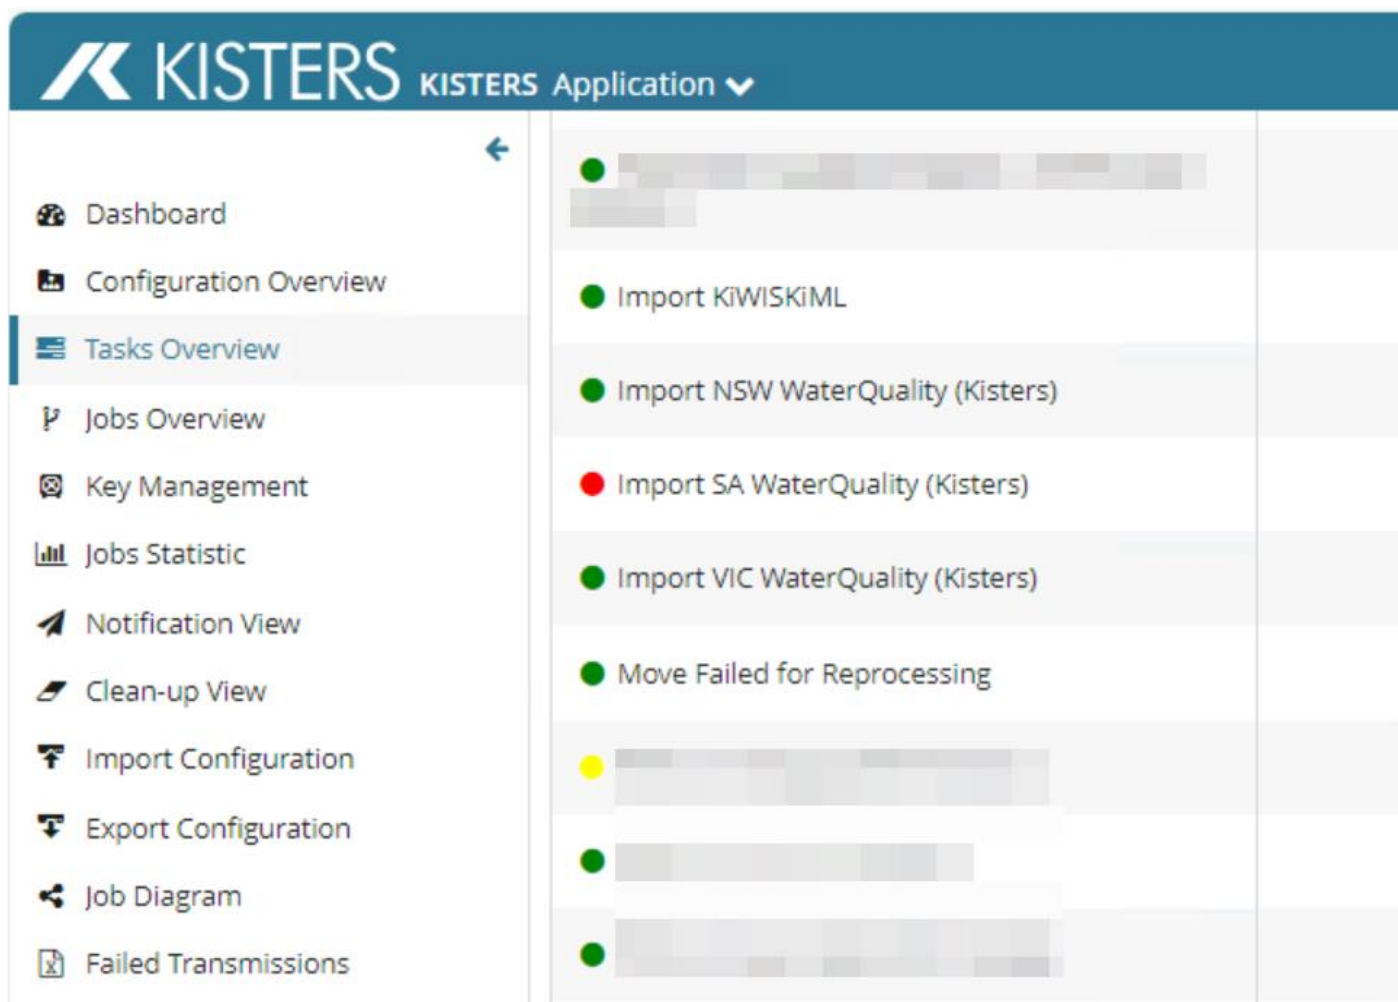
- Removing dud-samples that had no values or datetime
- Generating unique sampling event IDs based on station and date
- Generating unique sample IDs based on sample and replicate IDs from different sources
- Delineating sampling events into different measuring programs

Automation

A migration in a live system is only as good as its facility to accommodate new data. Each rule developed needed to be implemented as functions to be applied to raw files from the systems providing the MDBA with data.

A series of Python scripts were developed relative to each provider. These scripts retrieve raw data, then systematically apply each rule to transform the data into a common format, predicated on the creation of backwards-compatible standardised IDs. Once transformed, the scripts process the files through the WISKI batch interface, importing data into KiWQM as water quality samplings. Should the same data arrive 6 months later with enhanced metadata information, the original data will be updated with an auditable record through KiWQM's audit trail functionality.

This entire process is triggered through the **KISTERS Distributed Service Management (KiDSM)**. All components in the process are scheduled over KiDSM, meaning the data manager need not manually trigger anything, with exceptions and errata being captured through additional logging and row-by-row failed file extracts.



The KiDSM interface's Task Overview panel.

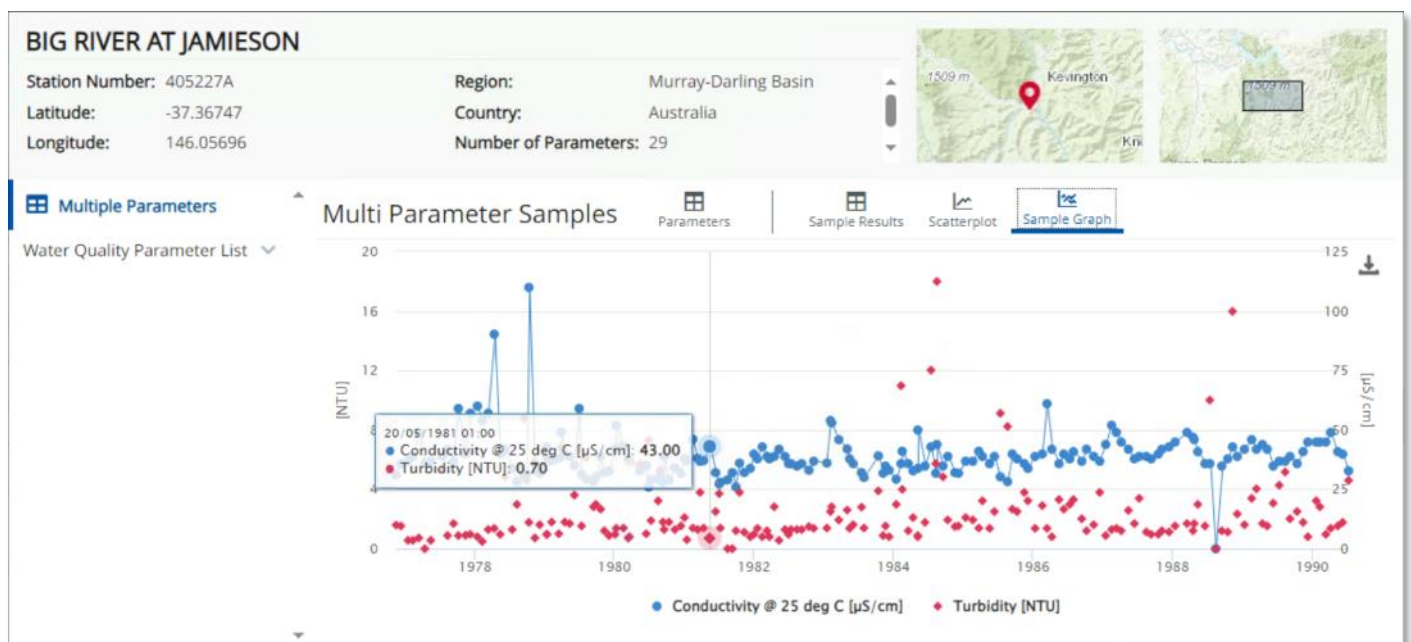
Publishing

Data can be accessed/delivered in a number of ways:

- The full complement of station, sampling-event, sample and result data can be explored through the KiWQM client with a range of analysis, validation and export functions
- Data can be accessed over the **KISTERS Web Inter-operability Solution (KiWIS)** API, allowing broad access to complex data for anyone in the organisation in a variety of formats, including virtual time series based on selected sampling parameters
- The **Water Quality Viewer** Portal application provides a web-based front end for the intuitive exploration of data, also providing export functions for users

ts_path	DissolvedOxygen	
rows	157	
<hr/>		
timestamp	value	value_quality
1976-11-25T09:00:00.000+10:00	8.4	65
1976-12-09T09:00:00.000+10:00	8.5	65
1977-01-07T09:00:00.000+10:00	8.4	65
1977-02-11T09:00:00.000+10:00	0.0	180
1977-03-17T09:00:00.000+10:00	6.6	65
1977-03-22T09:00:00.000+10:00	8.1	65
1977-04-07T09:00:00.000+10:00	11.0	65
1977-05-03T09:00:00.000+10:00	0.0	180
1977-06-01T09:00:00.000+10:00	10.1	65
1977-06-16T09:00:00.000+10:00	10.6	65
1977-06-30T09:00:00.000+10:00	11.3	65
1977-07-26T09:00:00.000+10:00	11.5	65
1977-08-04T09:00:00.000+10:00	10.0	65
1977-09-07T09:00:00.000+10:00	10.4	65

An HTML output of the KiWIS for KiWQM extension, building a virtual time series from discrete sampling data.



A plot of discrete sampling data in the Water Quality Viewer (Turbidity and EC).

Outcome

The new system provides MDBA with a way to import, explore and publish a highly diverse array of discrete water quality data from different sources on an automated basis, but does this without placing demands on data managers to curate this data. Importantly, through the extensive rulebook developed, it also doesn't brush errata under the rug to surprise an unwitting data user later on.

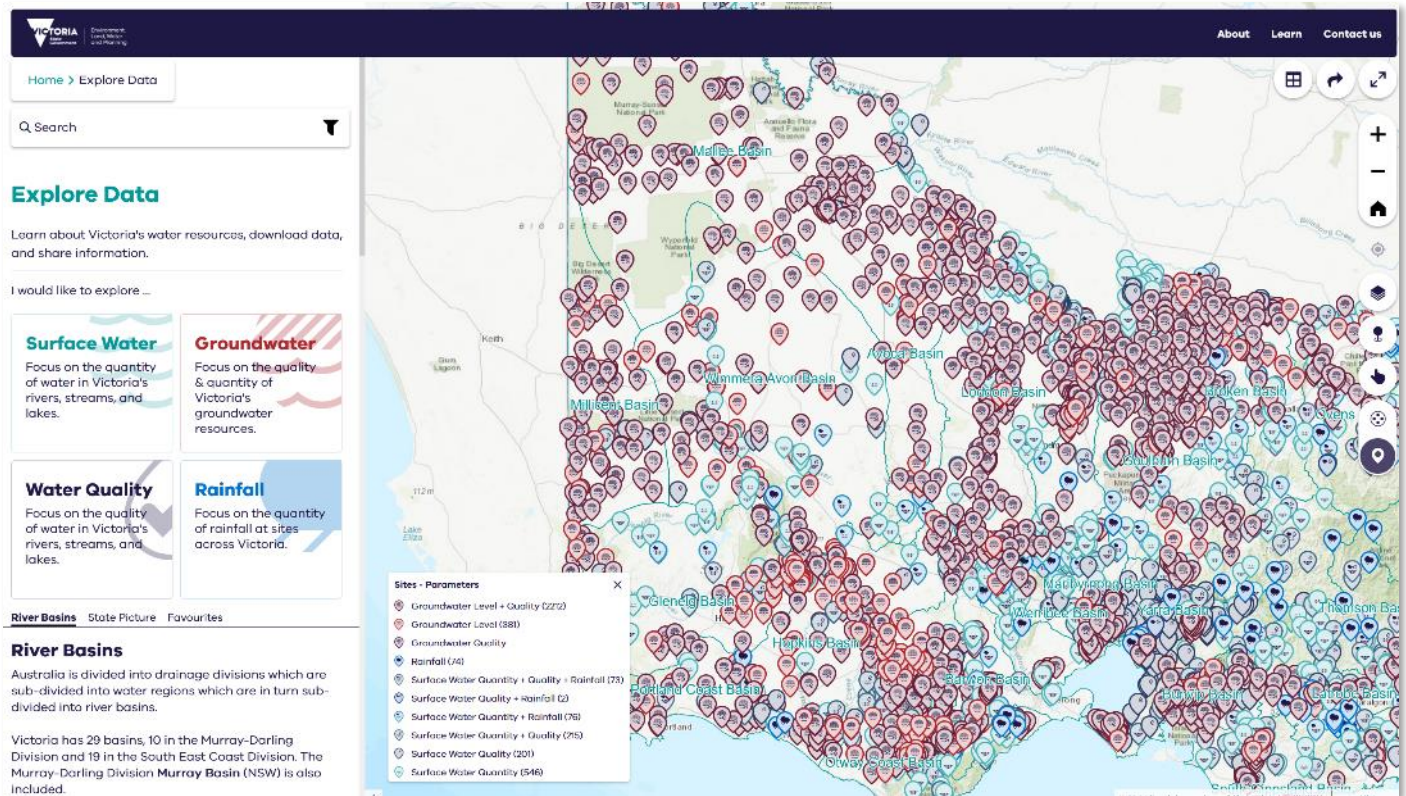
KISTERS Product News

KISTERS Raven

KISTERS Raven is a new digital solution for delivering innovative data visualisations on the web.

To create Raven we reengineered the best elements of WISKI Web and Hydstra Web into a single integrated Web interface designed to run off any data source, Hydstra, WISKI, Hydrotel, Datasphere, or other.

The first implementation of Raven has been designed collaboratively with the Department of Energy, Environment and Climate Action (DEECA). Their implementation is deployed over Hydstra.



Our colleagues in New Zealand delivered a presentation at the NZ User Group meeting in August 2023 demonstrating a Raven Web interface running over live Hydrotel data.

We are offering KISTERS Raven to existing Hydstra Web and WISKI Web clients, for discussions now, and with implementation delivery schedules starting first quarter 2024 for Hydstra Web clients and third quarter 2024 for WISKI Web clients.

To request an upgrade from Hydstra Web or WISKI Web to KISTERS Raven, get in touch.

Hydstra Product News

Hydstra 14 - call for feature requests

We are looking to finalise and lock-down what new features will be available in Hydstra 14 when it is released in the first quarter of 2024.

What new features would you most like to see in Hydstra - more analysis and reports, extra features in the Workbench, anything about ratings and gaugings, more fields (spare or otherwise) in tables, additional tables?

Now is the time to get your ideas in if you haven't already. Please log suggestions and ideas in Jira [Service Desk](#) or by email to hydstra.support@kisters.com.au.

Hydstra Releases

- v13 - Current release - fully supported, patched weekly.
- v12 - Previous release - serious bugfix patches only, please schedule an upgrade.
- v11 and below – Unsupported, please upgrade.

Hydstra Systems and Patches (v12, v13) are available

Hydstra is upgraded continuously through frequent delivery of patch updates. We recommend you apply patch updates on a regular schedule.

The latest Hydstra full system versions and patch updates are available to licensed users at resources.kisters.com.au/resources/hydstra-releases/.

Contact the [KISTERS Service Desk](#) if you don't yet have access to <https://resources.kisters.com.au>.

We're Going to Need a Bigger Boat!

Fans of the movie Jaws will never forget that classic line! Heavy users of SVRRUN might find that their server is running too hard – you don't really want to see more than 80% usage for long periods. SVRRUN can support up to 32 concurrent processes on a 32-core server, so don't scrimp on the main server. If you do update the server you should update the number of threads in SVR.INI, and restart SVRRUN to start using all the additional processes.

HYZIPDIRS support SVRIMP1 though SVRIMP4

Some users run multiple versions of SVRIMP to spread the import load, using SVRIMP1.EXE through SVRIMP4.EXE. Each version consults its own INI file (SVRIMP1.INI for example) and should have its own import folder. As SVRIMP processes incoming files they are moved to dated subfolders for future reference. These folders can quickly build up to hold potentially millions of files, and to keep the disk size down and backup processes fast, they should be zipped up after a while. Program HYZIPDIRS should be scheduled to run nightly, either in AUTOJOB or via a recurrent SVRRUN task.

V14.00 - heads up: Hydstra will no longer override the TEMP and TMP environment variables

Windows maintains two environment variables called TEMP and TMP, which both point to the user's temporary files folder. Why are there two? Complicated answer that goes back to the 1970s and CP/M (the precursor to DOS). Previously Hydstra programs have overridden these variables to point to the user's TEMPPATH Hydstra folder.

Recently, this has started causing some unwanted consequences.

Any non-Hydstra process that is either launched by Hydstra, or starts up from within a Hydstra environment, will inherit these overridden values, and use your TEMPPATH folder for their own temporary files. And since Windows doesn't know that TEMP and TMP have been overridden, it doesn't know to clean up these files to stop clutter and wasted disk space.

Some users have noticed that their TEMPPATH folders accumulate files/sub-folders like these:

- Edge_BITS_xxx and Chrome_BITS_xxx folders
- TCDxxx.tmp folders
- aria-debug-xxx.log files

In the first two cases, these can seriously chew up disk space, and it won't be released by the Windows Explorer's Properties / Disk Cleanup procedure.

For reasons like this, in v14.00 we are no longer overriding those two environment variables, so if you have any scripts that assume %TEMPPATH%==%TEMP%==%TMP%, you should review them.

Distributing new HYACCESS files

When you are provided with an upgraded set of HYACCESS files you may need to distribute a new HYACCESS.INI to a number of places in this potentially complex Hydstra world:

- To the top of your Hydstra tree, e.g. *h:\hydstra\prod\hyd\hyaccess.ini*
- To any service folders for SVRIMP and SVRRUN, e.g.
h:\hydstra\svr_services\svrrun and *h:\hydstra\svr_services\svrimp*
- To web servers (if any) in equivalent positions

Service jobs like SVRIMP and SVRRUN will not automatically see the new HYACCESS.INI until you stop the services and restart them.

If you still see a flurry of warning messages about expiring HYACCESS files, remove PTMPPATH\hyxmit.txt and the messages should go away. If they keep re-appearing then you haven't completed the job.

In a complex world it is best to set up Syncovery jobs to push the current HYACCESS file to all the other places it needs to go.

V14.00 - heads up: changes to how ARI axes are labelled on probability plots

We have modified our graphics engine to correctly display ARI values in probability plots. HYDIST, HYLP3 and HYIFD will have an INI file setting called **UseARI**. If you set this to Yes (and you are plotting a PARTIAL or MONTHLY series type), the bottom X-axis display of the probability plots will show ARI values. In all other case, the bottom X-axis will be probability expressed as an inverse, ie 1 in Y.

Actual ARI is calculated from the probability value by the following equation:

$$ARI = 1 / -\ln(1 - AEP)$$

Where AEP is the probability and $\ln()$ is the natural logarithm.

Currently, these programs display inverse probability, but incorrectly labelled is as ARI.

Experts As A Service

Would you like your Hydstra system updated to the latest patches every month or quarter? Want better data hygiene?

Experts as a Service is available starting from half a day per month to allocate KISTERS support staff to routinely check your Hydstra system, review hygiene procedures and processes, implement best practice, offer advice on fixes to data where needed, and implement system patching and version upgrades.

If you're interested in discussing this please [contact us](#).

Clearer user attribution for SVRRUN jobs

We're tweaking the way Hydstra user IDs are reported in a few places for jobs run by SVRRUN, especially for scheduled / recurring jobs such as created by SVRCRON.HSC. In HYDLOG logs, each job's sign-on will be a bit clearer.

For most jobs (ie, triggered jobs), the logged text will be something like:

Starting task UPDATE_STATS **submitted by** TGM **from** VM-HYDSTRA-PROD

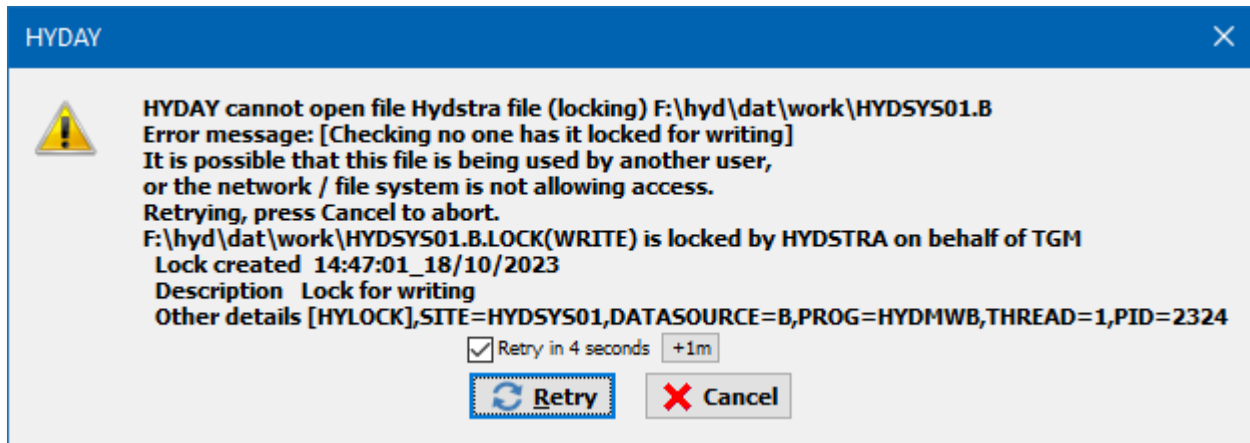
If the job is a recurring one, it will be:

Starting task UPDATE_STATS **originally submitted by** TGM **from** VM-HYDSTRA-PROD

Even more useful is that when a job locks a TS file, the TS lock file will report an enhanced version of the user ID, reporting the user ID used by SVRUN, as well as that of the person who triggered / submitted the task:

```
[HYLOCK]
OWNER=HYDSTRA on behalf of TGM
CREATED=14:47:01_18/10/2023
DESCRIPTION=Lock for writing
SITE=HYDSYS01
DATASOURCE=TELEM
PROG=HYFILER
THREAD=1
PID=2324
```

This means that if someone is denied access to a TS file, the interactive message they are shown should make it obvious that this is a server job, rather than the actual person running a program to lock the file:





Time Series Images

Recently we received positive feedback on the convenience and simplicity of using the TS Images system, and yet it still seems that only a handful of clients have set up the use of this feature since it was introduced two years ago.

The TS Images system has two main components:

- HYIMAGES, a program which allows for the viewing and importing of a time-series of images associated with a given site
- Integration with the workbench - image icons can be configured to appear in the same way as comments in the workbench, allowing for easy inspection in HYIMAGES using right-click





More information can be found under the HYIMAGES and 'TS Images system' help files. On the other hand, feel free to get in touch if you would like assistance with setting this up in the first place.

Tailor HYPERUPD to Your Needs

If you rely on the PERIOD table for having various time series properties at your fingertips, then HYPERUPD will be running on a regular basis. Most people have it run every night in 'full' mode (which tends to be a somewhat hefty process), with updates on trigger in the nimbler 'fast mode' if more current information is needed. But there is a trade-off - 'fast mode' only updates the period bounds, and leaves the maximum value, minimum value, number of points, etc. unchanged.

If in fact you do need all fields in the PERIOD table to be constantly up to date, then use the next fastest option - HYPERUPD has the command line switch '/UPDATEPTR=NO' which suppresses updating the PeriodTraces.txt file. This is a particularly effective alternative to 'fast mode' jobs which run against individual sites.

Checking Your HYDLOG Files for Errors and Warnings

Everything that Hydstra does is logged into HYDLOG files in dated folders under \hyd\adm\hydlog. In particular every program *start* and *end* is logged (with timing) along with any errors and warnings.

As a matter of course you should configure HYWOTSUP to run nightly just after midnight and save the report to the dated PTMPPATH\reports folder so you don't have to wait while HYWOTSUP runs. However HYWOTSUP only shows details of jobs with errors, and doesn't deal with warnings.

Another quick way of looking for warnings and errors is to use *gnuegrep*. Start a CMD box and go to today's HYDLOG folder:

```
hydlog
gnuegrep -i "\\|err\\|" *.txt>c:\temp\temp\errors.txt
gnuegrep -i "\\|wrn\\|" *.txt>c:\temp\temp\warnings.txt
```

When you see many similar errors or warnings you need to dig into the cause and rectify them if possible.

Ignoring Warnings about Data Discontinuities

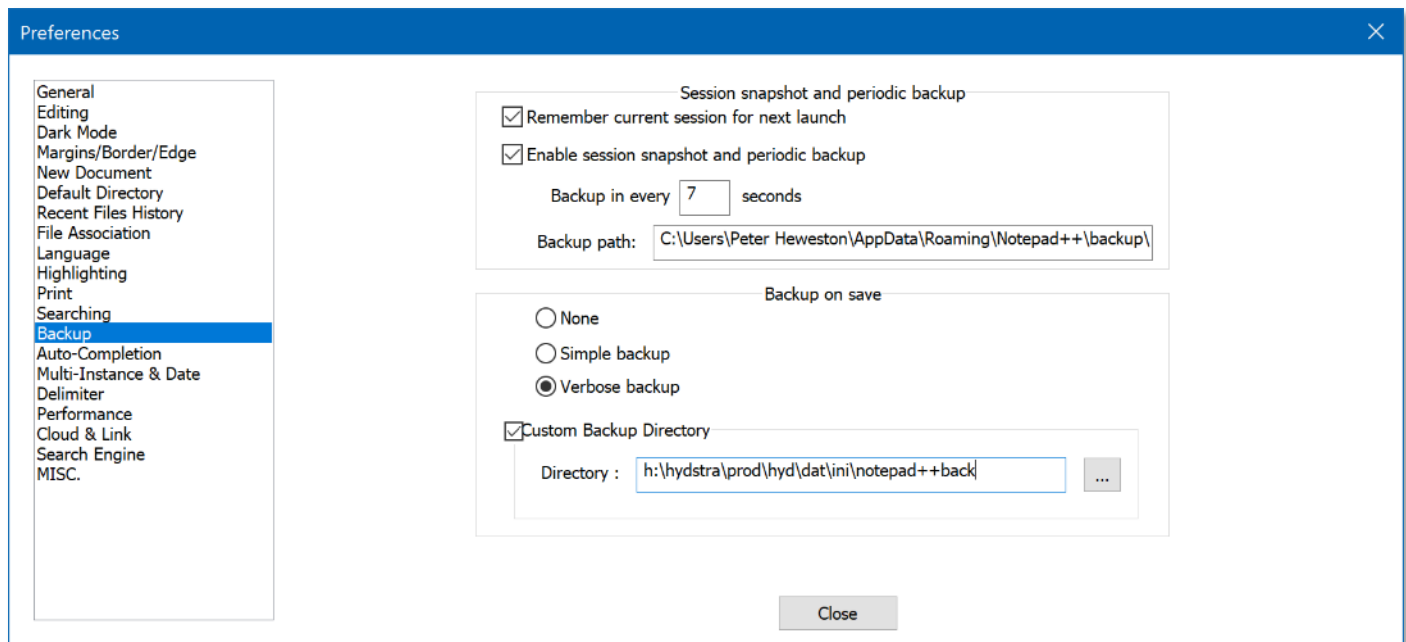
If your HYPXLORE and HYDLOG is littered with unwanted warnings about data discontinuities (value at end of one block differs from start of next block at the same time) and you don't care, you can turn the warnings off in HYCONFIG in the [Time Series] section with:

```
WARNDISCONTINUITY=NO
```

Organising Backups in Notepad++

If you use Notepad++ as your text editor for Hydstra scripts and batch jobs (and you should) you need to configure automatic verbose backups, so you can always go back and see what has changed. In Settings/Preferences/Backup you will find the backup configuration. Our recommendation is that everyone should configure backups to go to the

same Notepad++Back folder under INIPATH (be sure to make the folder first). That way no matter who changes a file, the previous version is saved for posterity:



Each time you save a file the previous version will be saved to the backup folder with a name of the form:

```
autojob.bat.2023-06-27_113718.bak
```

with the date and time the previous version was backed up as part of the file name. You can subsequently use Beyond Compare (another highly recommended tool) to find out what was changed and when.

It is important that everyone who edits Hydstra files configures their backups the same way, and into the same backup folder. The result is like a poor man's source code control, in that you have access to every past version of the file and when it was changed.

Importing Documents into DOCPATH

Hydstra provide a document store under DOCPATH that theoretically allows you to keep documents associated with the keys of any record in any table in Hydstra. In practice you probably don't want to allow the proliferation of documents without some thought, and DOCUMENTS.INI controls which tables can have associated documents.

The folder structure that Hydstra uses to store documents is quite complex, and with the possible exception of SITE, users should never attempt to guess the structure and manually import documents - they should always be copied in by a script or by the HYMANAGE documents tool.

If you are scripting imports the DLL call *table_doc_location* returns the folder name associated with a table record and its keys. All folder names for a record MUST be in upper case, and lower-case items in a folder name are a fail right from the start.

HYDBUTIL call DOCHECK will check that every document is correctly associated with a database table record, and report problems. HYGIENE test 24 simply calls HYDBUTIL DOCHECK.

We are developing a script HYDOCFIX that will fix a few upgrade issues that have slipped through our guards. For example at some time in the past we changed to TIME field in a HISTORY record to include seconds, so a folder name should have been upgraded from something like:

HYDSYS01-20100223-0-REMARK

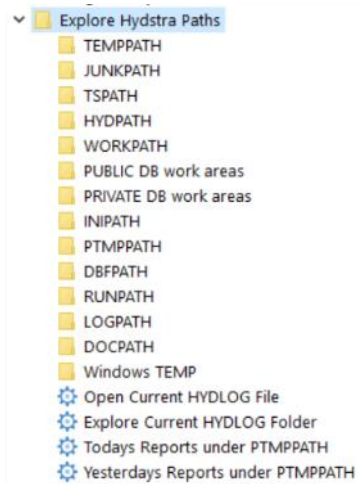
to

HYDSYS01-20100223-0.00-REMARK

since times with seconds are represented as having 2 decimal places. HYDOCFIX will find these problems and fix them (if possible). We will release it in a patch shortly.

HYXPLORE Handy Shortcuts

HYXPLORE has some handy shortcuts to key Hydstra folders under the Explore Hydstra Paths menu:



Use these to quickly open a file explorer in a specific folder, or look at your own HYDLOG file.

HYRATED – Global Mouse Mode

There is a new setting available in the Options / Program options dialog, Global Mouse Mode. If you enable the setting, HYRATED will attempt to remember your current mouse mode as you switch from view-window to view-window. The practical benefit is that if you are selecting gaugings, you can perform your mouse-drag in multiple windows in quick succession, without having to click on the window header, then choose “Select” from the toolbar, each time.

There are some limitations with this feature:

- If you select a mode that is not available in all view (such as “Move Rating Point”), then switch to a view that doesn’t support that option, it will fall back to the view you previously used in that view (or “Zoom” if it is the first time). This new mode becomes the global mouse mode, and when you switch back to the first view, it will remain selected.
- If you are using the Multiple Log Offsets feature (so that there are four views, including the log offsets view), when you click on the log offsets window, since that view does not support *any* mouse modes, the global mouse mode will reset to “Zoom”.

HYAUDIT Default Datasource

It's worth noting that in HYAUDIT.INI you can specify a datasource for many tests, and if you do, it will override whatever datasource you specify on the parameter screen. Since it is rarely the case that you want to do this, our general recommendation would be to not specify explicit datasources in HYAUDIT.INI and let the HYAUDIT parameter screen defined what datasource you want to run across. This means that the same test can be run across the whole archive , or the last two weeks of telemetry data, just from the HYAUDIT parameter screen, without changing HYAUDIT.INI in any way.

Without giving too much away, we can foreshadow that in V14 you can expect exciting developments in the area of HYAUDIT configuration and use.

V14.00 – changes to HYAUDIT Period Data Within Range test

We have changed the way the Period Data Within Range test operates. Previously it would extract intervals of data, one after the other. In other words, for a one-day interval, it would extract Jan 01, then Jan 02, Jan 03 etc.

This was a problem – what if an anomalous event was straddling two days? Each daily retrieval would only see “half” the event’s true size, and it might escape detection. To counter this, you might have to set the detection range STNINI entry artificially low – which then raises the prospect of false matches!

We have therefore reengineered the test to use a *moving window* – we scan the stored data point-by-point, but extracting one-day (in this example) intervals that *end* on each point in the trace. This change will have two effects:

- You won't miss any anomalous events due to splitting them between two intervals
- You might need to adjust your STNINI entries to a higher upper-range threshold, as the test will now find more exceptions than it used to.

Hydstra V14 - What do you want in it?

As the end of 2023 looms closer we are looking to bed down V14 ready for release. What features would you most like to see in the next release? More analyses and reports? Extra features in the Workbench? Anything about ratings and gaugings? More fields, spare or otherwise in tables? Additional tables? Now is the time to get your ideas in to us. Please log your suggestions and ideas in Jira [Service Desk](#) or email to hydstra.support@kisters.com.au.

Some General Observations on Password Policies

I know I'm shouting into the wind, but I'd like to remind you of a few observations on password policies, as published in Microsoft's document <https://learn.microsoft.com/en-us/microsoft-365/admin/misc/password-policy-recommendations?view=o365-worldwide> . In short:

- Maintain an 8-character minimum length requirement
- **Don't require character composition requirements. For example, *&(^%\$**
- **Don't require mandatory periodic password resets for user accounts**
- Ban common passwords, to keep the most vulnerable passwords out of your system
- Educate your users to not reuse their organization passwords for non-work related purposes
- Enforce registration for multi-factor authentication
- Enable risk based multi-factor authentication challenges

Furthermore, in case you spend a lot of time battling with blurry pictures of buses, traffic lights and other crappy CAPTHAS, see <https://au.pcmag.com/it-security/101213/bots-are-better-at-solving-captchas-than-humans-research-shows> . CAPTCHA stands for Completely Automated Public Turning Test. In summary, **bots can solve CAPTHAS faster and more reliably than humans**, so don't bother with them.

However many IT departments still refuse to acknowledge the futility of these techniques, thereby making our lives a misery. I personally have 440 entries in my password wallet. If each of these enforced a password change monthly (which thank heavens they don't), I would spend most of my waking life changing passwords.

WEBSERVICE.INI feature to control default values for optional parameters to webservice calls

You can override the default values for optional JSonCall functions using WEBSERVICE.INI. Here is a sample:

```
[Version]
Version=1

[get_ts_traces]
bad_value={missing}
```

The sections in the INI file must correspond to JSonCall functions, while the keywords are *optional* parameters (parameters that are not mandatory).

There are a number of restrictions to this feature:

- The values will only be used if a corresponding call is made, and that parameter is **omitted** from the call string. Any parameters provided by the caller will take priority.
- You can only specify one value per parameter per function, there is no way of having “if this then that otherwise something else” conditional logic.
- You cannot specify an empty string, but you can specify any number of spaces – just make sure to enclose them

in double-quote quotes, eg " ".

- For the get_ts_traces function, the results also apply to *format=csv* return data.
- This applies to webservice.exe only, not HYDLLP etc.

V14.00 - HYAUDITOR replaces HYAUDIT.SUMREP

In V14 we deliver a new CACHE table, and use it to cache HYAUDIT test results. New program HYAUDITOR replaces HYAUDIT.SUMREP with a more consistent looking output:

HYAUDITOR Summary Report								
Site	Data Present	Data Current	Data Within Range	Period Data Within Range	Data Within Calibration	Data Within Rating	Rate of Change Limit	Rainfall Intensity Limit
HYDSTRA_TESTING	n/a	acked	n/a	n/a	n/a	pass	n/a	fail
HYDSYS01	fail	fail	fail	n/a	fail	pass	pass	fail
HYDSYS01X	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a
HYDSYS02	fail	fail	n/a	n/a	fail	fail	n/a	fail
HYDSYS03	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a
HYDSYS04	n/a	fail	n/a	n/a	n/a	n/a	n/a	n/a
HYDSYS05	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a
HYDSYS06	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a
HYDSYS07	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a
HYDTEST	n/a	fail	n/a	n/a	n/a	pass	n/a	pass
% Pass	0	0	0	0	0	30	10	10
% Fail	20	40	10	0	20	10	0	30
% N/A	80	50	90	100	80	60	90	60
% Ack	0	10	0	0	0	0	0	0

Clicking on a test fail cell leads to a widget showing the cause of the failure, an offering various options for remediation, such as HYDMWB for editing the time-series, HYMANAGE for starting the SITE manage, where you may want to change a STNINI constant, HYAUDIT to re-run the test, and ACK to acknowledge that the task is failing for a known reason, such as equipment failure, and not to keep failing until further notice.

HYDSTRA_TESTING failed Bad Quality Percentage	Action	Action	Action	Action
* Bad quality data = 38.40%	HYDMWB	HYAUDIT	HYMANAGE	ACK

V14.00 - Style Upgrade for HTML Reports

We have embarked on a major style refresh on most Perl and Python-based HTML reports and post-processors. The resulting outputs will have a more uniform KISTERS style and layout. Here's an example of the new V14 HYLATEST output:

Archive Status by Site									
CSV File: hylatest.archive.csv									
<15000d		15000d-20000d			>20000d		No Data		Excluded
HYDSTRA_TESTING		HYDSYS01		HYDSYS02		HYDSYS04		HYDSYS06	
								HYDTEST	

Archive Status by Variable									
Grid displays the age of the data in days contained in the HYDSTRA A data file Site List=filter(tsfiles(a),match(hyd*)), Datasource=A, Run on 2023/07/18 17:05:33 CSV File: hylatest.archive.csv									
<15000d		15000d-20000d			>20000d		No Data		Excluded
Site	Site Name	QMIN	10.00 Rainfall (mm.)	100.00 Level (Metres)	11.00 Rainfall (in.)	141.00 Discharge (MI/Day)	150.00 Discharge Cu. M.	170.00 Sewer (Metres)	175.00 Velocity (M/s)
HYDSTRA_TESTING	Hydstra Test Data	?	12982d	13291d		10060d	13657d	14796d	
HYDSYS01	Hydstra Test Station	30.000	13167d	13656d			13657d	14796d	
HYDSYS02	Hydstra Test Station	50.000	8955d	15539d				9870d	9870d
HYDSYS03	Hydstra Test Station	50.000							
HYDSYS04	Test Geometrical	0.000		12252d					
HYDSYS05	Hydstra Test	0.000							
HYDSYS06		0.000			9149d				
HYDSYS07	Hydstra Test Station	0.000							
HYDTEST	Hydstra Test Data	0.000	15448d	Excluded					

Summary Report				
Variable	Variable Name	<15000d	15000d-20000d	>20000d
10.00	Rainfall (mm.)	75.0%	25.0%	0.0%
100.00	Level (Metres)	75.0%	25.0%	0.0%
11.00	Rainfall (in.)	100.0%	0.0%	0.0%
141.00	Discharge (MI/Day)	100.0%	0.0%	0.0%
150.00	Discharge (Cu. M.)	100.0%	0.0%	0.0%
170.00	Sewer (Metres)	100.0%	0.0%	0.0%
175.00	Velocity (M/s)	100.0%	0.0%	0.0%

All major post-processors on Delphi programs like HYDAY, HYMONTH etc have been similarly updated:

Kisters Pty Ltd.

Site HYDSYS01	Hydstra Test Station - Composite data												Year	1965
VarFrom 100.00	Stream Water Level in Metres												Table Type	Rate
VarTo 140.00	Stream Discharge in Cubic Metres/Second													
Figures are for period ending 24:00														
Day	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	Day	
1	0.134	0.020	0.000	0.000	0.042	0.062	0.092	0.078	0.055	0.043	0.432	0.092	1	
2	0.121	0.018	0.000	0.000	0.042	0.067	0.094	0.080	0.050	0.043	0.327	0.185	2	
3	0.123	0.017	0.000	0.000	0.040	0.073	0.101	0.083	0.053	0.038	0.266	1.032	3	
4	0.127	0.013	0.000	0.000	0.037	0.074	0.101	0.087	0.072	0.030	0.215	0.669	4	
5	0.121	0.009	0.000	0.000	0.036	0.074	0.096	0.092	0.153	0.029	0.159	0.389	5	
6	0.112	0.009	0.000	0.000	0.036	0.074	0.089	0.092	0.288	0.029	0.158	1.382	6	
7	0.108	0.008	0.000	0.000	0.036	0.075	0.086	0.113	0.328	0.029	0.156	0.727	7	
8	0.084	0.007	0.000	0.000	0.036	0.069	0.083	0.144	0.501	0.028	0.128	0.514	8	

Minor details like program version and run date re now delivered as a hover tip on the title line:

Kisters Pty Ltd.				
Site HYDSYS01	Hydstra Test Station - Composite data	HYDAY V146 12:03_03/08/2023	Year	1965
VarFrom 100.00	Stream Water Level in Metres		Table Type	Rate
VarTo 140.00	Stream Discharge in Cubic Metres/Second			
Figures are for period ending 24:00				

V14.00 – Changes to default Windows font for HYXPLORE, parameter screens, HYMANAGE etc

In version 14.00, we've adjusted the default HYPRMFNT.INI settings to use the current Windows recommended fonts for HYXPLORE, all parameter screen programs, HYMANAGE and MODSYN. The default font has been changed from Arial to Segoe UI, and the fixed-spacing font (used when we need things to line up) from Lucida Console to Consolas.

If you have a copy of this INI file in your INIPATH/TEMPPATH, the changes will not appear – you can either delete those files, or manually update them after looking at the one in MISCPATH.

Likewise, a couple of fonts (like the one used for text reports in HYXPLORE) are managed by “Settings” menu options in those programs – if you care, you'll have to change them from inside those programs.

V14.00 - HYMANAGE Delete Transaction Logs

In version 14.00, HYMANAGE will start writing recoverable transaction logs of all records deleted, to a folder under PTMPPATH.

This is configured by a new HYCONFIG.INI keyword: HYMANAGEDELETEDLOG, which is a number controlling how many *days* to keep these files.

- A value of -1 means no logs are written
- A value of 0 means they are kept forever
- Any other value means that transaction logs older than that many days will be deleted

Since deleting records in HYMANAGE is relatively rare, we do not anticipate the log files will grow uncontrollably, so the recommended value is zero – you may not discover something has been deleted until much later! A HYSRIPT job called HYMANAGE.RECOVER.HSC will restore deleted records for which logs still exist, to a work area.

V14 - User Constraints in PASSWD

In Version 14.00, the PASSWD table will contain a new field: CNSTRAINTS (Constraints). This is a comma-separated list of keywords, or keyword=value pairs. Currently supported keywords are:

A comma-separated list of constraint keywords, or keyword=value pairs. Currently, the following constraints are supported:

NODELETE

This user is not permitted to delete records for any table, using HYMANAGE.

Depending on their permissions and the table, they may be allowed to add or modify records.

Note that this constraint can only take away what would otherwise be granted

NOMULTIDEL

If allowed to modify a database table, this user may only delete single, individual records

Whenever they choose the delete option in HYMANAGE, they will only be allowed to delete a single record.

They will not be offered the choice of "this record / all records in view / etc", the choice will be fixed.

Likewise, it also means that if a record has "child records" underneath it, they will not be able to delete the parent record

Hydstra Administration Service

"Experts as a Service"

In recent times some Hydstra client agencies are finding that they no longer have either the skills or the manpower to maintain their Hydstra systems in good order (a review of a full HYGIENE output is a likely indicator of how well your Hydstra system is being maintained, with lots of green being good and lots of red being poor).

In response to client requests KISTERS have been evaluating a new advanced support contract which allocates KISTERS support staff to regularly check your Hydstra system to review procedures and processes, implement best practice where possible, and offer advice on fixes to data where needed.

This service is available starting from one day per month, and is dedicated to system monitoring and improvement.

Common inclusions:

- Regular log file analysis and reporting using HYWOTSUP, HYDLOGEX etc to identify problems.
- Running HYGIENE and reporting on issues raised. Where possible we will rectify the issue if caused by a system problem, or report to the client if a data problem is encountered.
- Automation of routine management tasks in AUTOJOB and/or SVRRUN (if licenced). The Hydstra Help topic 'Routine System Administrator Tasks' will be used to guide development.
- Configure and review spinning backups at various levels, subject to space availability.
- Manage INIPATH, removing obsolete files and controlling access.
- Review PASSWD to control who has access to Hydstra. Expire users who no longer require Hydstra access.
- Review file space usage, clean up old log and report files, zip up raw data folders as needed, report on stray folders, etc.
- If required, configure operating system file permissions to various folders under Hydstra
- If required, control stray files and folders in the Hydstra tree using HYDIRCHK.
- Review IT backup procedures, particularly as they relate to very long-term data retrieval from backups (years or decades later).
- If applicable, configure and run WEBTESTALL to confirm that web services and web sites are running and up to date.

If you're interested in discussing how a Service Agreement could be customised for your Hydstra system, please contact us through our [Service Desk](#).


Proposed DEFAULT User Expiry in V14

For a variety of security-related reasons we propose to expire in V14 the little-used DEFAULT user id in PASSWD. In V14 every Hydstra user will need to be explicitly registered in PASSWD, and hopefully most will auto-login via the WINUSERID field. The only reason a PASSWD entry should not have a WINUSERID entry is if a user needs to have two different Hydstra levels and switch between them - for example by default they log in with TS=2, but occasionally they need TS=3 then the TS=3 entry might need to have a manual login and a blank WINUSERID field, while the normal TS=2 entry will have WINUSERID set.

For the same reasons we also encourage you to expire generic users like USER0, USER1 etc. They are dangerous in that they don't track who exactly made what change, or indeed who was using them at any point in time.

Marking Missing Data in the Data Manager's Workbench

Typically, missing data is automatically detected and identified (with a bad quality point or a gap between blocks) on import into Hydstra. However, if you find data that has not been imported correctly, there are a number of ways to fix it:




1 - Use the Points – Insert... menu to insert a point, edit the point and change its quality code to a bad quality code. This will insert the bad quality point in the middle of the missing period, which is fine, but if you want to replicate the standard behaviour you would...

2 - Use the Blocks – Add New Point... menu to add a zero value, bad quality, point one second before the end of the missing period. However, for longer periods of missing data, you might like to create a gap between blocks, so you would...

3 - Use the Blocks – Split menu to split the block at the start of the missing period, split again at the end of the period, then delete the 2 point block that was created by the double split operation.

While these are all valid methods, they are somewhat labour intensive, so we have added a Blocks – Create Gap... menu to the Workbench. Simply select the two points either side of the missing period and choose your desired gap method:



Some new keyboard shortcuts in the Data Managers' Workbench

I suppose only boomers will care about keyboard shortcuts? I (Trevor M) seem to be the only one in the company who regularly uses them... but anyway, we've added a few more to the workbench (the Editor windows).

- Ctrl+Shift+A = clear selection

- Ctrl+Y = Redo last (undone) modification
- Ctrl+M = open Quality Map dialog (for currently selected points)

Please note that if you see a button or menu item where one of the characters in the text is underlined, that is a keyboard accelerator (which is technically not the same as a shortcut, but practically very similar). For an already-open menu, you can just press the corresponding key. For buttons, you can press Alt+{key}.

Hydstra V11 Unsupported

If you're on V11 you won't read this anyway, but if you have any friends on V11, ask them to upgrade ASAP. V11 is already unsupported, and unsupportable, as we don't have clean V11 source code after the crash, and with V14 release possible before the end of the year, V11 will become positively antique.

Who Deleted My Database Records?

It seems some agencies are having issues with Hydstra users accidentally deleting database records by mistake. We are looking into possible mechanisms for making accidental deletions more difficult, but in the meantime, you can hunt down and assassinate (or perhaps re-train) the offending users by searching back through HYDLOG. Whenever HYMANAGE is closed it writes a summary of record deletions to HYDLOG, containing a summary as follows:

```
HYMANAGE|2708|*|||0||HSY|*|||0|msg|N|DELETE SITE=2, HISTORY=8, PERIOD=79, STATION=1, STNINI=21,
SERIES=484, PEAKTIME=6, BENCH=2, RATEPER=18, TTABPTS=2, TTABHD=1, GAUGINGS=213, SECTIONS=493, SECTHD=5,
SECTSURV=1, STNVISIT=4
```

You can find them using HYDLOGEX or *gnuegrep* to search for 'hymanage.*delete ':

```
h:\ \hyd\adm\hydlog\2023\07>gnuegrep -ir :hymanage.*delete" *.*
```

The issue of getting back the deleted records is rather more complex, and relies on good backups. We are reviewing our options there as well.

HYLATEST enhancements

HYLATEST has been enhanced in a number of ways: the *varlist* INI file entry has become a full variable filter, so can accept expressions like *not(anyof(*.01, *.02, *.03, *.10, *.20, 184-999))* to exclude variables you aren't interested in. A new keyword *stname* allows you to choose *short_name* or *name* to label your sites. A new keyword *ignoregood* allows you to ignore report lines that don't have anything less than *good* class. See updated HYSTNS documentation in Hydstra Help for an elaboration of Variable Filters.

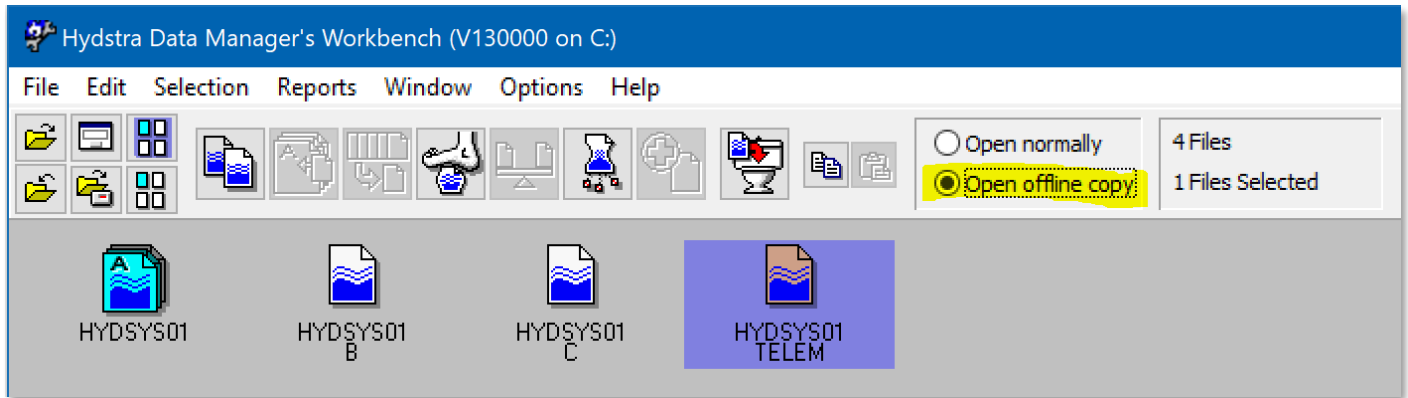
HYMANAGE Date Format

Dates in all Hydstra tables are stored with zero-padding for single digit days and months - however, you may still find that this format is not respected when viewing the tables in HYMANAGE. This is because in the manage, dates are presented according to the Windows 'Regional format' settings. See the Hydstra News article dated 2023/02/23 for how to update these settings.

Reminder to Open Files in Offline Mode When Possible

If you just want to look at live incoming data, but not edit it, there's a feature in the workbench to open a file in offline (read-only) mode, which opens a copy and doesn't inhibit write access. You can't edit files opened in offline

mode.



You should use offline mode whenever possible!

V14 - Storing the Analysis Date and Time of WQ Results

In version 14 the RESULTS table (as well as WREHOUSE, RESDEFLT, RESEXTR and RESENTRY) will have two new fields, RESADATE and RESATIME. This is for the analysis date and time for each specific result. You can of course leave them blank, as analysis date and time is already stored in the SAMPLES table. However some results are analysed independently of the sample as a whole, so now you can store information correctly.

When compiling samples and results into the WREHOUSE table, these new values are copied correctly, and when importing WQ data in batch mode using WQIMP, you can configure the rules to extract this information if it is present in the incoming data files.

Keeping SVRIMP Folders Under Control

We recommend you configure SVRIMP.INI to say *Daily Raw Folders = Yes*, particularly if you are processing a lot of files. Any Windows folder that ends up with 100,000 files or more slows down every process that accesses it. And don't forget to run HYZIPDIRS regularly to zip up old SVRIMP folders of processed data. There's really little need to keep more than the last month or so unzipped.

If your SVRIMP data is in turn coming from somewhere else, like a SCADA system or data lake, perhaps you don't need to keep yet another copy in Hydstra for very long. HYZIPDAT can be used to delete files and folders older than a specified age.

Modernising your Hydstra Fonts

In Windows 11, Microsoft have started using some new fonts to make things look more 'modern'. It's a standard play for Microsoft from time to time: change the fonts, round the windows corners, square the windows corners, make the headers transparent, make them opaque, change the colours, etc to provide the illusion of progress.

If you want Hydstra to look more 'modern' you too can play with fonts to a considerable extent. Take a copy of the following files from MISCPATH and put them in TEMPPATH:

HYPRMFNT.INI, HYVIEW.INI, WIP.INI

Then change all references to *Microsoft Sans Serif* to *Segoe UI*, and change *Lucida Console* to *Consolas*. At the same time you might like to bump up the font sizes a little, particularly if people have high-resolution screens.

If you're a Hydstra administrator, and you like the new look, move the files to INIPATH and everyone will get them.

It's worth noting that a number of fonts are built-in to compiled Delphi programs (e.g. HYDMWB), and there's no easy way of changing them unfortunately.

If you're setting this up as a system administrator, don't forget to manage HYINILOC.INI in INIPATH which controls INI file overwrites. Run program HYINIFILES.EXE to see what INI files are being overwritten.

AUTOJOB.BAT routine to submit tasks to SVRRUN

In some cases overnight AUTOJOB.BAT jobs take too long because AUTOJOB executes tasks sequentially in a single thread. We have added the capability to submit tasks to SVRRUN via new subroutine in AUTOJOB called :SVRJOB (instead of :RUNJOB). You need to incorporate the :SVRJOB subroutine from AUTOJOB.BAT in MISCPATH into your AUTOJOB.BAT, then instead of calling

```
call :runjob TESTRUN "... whatever"
```

you can call

```
call :svrjob TESTSUB "...whatever"
```

Of course you need to have an SVR licence and be running SVRRUN as a service.

Give some thought to how parallel jobs may interact, for example competing for time-series or database table resources. And you don't have any control over job ordering once tasks are submitted to SVRRUN.

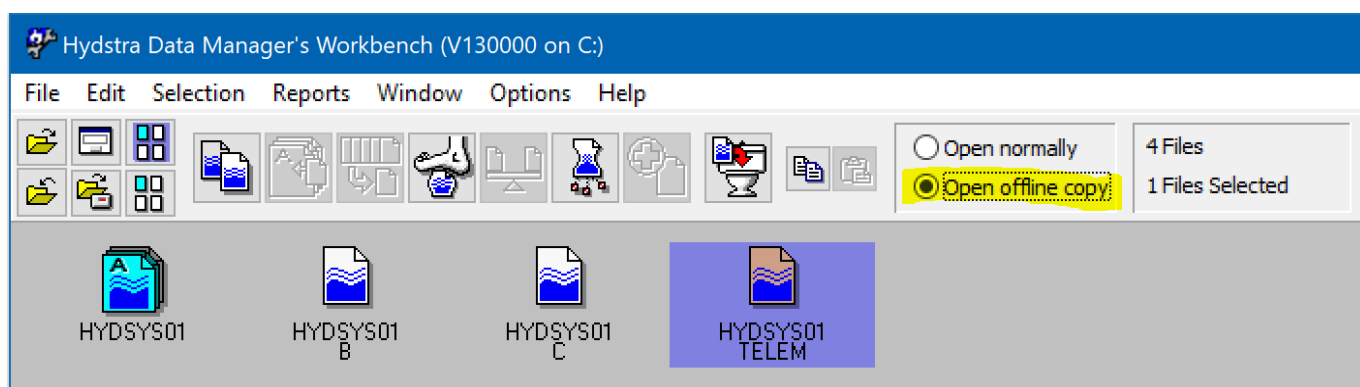
WEBSERVICE.PL withdrawn in V14

We propose to withdraw WEBSERVICE.PL in V14, as it no longer conforms to modern standards for interfacing with other languages, including Python and R. Furthermore WEBSERVICE.EXE, which more or less replaces it identically performs very much better. Please consider this in your future planning - you can switch to using WEBSERVICE.EXE fairly easily right now, if you haven't already. Contact Hydstra Support if you need assistance.

Avoid Locking Telemetry Files in HYDMWB

If you are running SVRIMP and data is coming in frequently you need to avoid locking up telemetry files by holding them open in the workbench or other analytical programs. A number of settings and strategies can help:

- Copy the incoming hot telemetry files to a second copy for general analysis and reporting using HYFILER MIRROR on a trigger. MIRROR can quietly skip locked files, and they will be updated next time around. However users need to be trained not to edit the copies as they are overwritten frequently and changes will be lost.
- You should limit in DATASRC.INI the number of people with workbench edit access to the incoming telemetry files to only those with a real need to edit them. Low-level TS=0 users shouldn't even see the files in the workbench. Keywords ULEVREAD and ULEVWRITE control the user levels that can see or edit a file in the workbench.
- If necessary set up a new role in the PASSWD *Levels* field like TELEDIT=1 and then further restrict edit ability to the TELEM datasource in DATASRC.INI to nominated users who have TELEDIT=1 in their PASSWD entries.
- If you just want to look at the live data, but not edit it, there's a feature in the workbench to open a file in offline (read-only) mode, which opens a copy and doesn't inhibit write access. You can't edit files opened in offline mode.



- In SVRIMP.INI you should set *Wait for locked files=no* to skip locked TS files and come back later.
- In SVRIMP.INI, the *Retry suspended files* (in minutes) is largely irrelevant – problems due to locked files are always retried, this option only controls the response to other, less-common problems.
- In SVRIMP.INI, if your processes are not to edit the raw, live telemetry data, you should set the *Preserve TS Edits=no*. This will further reduce the amount of time SVRIMP holds the telemetry file open, and thus opportunity

for access contention problems.

- In HYCONFIG.INI you should have a relatively short workbench timeout set to avoid someone opening a TS file and walking away leaving it locked, with WTIMEOUT=10.

Cleaning up IIS Log Files with HYZIPDAT

A relatively common cause of web servers failing is the disk becoming full due to IIS log files building up. Ideally IT should be looking after such issues, but if it falls on you to maintain the web server, Hydstra program HYZIPDAT can be configured to delete (or zip up) files older than a certain age. Set it up using the AUTOJOB framework and run it every night on the web server.

Backup up Hydstra using HYCLONE and ROBOCOPY

The Hydstra job HYCLONE wraps the Microsoft utility ROBOCOPY, making it easy to set up tasks like copying the whole of \hydstra\prod to \hydstra\back nightly (provided you have space). A nightly backup can be very handy for recovering from those oops moments without having to annoy IT.

In the past we sometimes set up ROBOCOPY.EXE in INIPATH, as it wasn't widely available, but modern Windows comes with a later version of ROBOCOPY built-in with additional interesting features.

To avail yourself of the more modern facilities you need to delete ROBOCOPY.* from INIPATH, and then edit HYCLONE.INI and change *\${inipath}robocopy* to just plain *robocopy*.

Then you can, for example, add /MT to the list of robocopy options. MT means Multi Threaded, with a default of 8 threads, which will make your copying go very much faster. You can even try /MT:16 to use even more threads, though sooner or later you will exhaust your IO capacity.

Run robocopy /? to see all the options available.

Chromium browser support in HYXPLORE coming in V14

We plan in V14 to optionally change the embedded browser in HYXPLORE from an IE based component to a Chromium based replacement (e.g. Edge). IE has been deprecated for some time now, and will likely be withdrawn completely soon.

HYXPLORE uses the embedded browser to display HTML outputs, amongst other things, and in particular HYXPLORE uses some tricks to allow it to capture clicked links in applications like HYAUDIT.SUMREP, where you can click on a HYAUDIT failure and start various Hydstra applications.

SVRRUN has a high-performance task ingestion method based on SVRSUB.PRM files

SVRRUN has a new feature – aimed primarily at KISTERS Pty Ltd developers – that can ingest tasks from a directory of files (a la SVRIMP). The file-based input queue must contain files of the same format as SVRSUB.PRM. The advantages/limitations are:

- The submitting process can simply write a file to the input queue folder instead of executing SVRSUB, resulting in large performance gains for high volumes
- SVRRUN reads these files, and itself appends them to the SVRTASKS table – this does involve double-handling, but it means there's a level playing field regarding priority and submission order.
- As this feature is primarily for internal use, there is slightly less hand-holding. In particular, it is highly recommended that processes write a file with a **non-job file extension**, then **rename** it when done, as SVRRUN does not perform any file-contention mitigation. The reason for this is to avoid SVRRUN seeing the file before you have finished writing (and closing) it – the rename step is far more atomic – the renamed file will not be seen until the rename step is complete, and therefore the file will be available for ingesting and deleting.

Dealing with rating table exceptions

As delivered, Hydstra raises an error under a number of rating table scenarios, including attempting to convert data before the first RATEPER record, below the lowest RATEPTS, above the highest RATEPTS, or outside the range of

RATEEQN. If you are running a lot of automation tasks with these default settings your HYDLOG files may fill with error messages.

We provide a number of HYCONFIG variables to control when happens where ratings are outranged in time or value, and we suggest you change the defaults to simply return bad data instead. The HYCONFIG keywords can be found in the [Time Series] section of HYCONFIG and we recommend settings of:

ABORTRAT=NO

RATEOVQ = 151

RATEUNQ = 151

ABRTNCTF = NO

ABRTRATPRE = NO

RATEPREQ = 151

See HYCONFIG.INI in MISCPATH for a description of these fields.

HYAUDIT can be configured to report on data outranging ratings using the *Data Within Ratings* test.

SVRIMP attempt to press on with other variables for the affected site if one of them has an invalid/missing data trans

Previously, if a missing or invalid datatrans was encountered while processing data for a site in SVRIMP, any variables after the affected one would be skipped, and everything written to the `\hyd\log\svrimp\{today}_ERRORS` folder. We have tweaked SVRIMP so that it will press on with any remaining variables that have come in. The file(s) copied to the error folders still contain all data for the affected site/run (including variables that processed OK), but now at least the good data will be processed to TS files.

HYCSV lets you configure a single header row, or no header row at all

You can configure the header row(s) written by HYCSV (using the HYCSV.INI file rather than new parameters). The default value for this new setting is to continue as usual - with three (or sometimes four) header rows. By changing the new *Header* keyword to SINGLE or NONE, you can condense that down to a single row, or none at all.

Staff News

The KISTERS team in Canberra is thanking David for his excellent work over the last 5 years supporting our WISKI/KiWQM and KiECO customers. David evolved over the last years into our expert in biological data and water quality. We will miss David for his expertise and surely as a very reliable and gentle person.



David is joining the NSW National Parks and Wildlife Service to assist in the coordination of invasive species (weeds) monitoring programs across the state's parks. He will continue working with environmental data in his new role to assist NPWS in getting the best biodiversity outcomes from the programs.

We think we gave David an excellent skill set for his coming challenge and luckily David and Lilly were prepared by Secret Santa with tools for life.

A big thank you to David from the KISTERS teams in Canberra and Hobart.

Information

This newsletter is published by KISTERS Pty Ltd and edited by Peter Heweston.

Canberra

Unit 4A, 24 Mahony Court, Weston ACT 2611
PO Box 3476 Weston Creek, ACT 2611, Australia
Phone: +61 2 6154 5200
Email: hydstra.support@kisters.com.au
Email: wiski.support@kisters.com.au
Website: www.kisters.com.au
Resources: resources.kisters.com.au
Service Desk: conflks.atlassian.net/servicedesk

Hobart

Level 8, 39 Murray St, Hobart Tas 7000
GPO Box 1390, Hobart Tas 7000
Email: hydstra.support@kisters.com.au

Sydney

PO Box 332, Liverpool BC, 1871, NSW, Australia
48-50 Scrivener St, Warwick Farm, NSW, 2170
Ph: +61 2 9601 2022
Fax: +61 2 9602 6971
Email: sales@kisters.com.au
Website: www.kisters.com.au

New Zealand

Core Facilities Building, Waikato Innovation Park
1 Melody Lane, Ruakura, Hamilton 3216
PO Box 9466, Waikato Mail Centre
Hamilton 3240 New Zealand
Phone: +64 7 857 0810
Fax: +64 7 857 0811
Mobile: +64 21 489 617
Email: sales@kisters.co.nz
Website: www.kisters.co.nz

Sacramento

KISTERS North America
2999 Douglas Blvd, Suite 180,
Roseville, CA 95661, USA
Phone: +1 916 723 1441
Fax: +1 916 723 1626
Website: www.kisters.net

Aachen

KISTERS AG
Pascalstrasse 8+10 52076 Aachen
Nordrhein-Westfalen, Germany
Phone: +49 (0)2408 9385 0
Fax: +49 2408 9385-555
Email: info@kisters.de
Website: www.kisters.eu/de

Africa

Gina Gaspar, AQUATRES
Mobile: +27 82 713 8491
Email: aquatres.sa@gmail.com

France

Kisters France SAS, 147 Avenue Paul Doumer
92500 Reuil-Malmaison, France
Phone: +33 1 41209-200
Fax: +33 1 41209-229
Email: info@kisters.fr

Italy

Temistocle Li Vigni
Certified Practicing Hydrographer
GEOSPHERA Hi-Tech Supplies Srl
Via Panoramica 85 – 80056 Ercolano (NA) – Italy
Phone: +39 081 7779059
Email: geosphera@geosphera.com
Web: www.geosphera.com

Spain

KISTERS Ibérica s.l.
Calle de Hernando de Acuna No. 34, Planta 1a,
Oficinas 3 y 4, 47014 Valladolid, Spain
Phone: +34 983-330744
Fax: +34 983-341502
Email: iberica@kisters.es

Shanghai

KISTERS Shanghai Software Development Co. Ltd.
6D, No. 438, Pudian Road
Pudong New Area, 200122 Shanghai, China
Phone: +86 21 68670119-801
Email: info@kisters.cn