



KISTERS Australia News

May 2022

Contents

Contents	1
From The GM's Desk	2
KISTERS Recovery Update	2
WISKI Product News	3
The new WISKI Release 7.4.13 is available	3
Hydstra Product News	4
Hydstra V12 and V13 Systems and Patches Available	4
Hydstra Version 11 Unpatchable and Unsupported	4
Introduction to Hydstra News	4
FLAG field in HYDMEAS becomes coded in V14	4
Finding and dealing with SVRIMP problems	5
Support for more flexible times and dates in HYGENLOG	5
HYADJUST additional commands	6
Basic Workbench Training Videos	6
More on LOG4J vulnerability	6
Best Practice for SVRIMP and SVRRUN	6
Interfacing Hydstra with APIs	6
The Dreaded C0000006 Error	7
New Perl Datasource Expression	7
Spring Vulnerability	7
Thoughts on DATASOURCES	7
HYCHANNELL Can Save to Work Areas	8
HYGAUGE Supports Hach FH950 Gaugings	8
Information	9

From The GM's Desk

By Paul Sheahan, General Manager, KISTERS Pty Ltd

Hi,

Save the date !

We are pleased to announce that KISTERS will be hosting face to face user group meetings in both Australia and South Africa this year.

The Australian user group meeting will be held in Canberra on **Thursday July 28th**. This meeting will be a full day including desktop and web product updates followed by a conference networking dinner. On the Friday we will host an open office at our Weston Creek office where you're invited to book a time to talk with specific staff or just roll up to talk things hydrology, data management and systems.

The South African user group meeting will be held in Pretoria on **Wednesday 14th September**. This meeting will be a full day including desktop and web product updates.

Further details on locations and costs will be circulated when they are finalised. Please contact Paul Sheahan if you are interested in presenting a paper at the conference.

We look forward to being able bring you all together again and engage in our shared passion.

Cheers,

Paul

Paul Sheahan
General Manager
KISTERS Pty Ltd



KISTERS Recovery Update

By Peter Heweston

The ransomware attack of last November hit KISTERS hard, but the rebuilding process is largely complete.

- In Australia all staff laptops have been re-imaged and rebuilt.
- All files on all computers have been triple-scanned with three different antivirus scanners to ensure they are free from corruption.
- The email system is up and running, and all staff can be contacted on their usual Kisters email addresses. Two support email addresses have been created for Australian clients – hydstra.support@kisters.com.au and wiski.support@kisters.com.au. The old support@kisters.com.au is redirected to hydstra.support@kisters.com.au.
- Hydstra systems and patches are rebuilt and uploaded regularly with fully signed executables.
- A new WISKI release is announced below.
- Lotus Notes is up and running for managing Hydstra System Reports. It will be replaced by Jira later in the year.
- Confluence is up and running for in-house procedural documentation.
- JIRA is up and running for selected clients, with a wider roll-out expected later in the year.
- The VOIP phone system in Australia is functional, and staff in Canberra and Hobart can be contacted on internal phone numbers if you know their number. The main support number is 02 6154 5200.

WISKI Product News

The new WISKI Release 7.4.13 is available

We are pleased to inform you that we have successfully released the announced 7.4.13 WISKI version, as planned on the 12th of May 2022.

As you may know, we have worked due to the security incident in November 2021 on the re-creation of the entire IT infrastructure, including the development environment systems, the remote maintenance access, the build, development and test systems in accordance with stricter security guidelines. We see these measures during the last months as key prerequisites for a sustainable and safe basis for all further work in the future.

With this newsletter, we would like to inform you about the current state of development, explain the release strategy for the next weeks and the procedure to plan installations on your systems.

The new WISKI Service Release 7.4.13 SR 7

The Service Release SR7 of WISKI comprises of the last status of the continuous development during the last years. Our development teams have worked on analysing and solving multiple security risks that are critical and relevant to our software. This includes extensive fixes and updates to all affected Log4j artefacts.

In addition to server-side enhancements, we have modernised the user interface of the WISKI Desktop Client and improved the stability, usability and accessibility. We look forward presenting further enhancements during the international user conference in June or as seminars during the water academy.

What are we working on at the moment?

Since the WISKI 7.4.13 SR7 is the first release after the reconstruction phase, we have decided to adapt our release strategy and time plan. Since last week, we have installed this version on various internal systems to undergo further practical testing, as far as possible also under real conditions. Additional to our internal systems, the version was delivered to "First customers". Currently we want to use this opportunity to collect and immediately react on initial feedbacks and experiences with multiple size and complexity of customer environments. Therefore, we will provide "interim" releases instead of patches and service packs in the next few weeks as follows:

Version and SR (Service Release)	Date
7.4.13 SR7.1	25.05.2022
7.4.13 SR7.2	08.06.2022

After this phase, we will get back as usual to our standard release plan with patches, service packs and official releases.

How can I book and organise an installation?

Please get in touch with your contact person at KISTERS to arrange an appointment with our support team. We highly recommend to follow the "short-term" release schedule described above as an orientation for the time plan.

On behalf of the entire KISTERS water team, we would like to thank you very much for your patience and understanding over the last difficult months.

Hydstra Product News

Hydstra V12 and V13 Systems and Patches Available

The Hydstra build system has been restored, and images for V12 and V13 are available on the KISTERS web site at <https://www.kisters.com.au/index.php?filename=downloads> . Please contact hydstra.support@kisters.com.au for the download login details, and please provide your HYACCESS.INI file with the request.

Hydstra Version 11 Unpatchable and Unsupported

Hydstra V11 was already officially unsupported, as only the last two version are supported (V12 and V13). However due to technical reasons it is impossible for us to do anything in Delphi for V11. Please consider upgrading to V13 (a two-step upgrade) as a matter of urgency. Contact us if you would like us do it for you as a consulting project.

Introduction to Hydstra News

Hydstra News is a new information channel for Hydstra users that automatically displays Hydstra News in HYXPLORE whenever you start it, provided you are running a recently patched V13 or V12 Hydstra system.

Any TS=3 user with access to the internet will automatically attempt to download the latest news from the KISTERS web site in the background every time they start HYXPLORE. Every Hydstra user regardless of user level will be shown the news every time they start HYXPLORE.

If nobody in your organisation has internet access then your administrator can manually download the news at home and bring it in to work. The details of how the news system works are documented in the Help topic *Hydstra News Distribution*, and you can suppress it if you wish.

Here below are an assortment of recent news, tips and techniques from Hydstra News, and we hope to bring you fairly frequent updates via this mechanism. Please let us know what you think.

Where we announce a new program or capability in Hydstra News, it is implied that it will become available in the next V13 patch after the announcement. Patches are typically created every Thursday evening and uploaded every Friday. Hence if you're reading about it, it's already in the current patch up on the KISTERS web site. Of course, that doesn't mean your organisation has downloaded and applied the patch! V12 users – sorry, you need to upgrade to receive the new facilities described in Hydstra News.

FLAG field in HYDMEAS becomes coded in V14

In Hydstra/GW the HYDMEAS table is used for manual readings of bore water level, and may also be used to record bore water quality readings. Consequently, it has the same FLAG field as the Hydstra/WQ RESULTS table, allowing you to flag values with flags like <, >, TN (Too numerous to count), ND (not detected) etc.

Through some oversight of ours it turns out that the groundwater HYDMEAS.FLAG is not a coded field, when it should be using the same codes as water quality RESULTES and WREHOUSE, namely codegroup RES.

In V14 we will bring HYDMEAS.FLAG into alignment by making it a coded field against code RES, but this may cause HYGIENE and validation errors unless you are prepared for it.

In order to explore current values, you run the GW HYMANAGE - unlink, and run Tools/Statistics on the FLAG field.

Alternatively, you can run HYDBSQL with the following command:

```
--timeout=0
select FLAG,count(*) from HYDMEAS
group by FLAG
order by FLAG
```

Do the same for WREHOUSE, and then align the HYDMEAS flags and the RES codegroup to make everything valid.

The originally distributed RES codegroup include codes LT and GT, though you probably prefer the more symbolic < and > and you may wish to change your data and the RES codegroup. You can change your data in HYMANAGE by unlinking, filtering, and then using the Replace function.

Once you have everything aligned you can test it out in V12 or V13 by putting in a USERDICT entry for HYDMEAS.FLAG making it a coded field looking up code RES and then running HYDBUI CHECK in full mode. IF everything passes then you are ready for V14 when it comes!

In V14 you can remove the redundant USERDICT entry, and HYGIENE will remind you if you forget to do it.

Finding and dealing with SVRIMP problems

SVRIMP is designed to carry on regardless of incoming data errors, and any errors are logged in HYDLOG and also reported in a dated error folder under the SVRIMP folder, e.g. :

```
H:\hydstra\prod\hyd\log\svrimp\202205
H:\hydstra\prod\hyd\log\svrimp\202205_ERRORS
H:\hydstra\prod\hyd\log\svrimp\holding
H:\hydstra\prod\hyd\log\svrimp\import
H:\hydstra\prod\hyd\log\svrimp\marshalled
```

In the 202205_ERRORS folder are two sorts of files – the CSV files that couldn’t be processed, and a series of files with names like ERROR_20220501013523.txt. Each of these files has the error message associated with a failed import. The sort of messages you might find include:

```
SVRIMP could not get locking access to D:\APPS\HYDSTRA\PROD\hyd\dat\ts\TELEMRAW\HYDSYS01.T
Context: Checking no one has it locked for reading
```

and

```
Serious error occurred while processing marshalled data for HYDSYS01
Error message:
Variable 53 not found in VARIABLE table in GetVarDetails (site HYDSYS01)
```

You really need to keep an eye on the error folders and deal with the issues promptly.

The first class of error usually indicates people are holding the T file open in the Workbench, or perhaps in long-running jobs like Excel models. One suggested solution is to immediately copy the incoming live telemetry file to a working file for read-only uses, but if data really needs to be edited in the workbench, get in quickly, do the job, and get out before more data arrives.

The second class of error is sometimes a misconfigured logger or import script.

Support for more flexible times and dates in HYGENLOG

The HYGENLOG logger processing system has been enhanced to process times and dates with no leading zeroes under certain limited conditions. Spreadsheets and loggers sometimes put out the first of January as 1/1/2022, making it difficult to parse with a template without a Perl or Python pre-processor to make everything line up. This new capability is *only* available for comma-separated data formats where the time and date are both in the first column, and is configured by adding time- and date-delimiters to the LOGMAST table, specifically the HMSTEMPLT and YMDTEMPLT fields.

For example, if a logger delivers data like this (note that the time and date is in the first comma-separated column):

```
1:15_1/1/2021, 123.456
10:1_1/10/2021, 234.567
10:1_10/1/2022, 345.678
```

You can process it with the following entries in LOGMAST (the [square brackets] are for clarity, and are not part of the entered values):

```
HMSVALID [9]
HMSTEMPLT [HH:II_]
YMDVALID [9]
YMDTEMPLT [ DD/MM/YYYY]
```

Please note the following qualifications:

- The feature is only available for delimited formats (i.e., comma-separated) where the time and date are both present in the first column.
- You have to include all the time- and date-separator characters in the “xxx Template” fields, including the one that joins time and date together (unless it is a space).
- You must specify the “ideal” form of time and date in the “xxx Template” fields, that is: describing a format where all elements are two digits (except the year, which is four). Hydstra will take this “ideal” template, and apply it

flexibly to incoming values that don't have leading zeroes.

- There is *no change* to how the “xxx Validation Line” fields work – you have to enter characters that accurately detect which lines contain date-and-time values, which is tricky if the formatting pushes date and time elements back and forth. In extreme cases, a single “9” in column one may be required.

HYADJUST additional commands

HYADJUST has new commands *precopy*, *preexec*, *postcopy* and *postexec* which allow your HYADJUST commands manipulate data before and afterwards, thereby enhancing your ability to do things with HYADJUST. We have also added a comment facility to comment out parts of your HYADJUST script, useful while testing. Check the HYADJUST doco for more details.

Basic Workbench Training Videos

Some time ago we ran a Workbench course for a client agency, and they have kindly allowed us to publish the recording of that course. The course was largely centred on using HYDMWB, though it did range over other basic topics as well. You can watch the videos at <https://www.kisters.com.au/trainingvideos.html>. You will need to contact hydstra.support@kisters.com.au for a user id and password to access the videos. There are over five hours of material in the workbench videos. Other available recordings include:

- Advanced Hydstra
- Advancing You Hydstra Skills
- Hydstra Administration
- MODSYN Modelling
- Hydstra ODBC

More on LOG4J vulnerability

As we have discussed previously, you can mitigate the recently described LOG4J security vulnerability by deleting the `\hyd\sys\run\LogViewer` folder, and if you are not licenced for Hydstra/ODB, the `\hyd\sys\run\ScriptServer` folder as well.

If you are licenced to use Hydstra/ODB then a recent V13, V12 or V11 patch will deliver a patch to log4j which removes the *JMSAppender* class, the cause of the vulnerability. When KISTERS Aachen delivers a patched version of *ScriptServer* which uses a later version of *log4j* we will install it and deliver it in a patch.

Best Practice for SVRIMP and SVRRUN

Our recommended ways of operating SVRIMP and SVRRUN have evolved over the years, and if you are updating, migrating or moving infrastructure it might be timely to do some work to update to best practice:

- SVRIMP and SVRRUN should be installed as services. The old method of a restarting batch job simply doesn't work reliably in the modern world, particularly where you want the processes to restart automatically after an outage.
- SVRIMP should be fed with SVRIMP #V2 CSV files. V2 files allow you to specify the datasource and direct different feeds to different datasources. For example, you might feed your mainstream telemetry to a TELEM datasource and your experimental IoT input to an IOT datasource. You may even choose to duplicate your incoming stream and send one half to a RAW datasource which is not be edited, but simply preserved for future reference.

Please don't use ZRXP format for any new work. We will continue to support ZRXP for the foreseeable future, but its use in SVRIMP is deprecated!

Interfacing Hydstra with APIs

An API (Application Programming Interface) is a way of interacting with an external system or source of data. Typically, an API is accessed by sending requests to specific URLs (endpoints) and then parsing the information returned to you from that request. Most modern APIs use JSON as the data transfer format – sometimes to pass

information in the request and almost always to return information in the response. Some older APIs use XML instead.

Many sources of data that you might want to import into Hydstra offer an API. Data logger manufacturers often build hardware that can send data directly to a cloud storage service, then you use the API to download the data into Hydstra. Examples of this include Campbell Scientific's Campbell Cloud (<https://docs.campbellcloud.io/api/>) and In-Situ's HydroVu (<https://www.hydrovu.com/public-api/docs/index.html>).

State or national archives of publicly accessible time series data often include an API for accessing the data. Examples of this include the USGS Water Services (<https://waterservices.usgs.gov/>) and the CDEC web services (<https://cdec.water.ca.gov/dynamicapp/wsSensorData>).

There are also a number of options for developing field data entry forms, that run on a mobile device and send their data to a cloud storage service. These are similar to the ones offered by data logger manufacturers, but are typically used to collect discrete data or site visit information, rather than continuous time series data. The data is then accessed via the corresponding API. Examples include GoCanvas (<https://www.gocanvas.com/>), FastField (<https://www.fastfieldforms.com/>) and ESRI's Survey123 (<https://survey123.arcgis.com/>).

We have significant experience with all the above-mentioned APIs, plus many more. Please contact Hydstra support if you have a source of data that includes an API that you would like to interface with Hydstra.

The Dreaded C0000006 Error

If you ever see a C0000006 error it is the results of network problems, and sadly there's nothing we can do about it except suggest you improve your network or reconfigure so that programs are closer to data. If you run Hydstra across a network to your workstation you may find a Citrix or RDP solution is more robust – but if your Citrix session has a C0000006 error then you really do have problems. If you Google for C0000006 you'll find plenty of discussion, and a few suggestions. Antivirus can sometimes be a problem, as can opportunistic locking. Bottom line is, on a good network, the problem never happens.

New Perl Datasource Expression

A few programs take a list of datasources as a parameter (e.g. HYDATSUM). If you pass in ALL datasources with something like DSOURCES(TSCARCH,TSCWORK,TSCAUX) you risk instantiating Perl and Python datasources when you don't want to. We have developed a Perl HYSTNS expression that returns a list of datasources excluding Perl and Python ones, so you can use them safely in HYDATSUM for example. The expression is PERLLIST(HYSTNS_ALLDS.PL).

Spring Vulnerability

The open source Spring Java framework is used in web development, and a serious vulnerability has been recently discovered – see <https://tanzu.vmware.com/security/cve-2022-22965> .

No Hydstra component uses Spring.

Thoughts on DATASOURCES

Keep datasources under TSPATH

When generating new time-series datasources to keep different sources of time-series apart we recommend you keep them in separate paths under TSPATH (\hyd\dat\ts), though keeping them under WORKPATH is fine also. For example, we recommend you keep incoming telemetry data under TSPATH\telem. If you're doing IoT experiments you may like to create an IOT datasource just for that data.

No dashes in datasources

Datasource names should follow the same rules as site names (alphanumerics and underscores). However, in the past we didn't fully enforce that rule, and in particular we let datasource names through with a dash in them, which subsequently caused problems. In more recent V13 patches we enforce that rule, so your system will break after a patch until you rectify invalid datasource names with dashes in them.

Keep raw data in a separate datasource

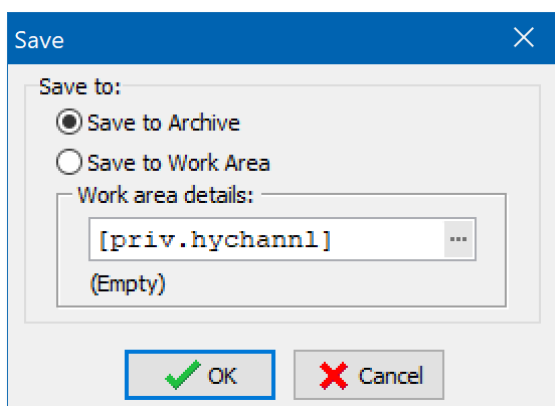
Some organisations like to keep a copy of raw data separate from edited data, and to do that they duplicate the feed into SVRIMP, or document manual or batch processes to copy data from raw to edited. Historically some agencies use a different subvariable to denote raw data, but these days we suggest you consider using a separate datasource instead – for example RAW and TELEM. The advantage of this scheme is that the files are smaller and less complicated when users look at them. Furthermore, you can use the access controls in DATSRC.INI to ensure that RAW files are only editable by TS=3 users, thereby preserving their integrity better.

Minimise holding TELEM files open

Another issue that sometimes causes problems is that users keep the TELEM files open for long periods, preventing SVRIMP from updating them. You can use datasources and SVRRUN triggers to ease that problem by importing data to TELEMIMP and then immediately copying to TELEM using a SVRRUN trigger and HYFILER MIRROR. Most reporting and viewing such as HYTELVIEW can be done on TELEM, and only users with editing permission are allowed to edit TELEMIMP, with strict training that they are to be quick and not linger.

HYCHANNL Can Save to Work Areas

HYCHANNL has been updated to allow save rating tables to a work area, as well as directly to the archive (if you have sufficient permissions). The feature will be available in the next V13patch.



HYGAUGE Supports Hach FH950 Gaugings

In the next V13 patch HYGAUGE will support gaugings made by the Hach FH950 Portable Velocity Meter. You will need to add the following line to HYGAUGE.INI in INIPATH:

```
Hach FH950 File = *.tsv, hygauge.hach.py
```



Information

This newsletter is published by KISTERS Pty Ltd and edited by Peter Heweston. It is distributed using MailChimp (www.mailchimp.com)

Homepage: <http://www.kisters.com.au>

All personal KISTERS Pty Ltd email addresses in Australia are of the form *firstname.lastname@kisters.com.au*, but all general support and accounting emails should be addressed to support@kisters.com.au .

Canberra

Unit 4A, 24 Mahony Court
Weston ACT 2611
Email: support@kisters.com.au

PO Box 3476
Weston Creek ACT 2611, Australia

Phone: +61 2 6154 5200
Fax: +61 2 6288 9061
Email: support@kisters.com.au

Hobart

Level 8, 39 Murray St
Hobart Tas 7000

GPO Box 1390
Hobart, Tas. 7001
Email: support@kisters.com.au

Hyquest Solutions Australia Pty Ltd

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@hyquestolutions.com.au

Hyquest Solutions New Zealand Ltd

PO Box 15 169
Dinsdale
Hamilton 3243 New Zealand

Core Facilities Building
Waikato Innovation Park
Ruakura Lane
Hamilton 3214 New Zealand

Phone: +64 7 857 0812 (DDI)
Fax: +64 7 857 0811
Mobile: +64 21 489 617

Email: sales@hyquestolutions.co.nz

Sacramento

KISTERS North America
1520 Eureka Road, Suite 102
Roseville, CA 95661 USA

Phone: +1 916 723 1441

Fax: +1 916 723 1626

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

Africa

Gina Gaspar

AQUATRES

Email: aquatres.sa@gmail.com

Mobile: +27 82 713 8491

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

E-mail: 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