KISTERS Australia News

April 2021

Contents

Contents 1
From The GM's Desk 2
Ransomware Threat Hots Up 2
Interfacing Survey 123 with KISTERS Products 3
We Speak MQTT (Somewhat) 3
Hydstra Product News 4
- Hydstra V13 In Production 4
- 64 Bit Hydstra? 4
- Hydstra as a Service? 4
- Time Series of Images 4
- Create a Rating from a Section in HYRATED 5
- PSC and HSC Jobs on the WEB 5
- Detailed Technical Advice on Running Hydstra over a Wide Area Network (WAN) 7
- HYTSEXPORT Exports Hydstra Time-Series to your Warehouse, Data Lake or Swamp 7
- Tidy up your old report, HYDLOG and SVRIMP Folders 9
- Antivirus in Data Folders 9
- HYGIENE Test 33 Looks for SVRIMP Errors 9
- Hydstra Creating Time-Series Data for Non-Registered Sites 9
- Copying HTML Hydstra Outputs to Excel 10
- Try HYMAILER_PY.HSC with Secure SMTP Servers 10
- Partial Aggregations with Missing Data 10
- HYGAUGE.INI and RATINGS.INI - Allow in TEMPPATH? 10
- New HYAUDIT test: Data Within Calibration 11
- HYTSXML Renamed to HY2WIXML 11
- Hydstra Licence Expiry Warnings 11

WISKI Product News 11
- Release Management and Client Base 11
- Monitoring WISKI Server Performance: Application Metrics 11
- Monitoring Bulk I/O Operations: Process Analytics 13
- Client-side Performance 15
- Alarming in KiWQM 16

News from KISTERS Overseas 18
Staying Connected Globally 18
Staff News 19
Information 20

From The GM's Desk

By Paul Sheahan, General Manager, KISTERS Pty Ltd

Australia is never boring! The last months have seen flooding, cyclones and exceptionally high Easter temperatures. I must admit to enjoying the higher Easter temperatures though.

KISTERS are moving towards the use of JIRA software for project and agile development management. We will also move to implement 'JIRA service desk' for system support over 2021 for both Hydstra and WISKI clients. This will give clients the ability to interact directly through the support tool and have a role in accepting closure of support reports.

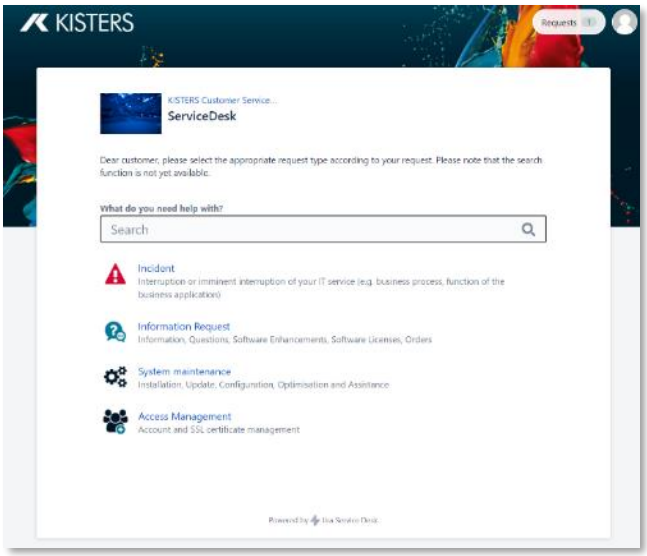

Our Hydstra developments continue with the implementation of theoretical rating development in HYRATED. This work enables ratings to be developed using section, slope and roughness coefficients. I'm looking to visit a number of regional Hydrographic teams over the next year and would like to be able to demonstrate this tool, please give me a call if you're interested in a presentation.

Moving forward with Hydstra development tools puts us firmly on the path for moving towards the next release possibly being a 64 bit application. This move is heavily favoured by IT departments but will have minimal effect on hydrometric data management work. We would really like to hear strong views for or against such a move, especially if you think that the move would be problematic in your organisation.

On behalf of the KISTERS Australia team, thanks for the opportunity continue working with you. We strive to be client focused and value your feedback.

Ransomware Threat Hots Up

In recent times an increasing threat has emerged in the computing world - ransomware. The typical ransomware attack generally follows the following pattern:

- Your IT system is penetrated via a phishing email, a weak password, or even so-called drive-by attacks where just visiting a web site from an old browser or unpatched operating system can infect you.
- Once the bad guys are in they may lurk for quite a while, downloading interesting files from your system, including personal and financial information, password files, contracts, etc.

- Eventually they spring the trap, and start encrypting your files. At this stage they will inform you of the attack, and provide a Bitcoin address where you can pay the ransom. Often they will publish some of your sensitive information on a public web site just to confirm that they really do have access to your data. They threaten to delete files if you don't pay up.

The general advice is never to pay a ransom, as there is no guarantee that they will actually provide the key to decrypt your data, or even if they do, that everything can be put back to normal. Good backups are the key to recovery.

Important components of ransomware protection include:

- Keep all operating systems and other software up to date and patched regularly
- Ensure you use the 3-2-1 backup strategy (3 backups on 2 different media with 1 offsite)
- Keep local backup drives offline except when in use
- Keep your antivirus software up to date
- Train staff to detect phishing emails and messages

A KISTERS client, the Scottish Environmental Protection Agency (SEPA), was subject to a serious ransomware cyber-attach just before Christmas last year, and the effects are still ongoing. Don't think that environmental agencies will be spared! You can read more about the SEPA incident at https://www.sepa.org.uk/about-us/cyber-attack/ and https://media.sepa.org.uk/media-releases/2021/sepa-cyber-attack-data-theft-service-delivery-and-recovery-update.aspx .

In 2020 there were many ransomware attacks on major corporations and agencies - nobody is immune. See https://www.pentasecurity.com/blog/8-most-catastrophic-ransomware-attacks-in-2020/ and https://www.itgovernance.co.uk/blog/list-of-data-breaches-and-cyber-attacks-in-february-2021-2-3-billion-records-breached for some disturbing reports of data breaches and ransomware attacks.

At the other end of the spectrum, we heard many years ago of a small environmental system running on a Pacific island which shall remain nameless where the sole server containing the system was stolen ... and of course there were no backups!

For small to medium sized Hydstra systems we still recommend writing your main time-series archive and database files to a DVD-ROM or USB stick from time to time, and keeping the media forever - don't recycle it.

Interfacing Survey 123 with KISTERS Products

If you are an ArcGIS Online subscriber (ESRI's software-as-a-service) or an ArcGIS Portal user (ESRI software installed on your own infrastructure), you may already have access to Survey123 (https://survey123.arcgis.com/).

Survey123 allows you to design your own data collection forms, deploy those forms to mobile devices, such as phones and tablets, and upload the information that is collected to your ArcGIS database. It is functionally similar to other form-based data collection services that we have experience with, such as GoCanvas (https://www.gocanvas.com/), FastField (https://www.fastfieldforms.com/) and others.

Where these third party services typically provide a cloud-based service for the data upload and retrieval, Survey123 stores the collected information in standard ArcGIS feature layers, so they can be queried using the ArcGIS REST API or integrated directly into your existing maps and visualization tools. To automate the download of data from Survey123, we have developed a Python script that will authenticate with ArcGIS, apply a date-based filter to the primary storage feature layer and download the data in CSV format. The script includes options for control of headers, date and time attribute merging and "other" attribute merging. The script will not only download data from the primary feature layer, but also from any related feature layers, if you have defined relationships between them. The default CSV format is typically suitable for direct import into KiWQM and KiECO. The script also supports a pluggable architecture, so if you need the data in a specific format (CSV for import into the Hydstra SAMPLES and RESULTS tables, for example), a format other than CSV, or need to define your own processing rules, you can do so.

We Speak MQTT (Somewhat)

MQTT (Message Queuing Telemetry Transport) is a lightweight protocol used by IoT (Internet of Things) devices to transport messages between devices. If you are dabbling in the world of IoT (and who isn't these days) KISTERS

already have some experience in subscribing to basic MQTT feeds, using both Hydstra and Hydrotel. In the Hydstra context we use a Python script to receive the MQTT stream from IBM Watson and pass the data directly to SVRIMP for processing. HyQuest supports their own ML-IoT protocol from Hyquest loggers via a Hydrotel driver.

It is worth noting that MQTT doesn't describe or constrain the content of IoT messages in any way, so each new source of IoT data is likely to require a new project to retrieve and process the data. IoT devices typically talk to a broker of some sort, and every big player in the IoT space also wants you to use their own broker, though you can run your own, with consequent further complications of authentication and log-in.

Talk to us if you're interested in feeding MQTT data to WISKI or Hydstra.

Hydstra Product News

Hydstra V13 In Production

Most of our users have upgraded to Version 13, which was released in May last year. Please consider upgrading soon, particularly if you are still running older V10 or V11 systems, as they have become unsupportable. Contact support@kisters.com.au for details on how to download the latest version. Note that you can only upgrade one version at a time, so if you are still on V11 for example, you need to go V11->V12, and then V12->V13.

If you don't feel confident to perform the upgrade yourself, KISTERS can assist as a small paid consulting project, provided of course that we can gain access to your server through VPN or Citrix. For very small systems another option is for you to send us the whole system either on a hard drive or as a zip file uploaded to Dropbox. We will upgrade it, and return it to you upgraded to the current release. This of course assumes you can live without it for a while.

64 Bit Hydstra?

In this 64-bit world we are considering our options regarding releasing Hydstra V14 in the next year or so as a 64-bit only application. From an internal development viewpoint there don't appear to be any impediments to going to 64-bit. From an end user perspective there wouldn't appear to be much difference, since Hydstra doesn't require the huge amounts of additional memory that 64-bit processes have available. It's doubtful there would be a noticeable performance difference either. If we went 64-bit we would include 64-bit Perl and Python in the distribution so your existing jobs in INIPATH would largely run as before.

The biggest potential problem would that existing external 32 bit software could not directly link to a 64 bit DLL, so if for example you were running 32-bit Excel you couldn't load the 64-bit HYDLLP. To that end we could possibly release a 32-bit DLL in tandem, with a slightly different name like HYDLLP32.DLL, but you would need to make minor changes to any 32-bit software that loaded the DLL. Those users using SQL Server as the backend wouldn't need to make any changes, as we don't care whether you we reusing 64-bit or 32-bit SQL Server.

We welcome your feedback on this subject. V14 is likely to be at least a year away, so there's no immediate hurry, but it would be helpful to have some guidance from our user community.

Hydstra as a Service?

We are in the position to offer Hydstra as a service to small Hydstra users in Australia and the Pacific region, hosted on our own servers. We see the initial target clients as the one and two person operations which need to use Hydstra to service their data collection contracts. You would access the system via a web browser to a Citrix front-end, which would then lead you to a session on our server with a dedicated Hydstra system available just to you. KISTERS would provide system upgrades and patches from time to time, we would set up routine scheduled tasks to groom your system and run HYGIENE, and we would also ensure that the system was safely backed up.

Please contact Paul Sheahan for more information in the first instance at *paul.sheahan@kisters.com.au*.

Time Series of Images

We are presently considering a new Hydstra component for storing time-series of images, such as you might receive from a webcam on a logger. Initially we plan something simple, with file-based storage, along the lines of the existing documents system, but without the need to create database records to hang the images off. Later we may look at

more sophisticated big-data stores. Either way, we expect to collect a lot of data quickly, and we will introduce a new HYCONFIG IMGPATH folder keyword, which may point to a large NAS perhaps.

If you are interested in the concept we would be pleased to hear from you with particular details such as:

- How do the images arrive? Are they pushed or do you pull?
- How many a day or month are taken?
- How big is each image, roughly?
- Do the images contain EXIF data internally with date and time taken?
- How many cameras are situated at a site? One only, or more than one?
- How many sites have cameras now, and how many in future do you expect?
- Do you take just static images, or movies as well?
- What do you expect Hydstra to do with the images? At least show them in the Workbench.
- Do you want to access them from the web somehow via Hydstra/WEB or webservice.exe?
- We expect that there will be some new DLL calls to access the images, and a pre-processor system to store them away.

Create a Rating from a Section in HYRATED

HYRATED now has a streamlined workflow for creating *theoretical* ratings – when you have a surveyed cross-section, but have no gauging measurements (or only a few, and they don't cover an adequate stage range).

You can divide the stage range up into one or more flow regimes, and for each subrange choose Manning's or Chezy's formula, estimate or measure the various coefficients (slope, roughness) and HYRATED will create a rating table that follows the resulting calculations.

By switching on the new Mannings Mode feature, these equation segments become ***live***, and by dragging inflection points, you can trigger HYRATED to calculate the equation parameters so that the curve goes where you want it. Which parameters? You can choose whether you are more confident in your chosen slope or roughness coefficient, and HYRATED will hold that one fixed and recalculate the other one.

When more gaugings become available, you can use Mannings Mode and drag the inflection points for your flow regimes (if there are more than one), to calibrate the theoretical rating so that it goes through the measured gaugings.

You can even configure that some flow regimes use a method of *None*, so that the live fitting by equation leaves them alone, preserving your manual edits.

PSC and HSC Jobs on the WEB

Any Hydstra .EXE or .HSC job that is driven by a parameter screen, such as HYPLOT, HYCSV, HYDAY, HYSITREP etc can be configured to be offered by the owner of a Hydstra/WEB site. We recognise that parameter screen jobs are for the more advanced user, so they may be relegated to an Applications tab in Hydstra/WEB, but it is a significant new offering for Hydstra/WEB providers and we encourage providers to discuss the possibilities with us. Each job you wish to make available requires a small amount of configuration work from us to set it up. Of course this facility is only available in V13.

You can try out examples of a few programs configured at http://vm-hyd-webdemo.canpa.kisters.net/static.htm :

KISTERS

help | login | contact | customise

Rivers And Streams » Surface Water Sites » Catchment 1

HYDSYS01 HYDSTRA TEST STATION - COMPOSITE DATA

bookmark this page

All data times are Eastern Standard Time

Rivers and Streams

favourites | search | find a site

- Surface Water Sites
  - Catchment 1
    - HYDSYS01 Hydst...
    - HYDSYS02 Hydst...

data use licence

glossary and metadata

Latest Values | Details | Prepared Outputs | Custom Outputs | Applications

GWHYPLOT	Groundwater Plot of all Data for a Site
HYANN	Report Annual Total Flows
HYCSV	Extract Comma Separated Data
HYDAY	Annual Table of Daily Values
HYEXTR	Produce Flat Files of Hydstra Data
HYEXTREM	Report Extremes of Hydstra Data
HYFLOW.PLOT	Flow Duration Analysis
HYFLOW.PRINT	Flow Duration Analysis
HYMONTH	Tabulate Monthly Values for Period of Record
HYPLOT	Plot Hydstra Data
HYQUAL	Quality Summary Report
HYSITREP	Site Summary Report
HYVALUE	Report Single Hydstra Data Value
WQX	Extract WQ Data, Build Warehouse, CSV Extract

Disclaimer: It is not guaranteed that the data is current nor that the data or data processes are free from corruption or error.

Last modified: 26/03/2021 14:55:51

HYANN - Report Annual Total Flows

Site List	HYDSYS01	
Data Source	A	Archive Files
VarFrom	100.00	Stream Water Level
VarTo	150	Stream Discharge Volume
Start Month	1	
Start Year	0	(0 = Period of Record)
End Year	0	(0 = Start Year / P.O.R.)
Each Day to	END	
At Time	2400	
Report Type	DESCEND	
Output File	TXT	

Run | Close

```
Hydstra Evaluation System                          HYANN V86  Output 26/03/2021

Site      HYDSYS01        Hydstra Test Station - Composite data
Variable 150.00           Stream Discharge Volume in Cubic metres

Year    Annual              Days            Rank
        Total               Missing
        (Cubic metres)
-----------------------------------------------
1974      186300000                          1
1975      133700000                          2
1978      120200000                          3
1961       84240000                          4
1976       84230000                          5
1960       58410000                          6
1963       54210000                          7
1984       50880000                          8
1966       43690000                          9
1964       41660000            14           10
1962       39480000                         11
1969       37540000                         12
1985       36080000                         13
1970       32420000                         14
1983       27950000                         15
1981       23840000                         16
1971       23670000                         17
1977       22340000                         18
1967       17980000                         19
1973       16880000                         20
1968        9940000                         21
1979        8715000                         22
1972        6616000                         23
```

HYANN

Detailed Technical Advice on Running Hydstra over a Wide Area Network (WAN)

Don't!

Use Citrix or RDP to run Hydstra running on a server on the same LAN segment and logically and physically close to the Hydstra file store. Hydstra is very file system intensive, and we find that WANs simply don't respond quickly enough, leading to all sorts of issues including the dreaded C000006 exceptions, which are in all cases caused by slow network responses.

Please contact Hydstra Support for advice on accessing your own Hydstra system from remote locations, and be aware that for very small users we may be able to host Hydstra for you.

HYTSEXPORT Exports Hydstra Time-Series to your Warehouse, Data Lake or Swamp

Agencies are increasingly looking to big-data solutions for publishing Hydstra data to the world. As part of that push we have developed a new program HYTSEXPORT, which is an all-Delphi successor to HYGENEXP. HYTSEXPORT is highly configurable, but basically it exports recently changed time-series data to a delimited text file. It maintains a cache of the latest block of each variable so if the only thing that has happened to a data file since last run is that a few more data points have been added, only those few points will be exported. Of course if edits are made earlier in the block, or to previous blocks, or to ratings (in the case of flow), then whatever data has been affected will be re-exported.

HYTSEXPORT is significantly faster and more flexible that HYGENEXP, and is well suited to being run on a trigger each time a TS file is changed. The features of HYTSEXPORT include:

- Flexible output format - selectable field list
- Configurable field and line delimiters
- Optional site and variable metadata
- Configurable date/time formats
- Full mode/Change mode/Partial Period Mode
- Run date on each point for infinite history
- Point mode option for irregularly spaced data
- Block indicators for point mode delete/replace
- Quality code recoding to arbitrary names
- Quality code description output from QUALITY table
- Special value for missing data, e.g.-999999
- Configurable file naming - all in one file, or site files, or site/variable files
- Configurable trace naming rather than just variable names
- Mapping variable numbers to sensible names
- Perl or Python post-processing built in
- Time of max and min for period based max/min
- Report format distinguishes between level type data and total type data
- Output compression option to reduce file sizes
- Output comments from time-series
- Optional column headings at the start of each file

If you are pushing time-series data to other systems please have a look at HYTSEXPORT. We recommend you download a recent patch before going into production with it.

A swamp (not necessarily containing useful data)

Tidy up your old report, HYDLOG and SVRIMP Folders

HYZIPDIRS is a program for zipping up old folders of data - in particular SVRIMP raw files, HYDLOG files, and report folders under PTMPPATH. Please consider running it routinely in AUTOJOB as a matter of course to keep your Hydstra system tidy. Any process that backs up your system, or copies it somewhere else, will run immensely faster copying one zip file per day rather than hundreds or thousands of small files per day.

There is little point in keeping HYDLOG and report files from years ago, so just go to the top of the folder and delete whole years. SVRIMP raw data is another matter, and as a general policy we recommend keeping raw incoming data forever.

Antivirus in Data Folders

In recent times the threat of phishing and other problems have increased the need for good antivirus protection. However particularly for Hydstra, which is file intensive, we recommend trying to keep antivirus away from folders that are accessed frequently, and which don't hold executables, such as DBFPATH, INIPATH, PTMPPATH, WORKPATH, HYDPATH, TSPATH etc. if at all possible we suggest you mark RUNPATH as read-only and writable only by a system administrator and then apply an antivirus exemption to RUNPATH as well.

We have certainly seen situations where corporate antivirus removed Hydstra programs as quickly as we installed them, and we have encountered other situations where a change of antivirus resulted in a 200% slowdown loading Hydstra programs and scripts. We also see issue where a Hydstra program , having just written and closed a file, is denied access to write to it again, and we suspect that's because antivirus got in the way first.

HYGIENE Test 33 Looks for SVRIMP Errors

If you are running SVRIMP you should be checking every day for a dated folder under \hid\log\svrimp (configured in SVRIMP.INI) containing errors. The folder will be named something like \hyd\log\svrimp\20210330_ERRORS. If there are files in this folder the .txt file will explain the error, and the .csv files contain the offending data. Typical problems include mal-formed SVRIMP files, invalid datatrans, bad variable numbers, etc.

HYGIENE test 33 now looks back for the last 30 days and reports a failure if any such error folders are encountered.

Hydstra Creating Time-Series Data for Non-Registered Sites

Hydstra can be a bit wild and loose creating time-series files for sites that aren't registered in SITE. By default you can import data for non-existent site and nobody complains. While HYGIENE will report on these (and many other) problems, most people just roll their eyes when we start talking about HYGIENE.

As a half-way house, there are a few HYCONFIG options in the [Time Series] section which will slowly help you get your house into order. By default in HYCONFIG we deliver

```
ABRTNOSITE = ALLOWWORK,ALLOWARCHIVE
```

which says we don't care if you create a time-series file for an unregistered site. We ***should*** care actually, since HYPERUPD won't update the PERIOD table without a parent site, for example.

To gently ease your way into a better world, you might like to try:

```
ABRTNOSITE = WARNWORK,WARNARCHIVE
```

which will issue a warning every time you save data for an unregistered site.

Next stage once everyone is well behaved might be to go for:

```
ABRTNOSITE = WARNWORK,ABORTARCHIVE
```

And finally:

```
ABRTNOSITE = ABORTWORK,ABORTARCHIVE
```

Note that temporary files (0-9) are not affected by this setting.

You might like to consider tightening up on ABRTNOQUAL and ABRTNOSUBV as well while you're at it.

Copying HTML Hydstra Outputs to Excel

Many HYSCRIPT jobs produce HTML output. While sending these outputs directly to destination EXCEL produces fairly unsatisfactory results, copying the result from HYXPLORE with Ctrl/C and pasting into Excel does more or less exactly what you expect - dividing tables neatly into columns, allowing you to do subsequent analysis on the results.

Try HYMAILER_PY.HSC with Secure SMTP Servers

HYMAILER.HSC is the generic Hydstra email program that uses an SMTP server to send email. HYMAILER has difficulties connecting to secure SMTP servers, which are increasingly becoming the norm, and we suggest you use the functionally equivalent HYMAILER_PY.HSC which is implemented in Python instead of Perl, it seems to deal better with secure SMTP. In V14 we will probably replace the Perl-based HYMAILER.PSC with the Python based version.

Partial Aggregations with Missing Data

The bane of everyone's life is missing data - it mucks up reports, and a tiny bit of missing data can leak forward to daily, monthly, and even annual totals or means, making a whole year go missing because of a 15 minute bad data point. Our general advice is to edit across gaps with the best estimate you can make, quality code the estimate accordingly, and hence avoid gaps and missing data.

There's a little understood INI file called PARTAGGR.INI, which you can configure to decide that for some variables a certain amount of missing data is acceptable - missing in this context meaning either completely absent or bad quality data. PARTAGGR.INI allows you to specify a percentage of missing data that you will accept when computing period means and totals.

The feature was originally introduced years ago specifically for air quality monitoring, where it is not uncommon to take 80% data as acceptable.

However you do need to think carefully before implementing anything in PARTAGGR.INI. Particularly for water level and flow, a common reason for missing data is that the instrument got washed away. In other words, the most interesting and rare data was lost. Ignoring missing data means you are intentionally distorting your statistics. Personally I wouldn't want to accept anything less than about 95% before I published an aggregation. Even then, 5% of a year is 18 days - a lot of exciting floods can be lost in 18 days!

A further complication is that a table of daily values may not exactly agree with the monthly or annual totals, particularly if the higher aggregations are computed separately.

Feel free to copy PARTAGGR.INI from MISCPATH to INIPATH and edit as you please, but please be aware of the consequences.

HYGAUGE.INI and RATINGS.INI - Allow in TEMPPATH?

There is a difficulty with allowing HYGAUGE.INI and RATINGS.INI to be in TEMPPATH, in that a gauging can be computed from local meter ratings in TEMPPATH, but someone else cannot see the rating from which the gauging was computed. Of course this problem mainly applies to old-fashioned manual gaugings, as opposed to modern gaugings which are imported from gaugings technology of some sort.

Hydstra administrators can set a policy using HYINILOC.INI to forbid HYGAUGE.INI and RATINGS.INI in TEMPPATH, although at present HYGAUGE copies those files to TEMPPATH on start-up if they are not present. This process is of dubious historical origin and we propose to remove the copy to TEMPPATH so administrators can forbid TEMPPATH copies if they wish (and we recommend they do). This of course means that an administrator with write permission to INIPATH will need to add new meter and fan ratings to RATINGS.INI for everyone to use.

Meter ratings can of course be stored in conventional rating tables, and that's another way of making them available more generally.

Please get back to us if you think the change will affect the way you use HYGAUGE.

New HYAUDIT test: Data Within Calibration

The new *Data Within Calibration* HYAUDIT test lets you check your continuously-measured time series data against a sparser calibration data source such as spot measurements or even gaugings. If the difference between your main trace and the calibration points is outside a specified range, the test will fail.

HYTSXML Renamed to HY2WIXML

For the (few) people using HYTSXML to export Hydstra time-series, ratings, shifts and sections to WISKI in XML format, we have renamed the program to HY2WIXML.EXE and for the life of V13 we will distribute both. In V14 HYTSXML will be withdrawn. Please patch your V13 system and switch to the new name when you can.

We have recently developed tools to import WISKI ratings, sections and gaugings back to Hydstra - please contact us if you're interested.

Hydstra Licence Expiry Warnings

Warnings about impending Hydstra licence expiry are delivered only to users who are marked as System Administrators in the PASSWD table. Please ensure you have at least one user who logs into HYXPLORE regularly marked as a System Administrator. You can have as many as you like, but it's essential that at least one of the administrators starts HYXPLORE regularly, otherwise nobody will see the warning until it's too late and things stop working.

WISKI Product News

Release Management and Client Base

As the middle of the year creeps up (time flies!), the current WISKI release version is already at 7.4.11 SR13. Most of you are on 7.4.11, and those of you still on older versions are looking ahead to an upgrade.

In most cases, unless there are specific enhancements that you're waiting on in an upcoming release, upgrades are generally only necessary once every 1-2 years, depending on the circumstances. However given the continual expansion and enhancement to KISTERS' biological and water-quality monitoring software (KiECO and KiWQM, respectively), you may see benefit in updating more often if you collect data in these products. If you're unsure about how regularly your organisation should be updating, or would like some more info in the other modules available, get in touch at wiski-support@kisters.com.au and we'll be happy to advise.

Chief among the changes in 7.4.11 is the integration of the Standard Portal and its applications. As part of a standard WISKI 7.4.11 installation users are able to utilise the modern KISTERS Portal web interface to access, graphically present and export data in the *Data Viewer* app, configure import/export and scheduled tasks with the slick new *KiDSM* app, as well as keep an eye on system performance and track I/O errors using the integrated diagnostic apps *Application Metrics* and *Process Analytics*.

In this newsletter, we decided to add a little more information about the latter two applications. They're particularly useful for those of you who may wish to know a little more about how your system is performing, but are not sure how to investigate calculation/import efficiency more closely.

Monitoring WISKI Server Performance: Application Metrics

Application Metrics is designed to visualise medium to long-term patterns in TSM server performance and graphically display how the system consumes large calculation queues over time. This application can provide administrators with big-picture analysis into server efficiency, and be used to inform whether, for instance, additional server resourcing may be required to optimise the burn-rate of TSM queues.

Like other timeseries data storage in the TSM, this information is stored in timeseries objects in the WISKI client. The Application Metrics app simply extracts this data (as the Water Data Viewer might for hydrological data), and displays it in graphs and heatmaps over a customisable period. These timeseries track metrics such as the Mark Invalid queue, Pending Calculation queue, as well as their associated result (Done/Delayed/Error).

If you want to see these metrics as regular timeseries in the WISKI client, you can display them via the System sites explorer view, which itself is added out of the System view. Drop us a line if you're not sure if you've got them set up, and we'll be happy to investigate and set them up for you.

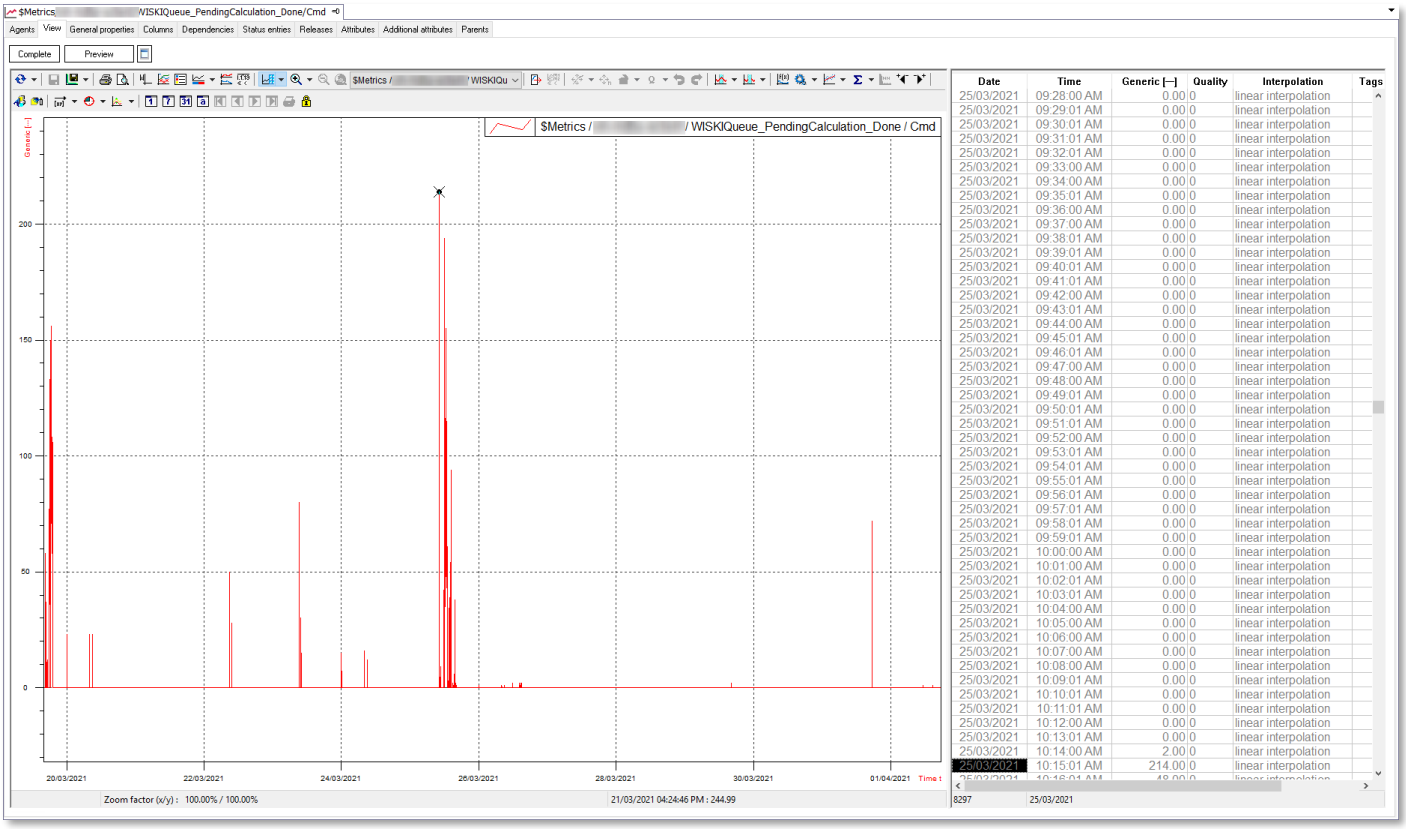

As standard timeseries objects, you're also free to export this data in any available format, should you have external uses for it.

When real-time information is required, such as when checking the burndown-rate of queue entries as a large amount of data is imported and recalculated, the WISKI Server Manager provides other tools, such as the Queue Monitor, to graphically analyse how low your backlog sticks around before being imported:

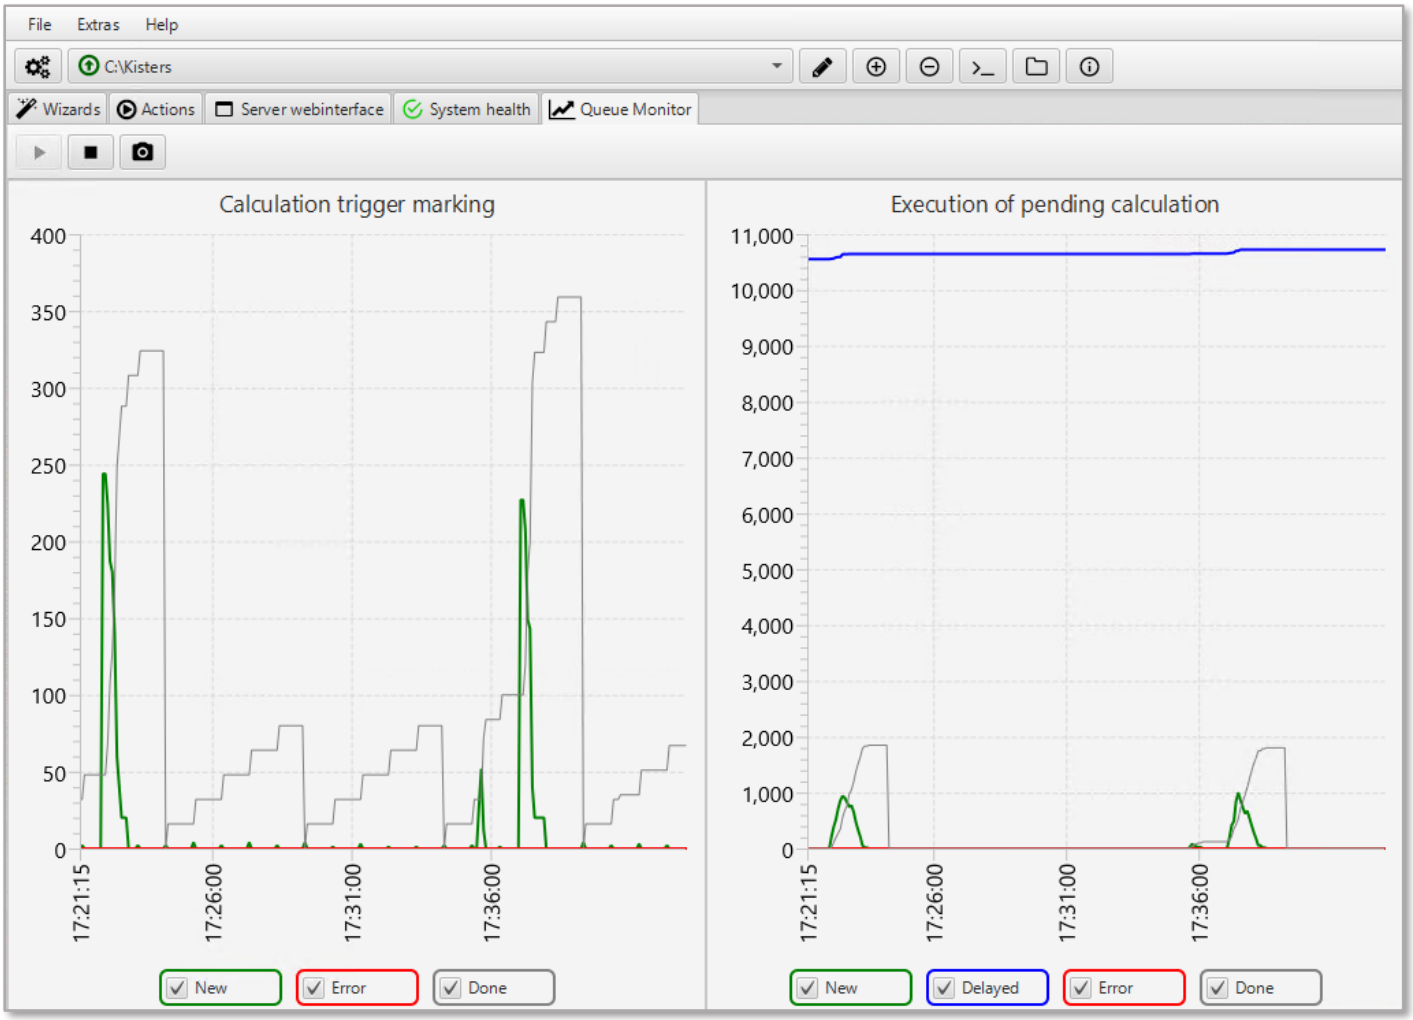

This leads on naturally to the other system administration application mentioned, Process Analytics.

Monitoring Bulk I/O Operations: Process Analytics

Process Analytics is a portal-based log viewer which is specifically designed to retrieve log entries associated with import and export processes. This app is designed to provide temporal analysis of success/failure rates of your import/export processes, but also provide a user-friendly log-viewer to isolate errors and track down the source.

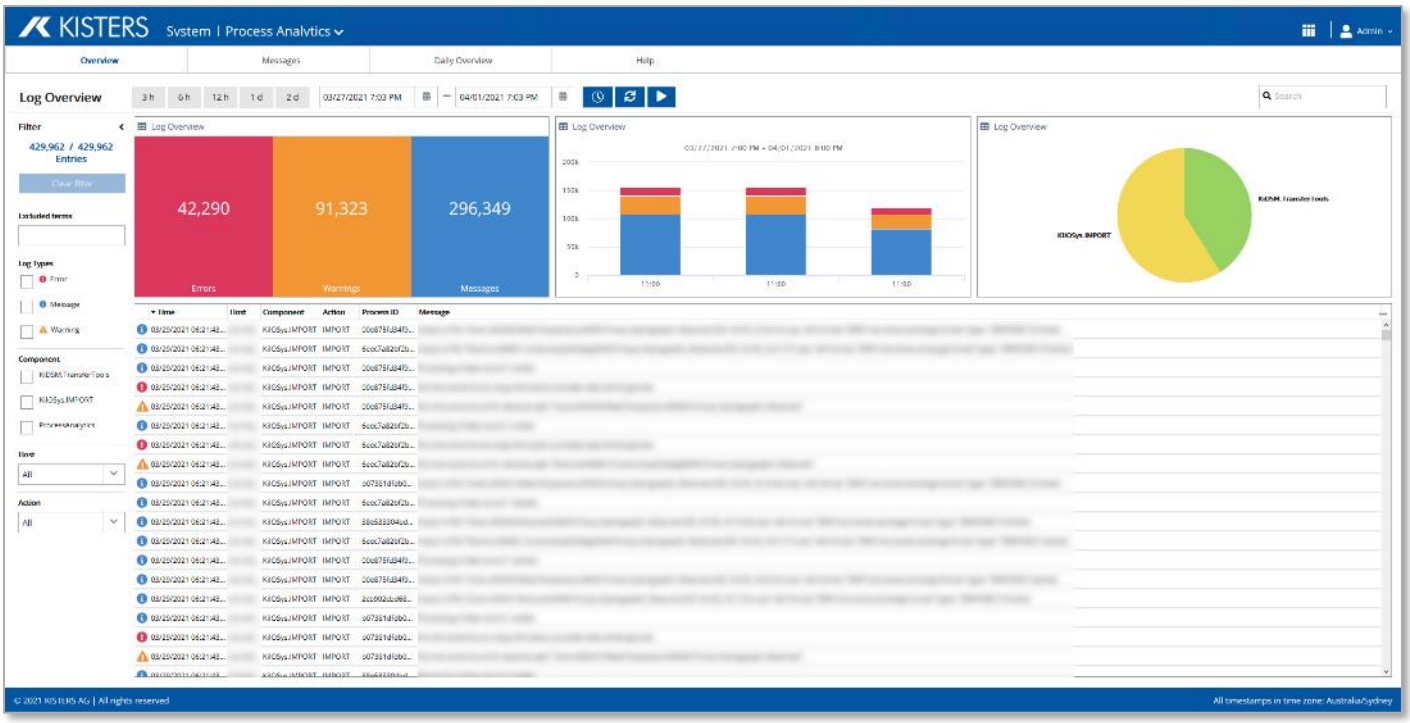

The Process Analytics importer (using Elasticsearch) will import and parse KiDSM and WISKI Server import/export logs (specifically from the KiIOSys node) and display them in a much more accessible way. Chasing down errors in logs – particularly in a system which sees large volumes of data coming and going – can be a messy business, and this app can illustrate the large-scale health of your IO processes, as well as easily pick out specific errors.

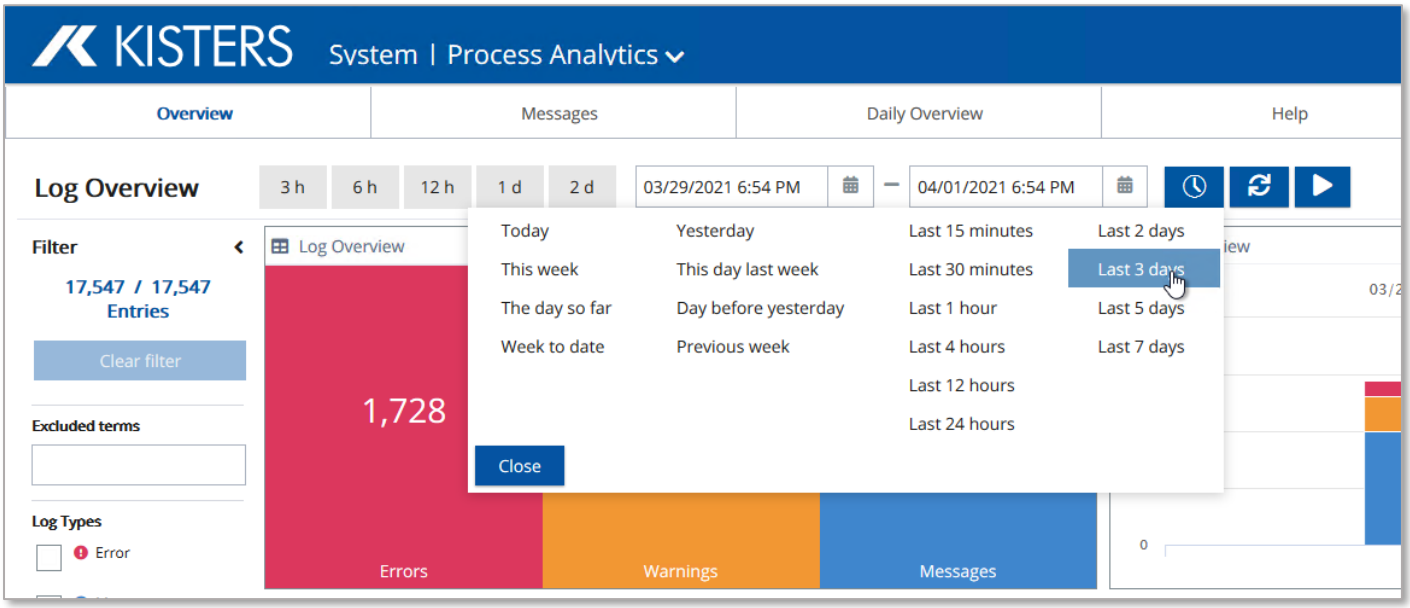

You can select the specific time period of interest, or apply a relative temporal filter. If you're interested in a particular station, for example, you can also use the search bar at the top-right to find log entries pertaining to your search criteria.

You can also interactively select and remove different criteria from the view. If you specifically want to isolate errors, you can simply click the *Warning* and *Messages* columns to deselect them, or uncheck them from the list.

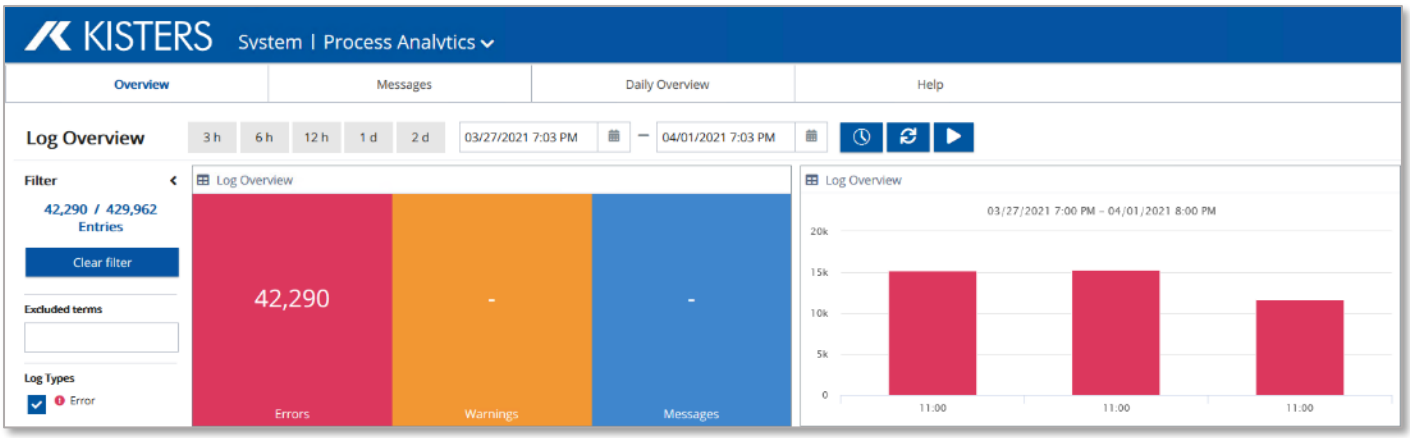

As above, if you're interested in learning more about these apps, or are having trouble accessing them in the Standard Portal, let us know and we'll gladly show you around.

It's worth noting that these apps are most effective (indeed Application Metrics is *only* effective) if you are importing/editing and thus calculating timeseries data. If you are dealing primarily with qualitative or sampling data that is not stored in timeseries, Process Analytics will still be useful if you're using KiDSM to automatically import/export data to/from your system.

Client-side Performance

For users who are not dealing with timeseries data, another quick tool that you might find useful is the *Performance…* metric:

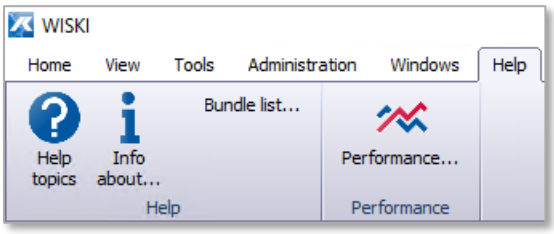

This tool is a useful indicator when you're establishing how the client performs relative to its connection to the database (or specifically, the database server). Take the below test results, for example:

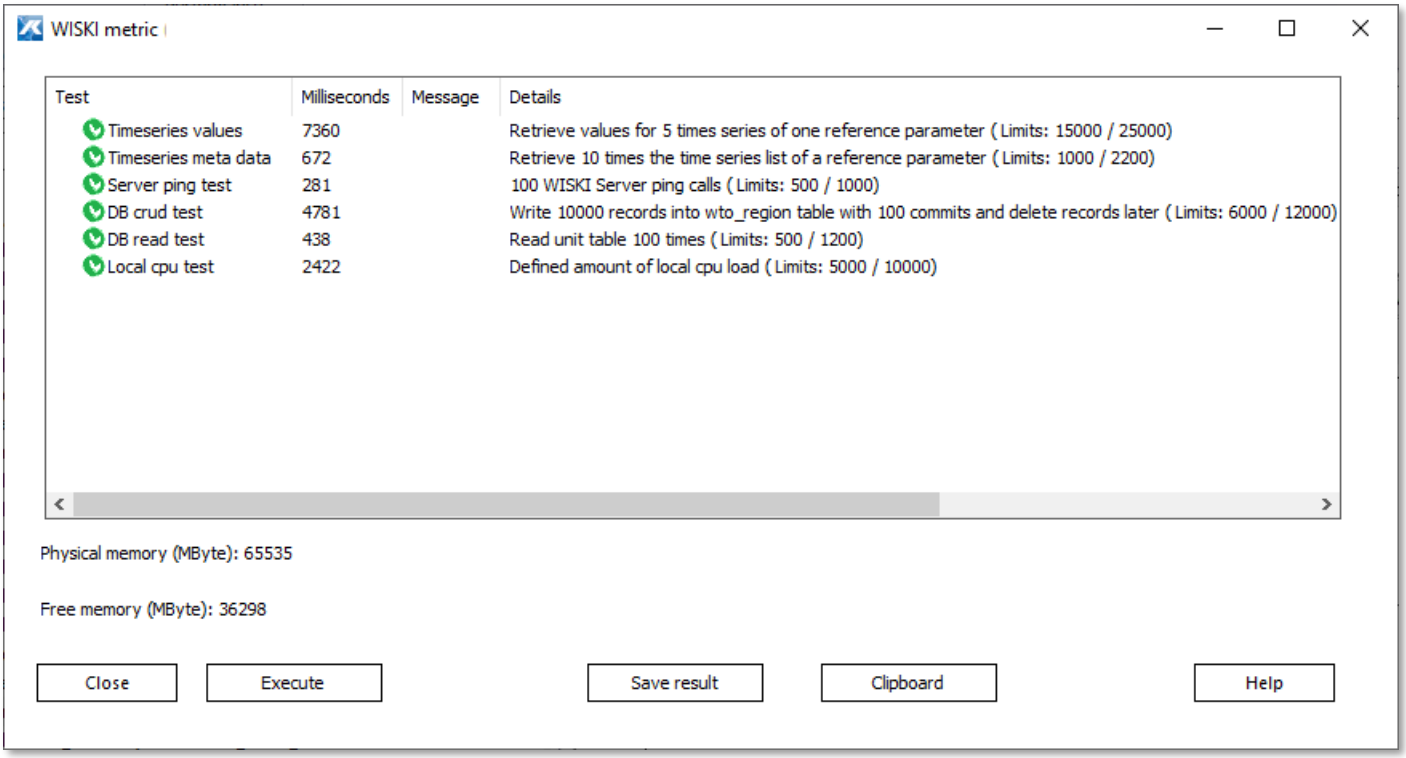

This feedback suggests that:

- The WISKI server (synonymous with TSM Server/TSM, responsible for timeseries data calculations) is resourced well and has a fast connection to the database
- The client from which the test is being run has a fast connection to the WISKI Server (which is, in most cases, hosted on a different machine)
- The WISKI Server is efficiently creating, reading, updating and deleting records from the database
- The WISKI Server is efficiently reading from the same table repeatedly
- The machine on which the client running the test is being run has enough CPU resources for the application

In all circumstances you want to reduce the latency between the machine running the client (which may be a Datacenter server, a Citrix server, other hosting platform, or indeed a user's notebook), the machine(s) on which the WISKI Server is installed and runs, and the database server which hosts the WISKI database. If in doubt, run this metric to see what the client thinks of your setup.

Alarming in KiWQM

Since WISKI 7.4.11 SR6, KiWQM now has the ability to send email notifications via Alarm Manager to recipients on import of data that exceeds specified parameter thresholds. This functionality is critically important for organisations that have strict water quality monitoring requirements such as water utilities that supply drinking water. The workflow is as follows:

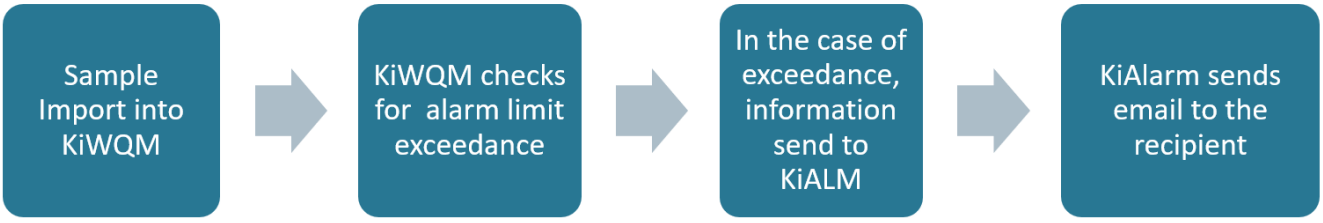

The parameter exceedance levels can be set up within Comparison lists which now allow minimum and maximum warning level thresholds as well as the upper and lower absolute thresholds as shown below.

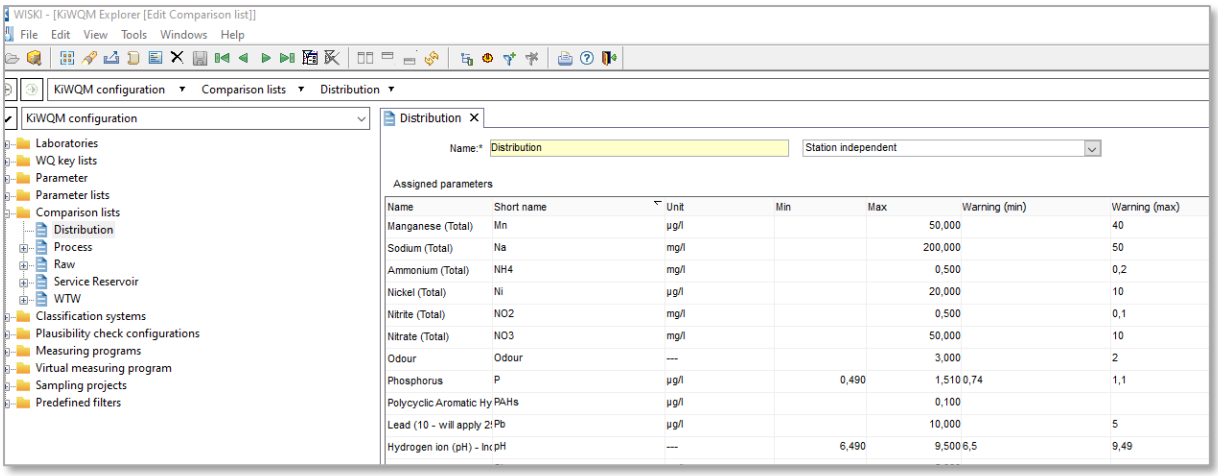

These new warning levels are only used within KiWQM to provide a visual indication that the value has exceeded the warning level as shown below. They will not generate an email notification.

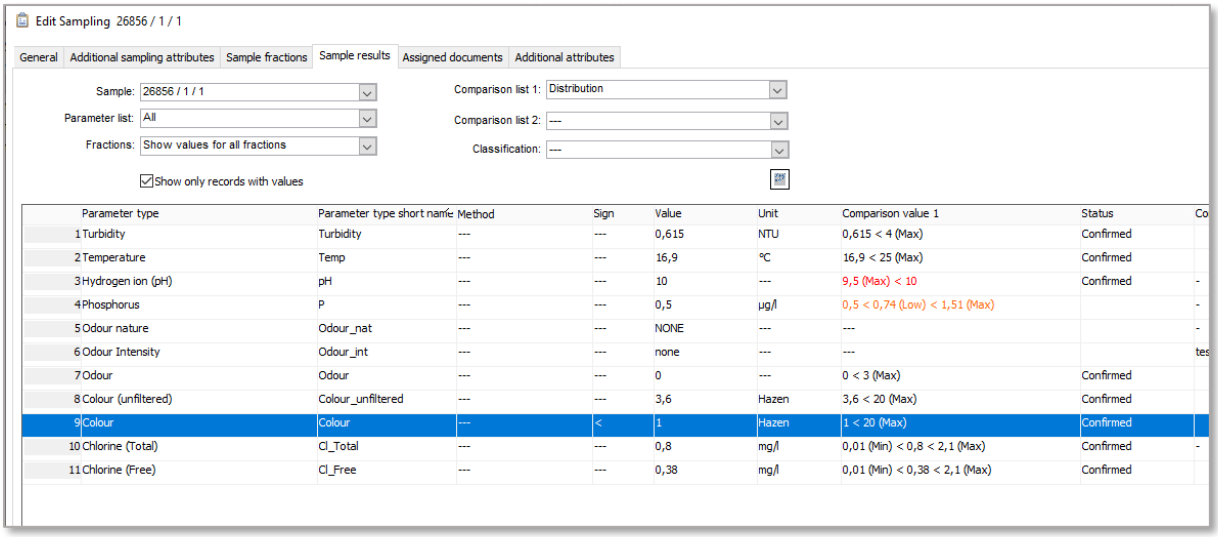

The comparison list that should be employed for alarming is configured within the system settings (KiWQM > Alarm Comparison List).

The alarm notification email is generated immediately on import of data that has exceeded either the maximum or minimum absolute thresholds of a given parameter, an example of which is shown below:

Alarm Manager Notification

KIWQM wants to make you aware that there are changes in your alarm status. Please see the list below for details

Station Nr.	Station Name	Measuring Program	Timestamp	Parameter	Method	Value	Unit	Limit Min.	Limit Max.
Douglas-WTW-Final	Douglas WTW	Water Treatment Works	2019-09-01T00:00:00	Nitrate	---	51	milligram per litre		50
01-0001	ABBEYLANDS	Distribution	2020-09-11T00:00:00	Phosphorus	---	0.3	milligram per litre	0.49	1.51
01-0001	ABBEYLANDS	Distribution	2020-09-12T00:00:00	Phosphorus	---	2	milligram per litre	0.49	1.51
01-0001	ABBEYLANDS	Distribution	2020-09-14T00:00:00	Hydrogen ion (pH)	---	3	scalar	6.5	9.5
01-0001	ABBEYLANDS	Distribution	2020-09-14T00:00:00	Ecoli	---	3	MPN/100ML		0
01-0002	ABBEYLANDS IM4 5ED	Distribution	2020-09-14T00:00:00	Ecoli	---	2	MPN/100ML		0
01-0003	ABBEYLANDS-IM4 5EE	Distribution	2020-09-12T00:00:00	Hydrogen ion (pH)	---	13	scalar	6.5	9.5
01-0003	ABBEYLANDS-IM4 5EE	Distribution	2020-09-25T00:00:00	Ecoli	---	2	MPN/100ML		0

KiALM manages the email generation, receiver's contact details, availability, acknowledgement and escalation processes.

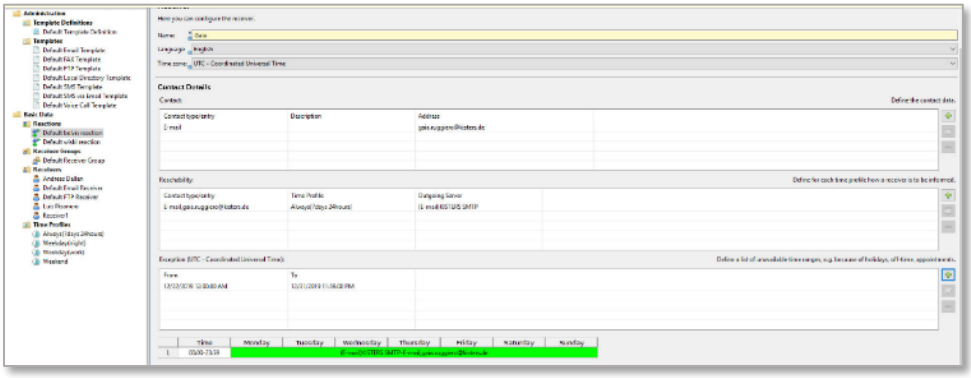

For more information or a demonstration, please contact us directly via the WISKI support email.

News from KISTERS Overseas

Staying Connected Globally

KISTERS offers a global Service Portal for our customers and other interested parties. The Service Portal supplies non-registered users with press releases and general brochures from the company and the Business Units Water and Energy. Further features are available for our licensed customers such as the latest WISKI releases and specific application information and other internal knowledge of your licensed software.

By clicking on the Newsletter section of the Service Portal you can subscribe to various newsletters where you can receive up-to-date information about the company and of the business activities in KISTERS water and monitoring solutions. These newsletters also provide information about access to various resources such as online webinars and training events.

KISTERS global Service Portal: https://serviceportal.kisters.de/en/

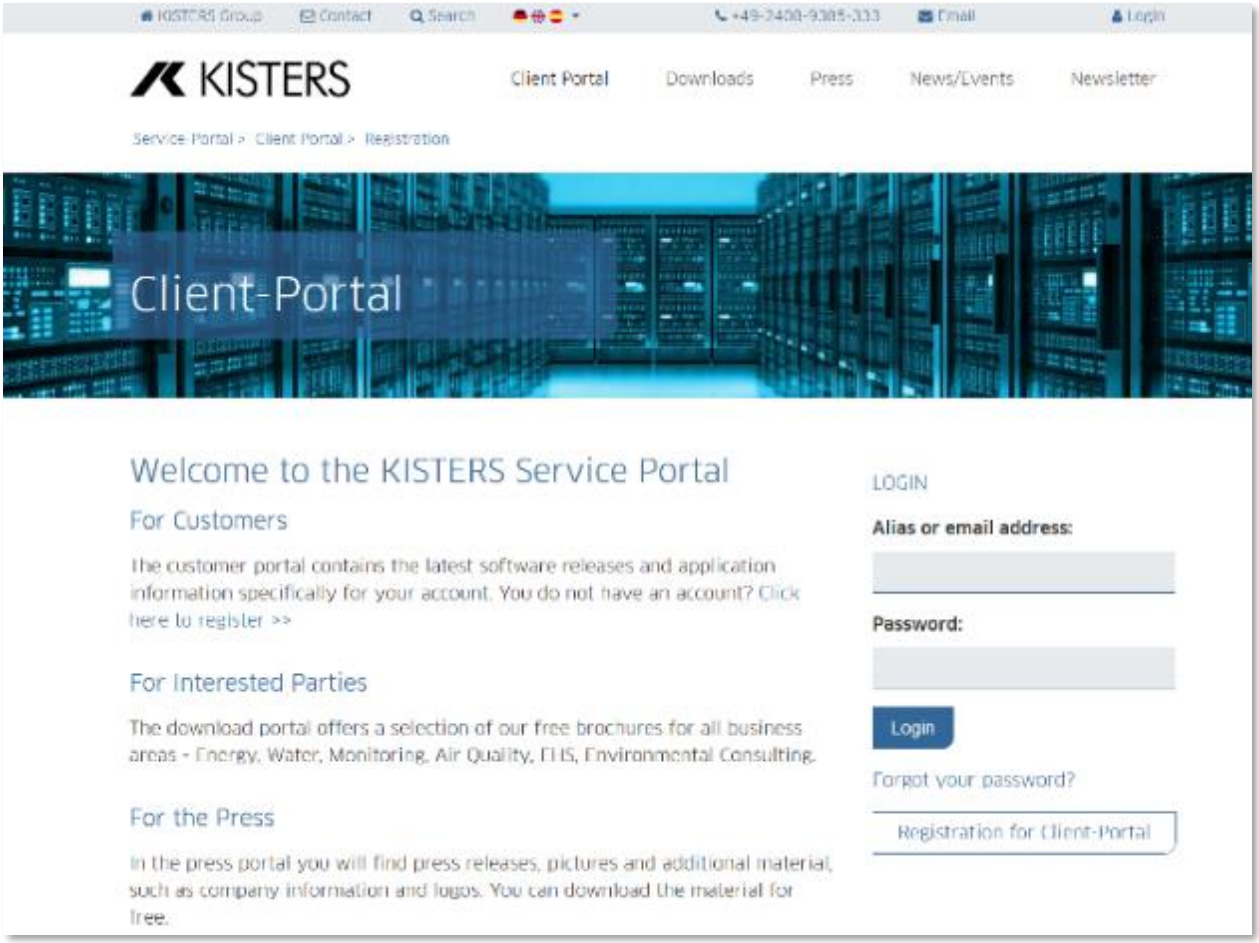

Please feel free to subscribe to the available newsletters or register for the client portal. For users of the client portal, we do however suggest, to use the software download option for WISKI from the Australian download portal at:

http://kisters.com.au/downloadswiski.html.

An overview of news from our offices in Europe and America is accessible via the following links:

https://water.kisters.de/en/press-room/

https://water.kisters.de/en/news/

https://www.kisters.net/NA/news/

Staff News

Anthony Skinner

KISTERS is pleased to welcome Anthony Skinner to the HyQuest Solutions Australia team as Business Development Manager.

Anthony has been involved in the Hydrographic industry for 30 years. Commencing his career as a trainee hydrographer in Sydney Water, Anthony worked on the Sydney metropolitan catchments as a field hydrographer, inclusive of all aspects of data collection, data review and editing, instrumentation and Installation. Later moving into the Sydney Water trading arm of AWT working in underground monitoring. Anthony was provided the opportunity to manage projects in I&I studies conducting monitoring programs throughout the APAC region.

He then moved to Ecowise to spend more time at home with young children, moving back into above ground catchments. Working closely with the development of water quality inset meters into standard hydrometric monitoring flow sites along with the development of more robust communications, to develop and standardise client expectations of today with online data platforms and client available real time data. Anthony also worked on FIFO projects in PNG with Ecowise as hydrographer for Ok Tedi Mining for several years.

Ecowise was acquired by ALS Global, and Anthony took over the role of NSW Operations manager with a team of professionals working on projects in the Government sector through to mining and private sector clients. The last few years saw an additional role of business development into the South American ALS core analytical business footprint with technical hydrometric consulting.

Anthony is current treasurer of the AHA (Australian Hydrographers Association), and a member of the WaMSTeC (Water Monitoring Standardisation Technical Committee) as the committee contributing to the review of Hydrometric guidelines.

Anthony will work alongside Mike Lysaght who this year will begin to transition towards retirement.

Mobile Phone: +61 425 309 193, E-Mail: anthony.skinner@hyquestsolutions.com.au

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@hyquestsolutions.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@hyquestsolutions.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