

KISTERS Australia News

April 2020

Contents

| | |
|--|----|
| Contents | 1 |
| From The GM's Desk | 1 |
| Business Continuity and Coronavirus | 2 |
| Working From Home | 3 |
| Hydstra Product News | 3 |
| Hydstra V13 Released Soon! | 3 |
| Hydstra V12 Released August 2017 | 4 |
| Hydstra/WEB Performance During Events | 4 |
| Managing INIPATH | 5 |
| Exporting Time Series to Data Lakes, Swamps and Morasses | 5 |
| Chained Perl Filters | 6 |
| Saving Hydstra Documentation to PDF | 8 |
| Computing Water Year in MODSYN | 8 |
| Hydstra Help Documentation | 8 |
| WISKI Product News | 9 |
| Release Management and Client Base | 9 |
| What's New in WISKI 7.4.11 | 9 |
| New GUI Features | 9 |
| Archive Data Agent | 11 |
| Groundwater Drilling Profiles | 11 |
| KeyCloak – Single Sign on and LDAP/Active Directory support in WISKI | 14 |
| KiWQM – Sample Planning, Scheduling & Monitoring | 16 |
| HyQuest News | 22 |
| PeciBal: Weighing Precipitation Gauge | 22 |
| HyQuest Web Site Refresh | 23 |
| KISTERS Training | 23 |
| News from KISTERS Overseas | 23 |
| Information | 23 |

From The GM's Desk

By Paul Sheahan, General Manager, KISTERS Pty Ltd

Welcome to 2020!

We, like yourselves, have had a hellish start to the year. Firstly, bush fires, then hailstorms, then flooding rains and now we are all affected by COVID-19 in one way or another. We are continuing to operate with 'business as (mostly) usual' with all our staff being able to work from home as are many clients. For us, supporting our users is our highest priority.

Earlier in the year we conducted an Africa user group meeting. I was a little daunted to run a user group meeting on my own but the great presentations from our users and Gina's excellent organisation made for a very enjoyable and informative event.



Given that the AHA conference scheduled for this year has been postponed to 2021, we are considering our options for a Kister User Group meeting this year. We will let you know once the dust has settled, but a webinar style meeting is one option for us.

With new releases of both Hydstra and WISKI available, we are ready to assist you with system upgrades and the adoption of the product lines new capabilities. Our KISTERS portal over Hydstra work is progressing and we aim to have a production demonstration of this at the user group meeting.

The NZ Hydrological Society recently held their technical workshop in Tauranga, themed 'The Future of Surface Velocity Measurement'. There was a significant amount of information sharing on all things surface velocity. New Zealand Hydrographers are in the great position where home driven R&D is almost written into their job descriptions, so there was a great deal of innovation in techniques and hardware which was demonstrated during the field day. The organisers excelled this year with the field day being rain free!

Paul Sheahan
General Manager

KISTERS Pty Ltd
0409 510 015



Business Continuity and Coronavirus

By Klaus and Hans Kisters

We are closely monitoring the spread of the novel corona virus and assessing the possible consequences for our national and international sites. We have taken the necessary precautions to minimize the potential risk of infection for colleagues and business partners. Herewith we would like to inform you about the measures we have already taken and ask for your understanding if jointly planned dates, activities etc. cannot take place.

Prohibition of business trips

Taking into account the risks of infection associated with personal business contacts, we have prohibited our employees from traveling to and from high-risk areas until further notice. We have instructed those employees who have been in these areas themselves or who have had contact with people who have been there and could possibly be infected to remain in domestic quarantine until a risk of infection can be ruled out.

Computing center

In our certified data center in Aachen/Germany, we operate both our own IT and the KISTERS cloud solutions for our customers. The highest level of security is required here and we implemented a comprehensive security concept, which is of course prepared for critical situations. The measures include, for example, a high-availability concept with parallel operation for the entire server and access architecture, as well as the possibility of secure remote access for our IT administration and service operation support.

Preparations for the maintenance of important processes

With general preparatory measures we safeguard critical functions, such as on-call duty, IT support and the company network, data center operation, customer support, consulting, software development and other internal processes. For this purpose, additional employees are equipped with notebooks so that they can perform the essential tasks from their home office via secure access to the company network. In addition, critical functions can also be taken over temporarily from other locations.

Hygiene measures

All employees were informed about the hygiene measures to be observed. In addition, all employees are instructed to hold internal company meetings only as telephone or Skype conferences and not in person.

The health and performance of our employees is our most valuable asset and we take responsibility for protecting it in the best possible way.

If you have any questions, please contact us by email at support@kisters.com.au, or by phone at +61 2 6154 5200.

Working From Home

By Peter Heweston

As the Covid-19 pandemic unfolds more and more of us are working from home. At KISTERS in Australia everyone is working in safe isolation, save Justin Millsteed who staunchly mans the Canberra office in solitude. He keeps the infrastructure running, and reboots servers as necessary, which hopefully isn't often. I'm reminded of the old joke about the manager showing someone over an automated factory. He remarks that it only takes a man and a dog to run the whole operation, and the dog is there to bite the man if he touches anything!

Most of the KISTERS AU group have been accustomed to working from homes, hotel rooms and cruise ships over many years, so the transition was quite seamless. Our phone system is a VOIP application running on a VM in Canberra, so it is directly accessible to everyone at home. For us it's pretty much business as usual, except we won't be visiting anyone for the foreseeable future. We may request direct access to your system if we need to provide more targeted support, such as assistance with installations and upgrades.

A number of our staff were caught overseas when the pandemic took off, and while everyone got home safely, some are still serving out their isolation in hotel rooms and at home.

Everyone at KISTERS AU has Skype, and we use that to keep in touch amongst ourselves. Clients with Skype for Business accounts can call us individually once you know our Skype addresses - contact individuals by email for Skype details. We have also acquired both GoToMeeting and Zoom licences for hosting larger groups. Of course many of our clients already use well established technologies like Microsoft Teams, Zoom, GoToMeeting etc., and we can respond to meeting invitations and participate via web clients.

So far nobody in KISTERS AU has been diagnosed with Covid-19, and we trust it remains thus.

Stay home, and stay safe!

Hydstra Product News

Hydstra V13 Released Soon!

We plan to release Hydstra V13 within the next month or so. When V13 is release, V11 will become virtually unsupported except for the most dire of emergencies. We will be unable to do further development on V11. If you haven't yet upgraded to V12 we suggest you plan to do so as soon as possible. A couple of larger agencies are considering a double hop, V11->v12->V13. Please contact us if you are still on V11 and find that idea appealing.

The V11->V12 upgrade takes quite a while, as every time-series file needs to be re-written. By contrast V12->V13 is relatively modest and should go through fairly quickly. However V13 involves a new version of Perl, and if you have Perl holdings in INIPATH you may be up for some work to fix a few issues. We provide a tool in recent V12 patches

called HYCHKPERL.HSC which points out various issues, and we can help you rectify the issues as a small consulting project.

We have released a Beta version of Hydstra V13, and you can download it from the KISTERS web site at the usual place. You will need to request new HYACCESS files using <http://kisters.com.au/hyaccess.html> and then download the software from <http://kisters.com.au/downloads.html>. Contact KISTERS Support for a login to the downloads page.

Under no circumstances should you go into production with the V13 beta yet - it is strictly for testing and evaluation purposes only at this stage. Keep an eye on our web site for a final release.

Hydstra V12 Released August 2017

Hydstra V12 was released quite a while ago, and is still the current release system available on the KISTERS web site. Please plan to upgrade to V12 soon if you haven't already, as V13 is already waiting to spread its wings.

Please contact support for a V12 HYACCESS file, and plan to do an offline test upgrade first if at all possible. The upgrade from V11 to V12 is quite routine now, but can be time consuming as all time-series files need to be upgraded.

Hydstra/WEB Performance During Events

Hydstra/WEB provides a web-based data delivery service for Hydstra clients, and several major state agencies run Hydstra/WEB. It operates in one of two modes, dynamic and static. The dynamic mode produces product for display on the fly, resulting in fairly heavy loads on the web server. We realised a long time ago that this resulted in unsustainable loads on the web server at the very time (during floods) that people most wanted to see the data. In response we developed static mode, in which produce is pre-computed on the back-end server, with very little load on the web server. The result is that the static mode server can easily deal with 1000 concurrent users browsing web pages.

The downside is that significantly higher back-end server load is required to build the product pages. However the back-end load is a product only of the number of sites and the frequency of data coming in, and hence is finite and more or less unchanging, whereas the dynamic server load is a product of the number of users hitting the web site, which is unknowable and potentially very large.

We recommend that Hydstra/WEB users consider moving to the static model, though they will need to provision a larger SVRRUN server to do the computations required. In one state agency a 36 core CPU runs 32 threads of SVRUN to keep the wheels going round.

Consumers of web services also put a heavy load on a Hydstra/WEB server, and they are rather more difficult to manage. I often remark that providing web services is akin to offering a blank cheque to the world - and people will cash it. It's the provider who pays for web services, in terms of the required infrastructure, and not the consumer. There is little incentive for consumers to be efficient.

In order to keep web service users under control we introduce a couple of new features in Hydstra v12. The first is the notion of registered users vs anonymous users. The web service can be throttled for anonymous users to limit the number of calls and the amount of data that can be retrieved per minute. Registered users can be given a different set of limits - or indeed none at all.

For major agencies we recommend that web services be provided from a different web server to the browsable web site. This prevents one class of users from overloading the other. You could even consider separating anonymous users from registered users onto different web servers.

In V12 we introduce a Delphi executable WEBSERVICE.EXE which is functionally identical to WEBSERVICE.PL, but which runs 4 or 5 times faster. We recommend that providers of web services encourage their user community to switch to the EXE version. If possible your IT people might be able to redirect all calls from .PL to .EXE without the users having to change the URL they use.

One problem with offering web services at the moment is that you don't really know who is using them, and whether they are using them intelligently. For example some users make repeated calls, one site/variable at a time. Educating users in the concept of site lists can reduce the number of calls considerably. Furthermore telling users

about the *multi_call* call, which allows multiple calls to be bundled up into one, can reduce overheads and latency considerably. For advanced users the *get_recalculation_periods* call can notify them when data changes, including recomputing flow when a rating changes, for example.

From analysing HYDLOG entries we see some external users web-scraping the Hydstra/WEB pages to extract current values, when it would be much more efficient for them to call the web service call *get_latest_ts_values*. However finding the offending users and asking them to change their ways can be quite difficult. At the moment blocking their IP address and waiting for them to contact you is probably the only mechanism you have.

Finally it's worth noting that in V12 a number of time-series web service calls can return CSV or WaterML2 XML data. Consult the HYDLLP documentation for more help. Only WEBSERVICE.EXE can return WaterML2.

Managing INIPATH

In this era of working from home and web conferencing we get to see into people's bedrooms, studies and lounge rooms. And frankly, sometimes we see things we'd rather not see. The same is true when we look into people's INIPATHs - we often see mess and chaos. Herein a few guidelines for what should be in INIPATH, and how to manage the mess:

- The Hydstra administrator should know what every file in INIPATH does, and why it's there.
- Sort files in reverse date order and in particular look at the very old files. Are they still needed? Old INI files might still be valid - if the application still exists. Old batch files might be history and could be consigned to the dustbin.
- Check all INI files against MISCPATH - if there's no MISCPATH version it's likely the INI file is obsolete, unless it belongs to a program written specifically for your orgcode, in which case the ini file name will begin with your orgcode.
- Run HYINIFILES and check for strays.
- Make an *obsolete* subfolder under INIPATH and move anything there that isn't being used. For files that are changed frequently, e.g. HYBATCH.INI or DATASRC.INI, we suggest renaming the old version prior to move, with names of the form HYBATCH.INI.2020-03-21.
- All files with names like .BAK, .OLD, .myname, .version1 etc. should be consigned to the obsolete subfolder (or deleted)
- Keep a *software* folder under INIPATH, and keep in it installers for software that Hydstra users might need to install, such a Notepad++, PathCopy, Beyond Compare, Process Explorer, Adobe Acrobat Reader, etc. Be sure to keep licence file details for any paid software in a licence.txt file.
- Write access to INIPATH should in general be restricted to TS=3 users. Where it is deemed desirable, specific INI files can be relaxed to TS=2 or TS=1 users.
- Batch files should all have HYLOGIT commands at the start and end announcing that they have run. For example:
hylogit /A=HYRUNPL /T=START "%*"
do stuff
hylogit /A=HYRUNPL /T=ENDT "%*"
Once all batch jobs are instrumented you can use HYDLOGEX to see how often (if ever) they are run.
- If you really want to keep a clean and tidy house, look at configuring and running HYDIRCHK, which allows you to tightly control which files are allowed in various folders.

Exporting Time Series to Data Lakes, Swamps and Morasses

A number of users are starting to think about data warehouses and are looking to populate some external system with Hydstra time-series data. There are a number of approaches available, including using the DLL, and exporting to text files on a routine basis. However to keep a true replica requires something rather more sophisticated. It's unlikely for example that external systems will be able to compute flow or storage volumes, so you should get Hydstra to do the heavy lifting and simply export the result to your warehouse.

This raises the issue of retrospective changes - what should you do if someone changes a rating table in the past, for example? Well the correct answer is that you should recompute whatever data was affected by that table. Hydstra is able to look backwards through the variable conversion system and work out what the impact of such a change

might be. The DLL has a call *get_recalculation_periods* which does exactly that, and HYWDTF_OUT and HYGENEXP use a similar technique.

For relatively infrequent updates we suggest HYGENEXP might handle your requirements, possibly in conjunction with a plugin aimed specifically at the format of the target system. We can assist in developing such a plugin. However HYGENEXP is architecturally not suitable for running against individual sites on a trigger, and for that purpose we have recently developed a Delphi framework that can be run site at a time on a trigger.

A further problem is that Hydstra time-series data is stored in blocks of up to 3000 points, so frequently changing files like telemetry files will constantly report that the whole block has changed. Our Delphi framework and HYWDTF_OUT (but not HYGENEXP) keep a cache file for each last block in a variable, and if the only change is the addition of a few points to the end of the block, then only those last few points are exported.

It's worth noting that a warehouse cannot rely on the time and date of a data point as being part of a key identifying the point, as block start and end times can be changed in the workbench, and points can be inserted and deleted. We don't keep a record of deleted points, only that the block has been modified. The only safe way to keep another store in sync is to send it delete commands for each changed period, followed by a replacement string of points. HYGENEXP has this capability, and you should look at the *hygenexp_mwcp* plugin for an example of such a system.

Any external data store must also preserve the notions of missing data or gaps, and of datatrans (how points are joined together), otherwise subsequent interpretation of the data will be incorrect, particularly if interpolation between values is required. Producing aggregate values such as daily means requires a fair bit of understanding of how interpolation between points should be carried out, particularly if there aren't points exactly at the start and end of the requested period. A further complication is that datatrans can change even within a Hydstra block - for example older rainfall was probably recorded periodically as DT5 while modern tipping bucket rainfall is recorded tip by tip as DT6.

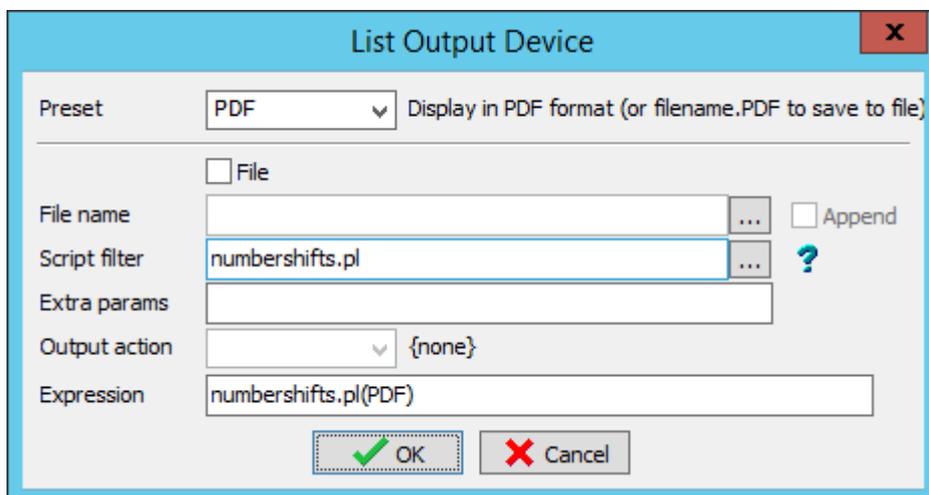
In terms of sizing a warehouse, run HYDATSUM (it takes a while) to find out how many raw data points you have stored in Hydstra. Multiply that by the point size in the new store , and add space for derived variables such as flows and daily aggregations. The result may surprise you!

If you are looking to replicate Hydstra time-series data in near real time to an external store you should probably consult us about developing a targetted application in Delphi for optimal performance. It's likely to take days of consulting work, depending on the complexity of the target system. You will also need to be running SVRRUN if you want trigger-based replication.

Chained Perl Filters

In some cases, it may be desirable to chain two or more output Perl or Python filters together. This is particularly useful when you want to modify the regular output of an application and have that modified output saved to a different file format, such as RTF, HTML or PDF.

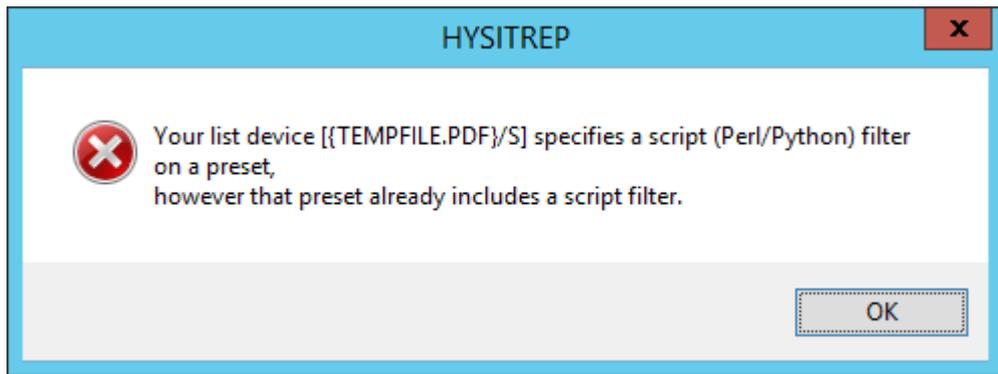
Since the file format conversion is done with a filter, you cannot specify your own filter to modify the output. For example:



In this case, the PDF preset, as defined in LISTDEV.INI, already uses a filter, hypdf.pl:

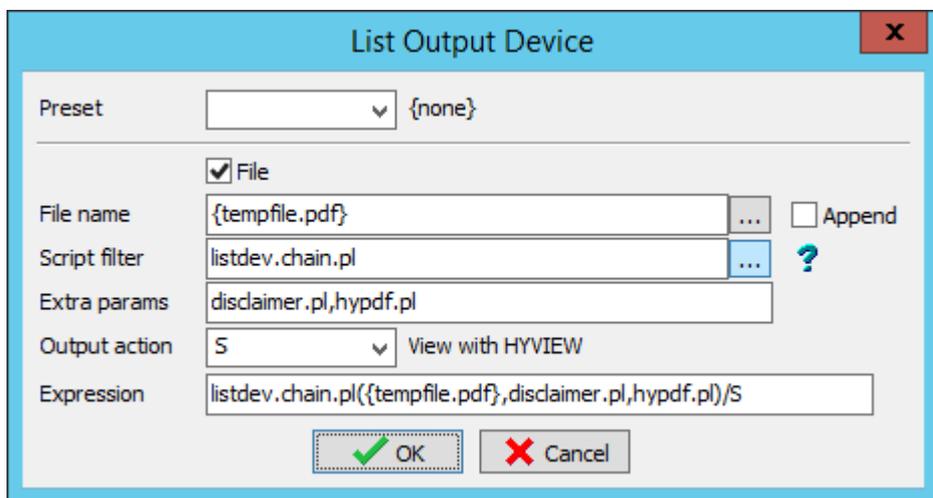
```
PDF = hypdf.pl({TEMPFILE.PDF})/S Convert to PDF and Display
```

Attempting to run the application with two filters will fail:

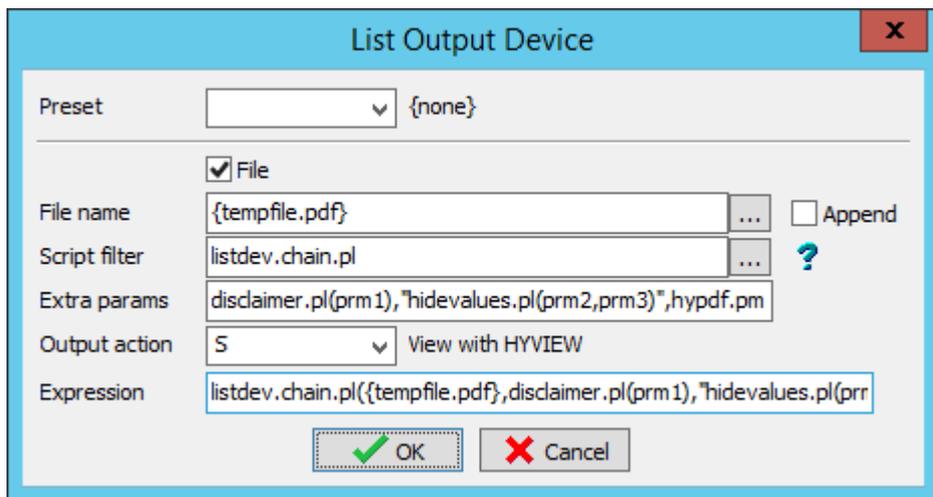


To run an application with multiple filters, you must use the *listdev.chain.pl* Perl script or *listdev.chain.py* Python script. These are special filter scripts whose parameters are a list of other filter scripts. Any number of filters can be chained together, as long as the output of one filter is expected as the input to the next filter.

For example, to modify a standard report and then convert the report to PDF:



If the chained filters require parameters, they can be specified as well. If a chained filter requires two or more parameters, enclose that complete filter expression in double quotes:



Chaining even two filters together results in a complex expression, so it is highly recommended that you save these expressions as jobs or as list device presets, in LISTDEV.INI.

Saving Hydstra Documentation to PDF

You can save sections of the Hydstra Help file in V12 to a single PDF file fairly easily using the following procedure:

1. Open HYPLORE but close any outputs in it
2. Open any number of Help pages in successive tabs
3. Run HYPLORE File/Save All to a PDF file

Computing Water Year in MODSYN

Sometimes you wish to model something in MODSYN based on water year rather than calendar year. It's a few lines of MODSYN code to derive the water year from the calendar year and the HYCONFIG value SMONTH, which is the start month in the water year:

```
scalar smonth
vector wateryear
smonth=hyconfn('smonth')
wateryear=iif(smonth>1,iif(month(>)=smonth,year(),year()-1),year())
```

If HYCONFIG.SMONTH=10 (water year starting in October), this will develop a vector containing 1997 for 1998-10-31, implying the 1997/98 water year, and 1998 for 1998-11-01, implying the 1998/99 water year. On the other hand, if HYCONFIG.SMONTH=1 then the water year and calendar year are the same.

Hydstra Help Documentation

Fun fact - if you assembled the whole Hydstra manual into a printed document, it would be 7000 pages long, and that's without any page breaks for chapters and topics. That's 7 reams of paper printed double sided. There are some 1.7 million words of Help documentation.

WISKI Product News

Release Management and Client Base

As mentioned in our previous newsletter, WISKI 7.4.11 has since rolled over to become our current product version and our recommended release. The most recent service release is WISKI 7.4.11 SR7, which includes a range of new features and GUI enhancements.

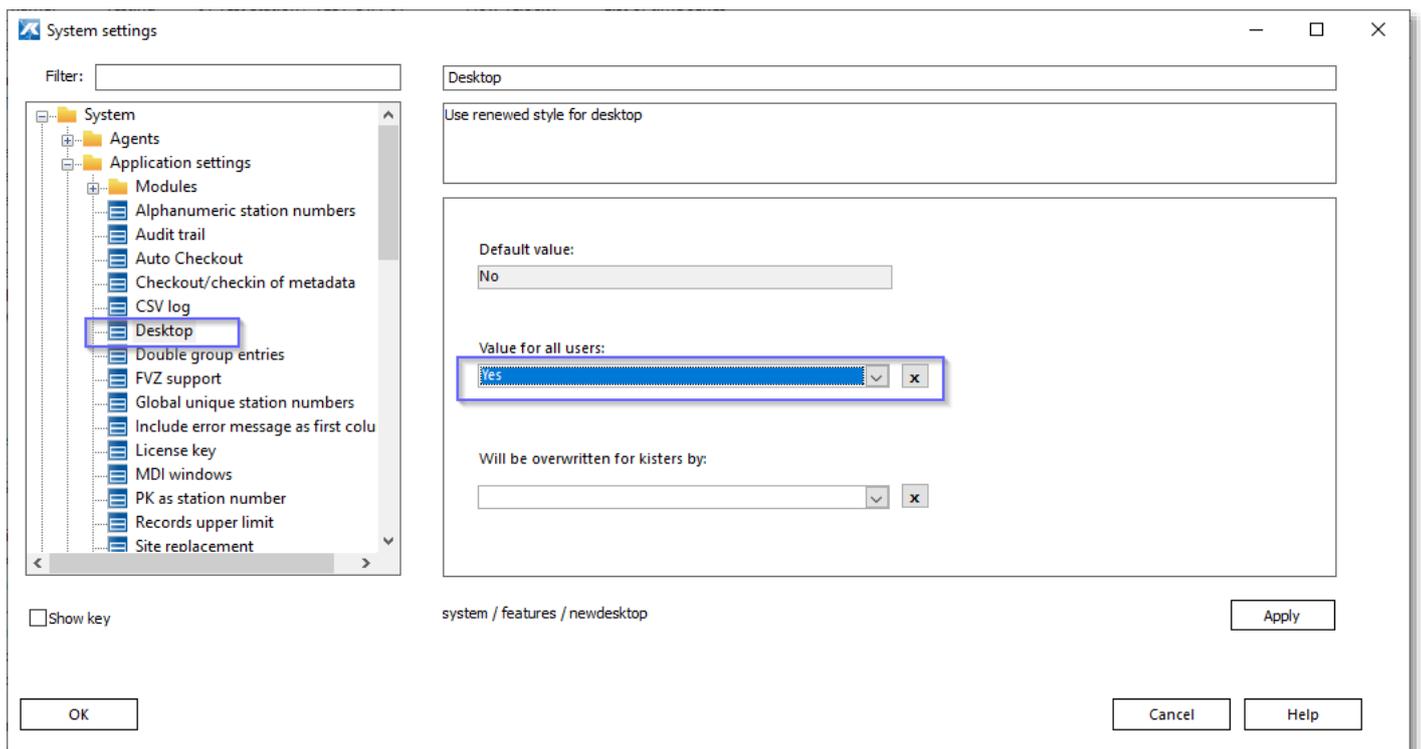
While older versions will always be supported in some capacity, upgrading to a newer version will enable access to new software enhancements, any and all bugfixes, and access to the continuing developments around KISTERS Portal based applications. We recommend planning an upgrade to 7.4.11 in the near future to take advantage of the exciting new features in the software. Feel free to get in touch at wiski-support@kisters.com.au for advice on coordinating the upgrade, or any queries about what's on offer in the latest release(s).

What's New in WISKI 7.4.11

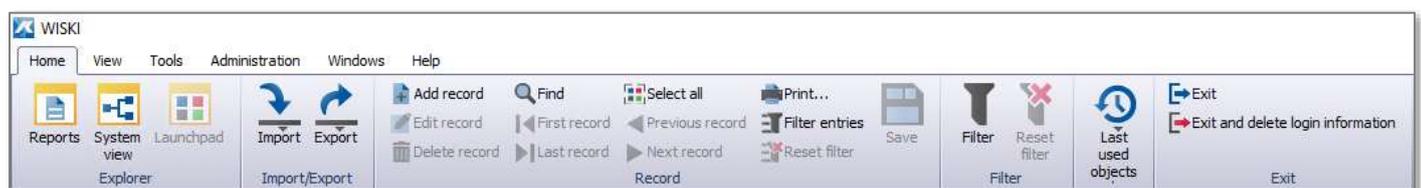
The new version has a number of new features within the client to streamline some of the most common processes, make the user experience more intuitive and refresh the GUI appearance with a more modern feel. This includes the introduction of a new toolbar with larger, more representative buttons as well as new menus, which will aid new users in finding their feet as they navigate through the application, as well as enable experienced users to carry out their regular processes with fewer clicks.

New GUI Features

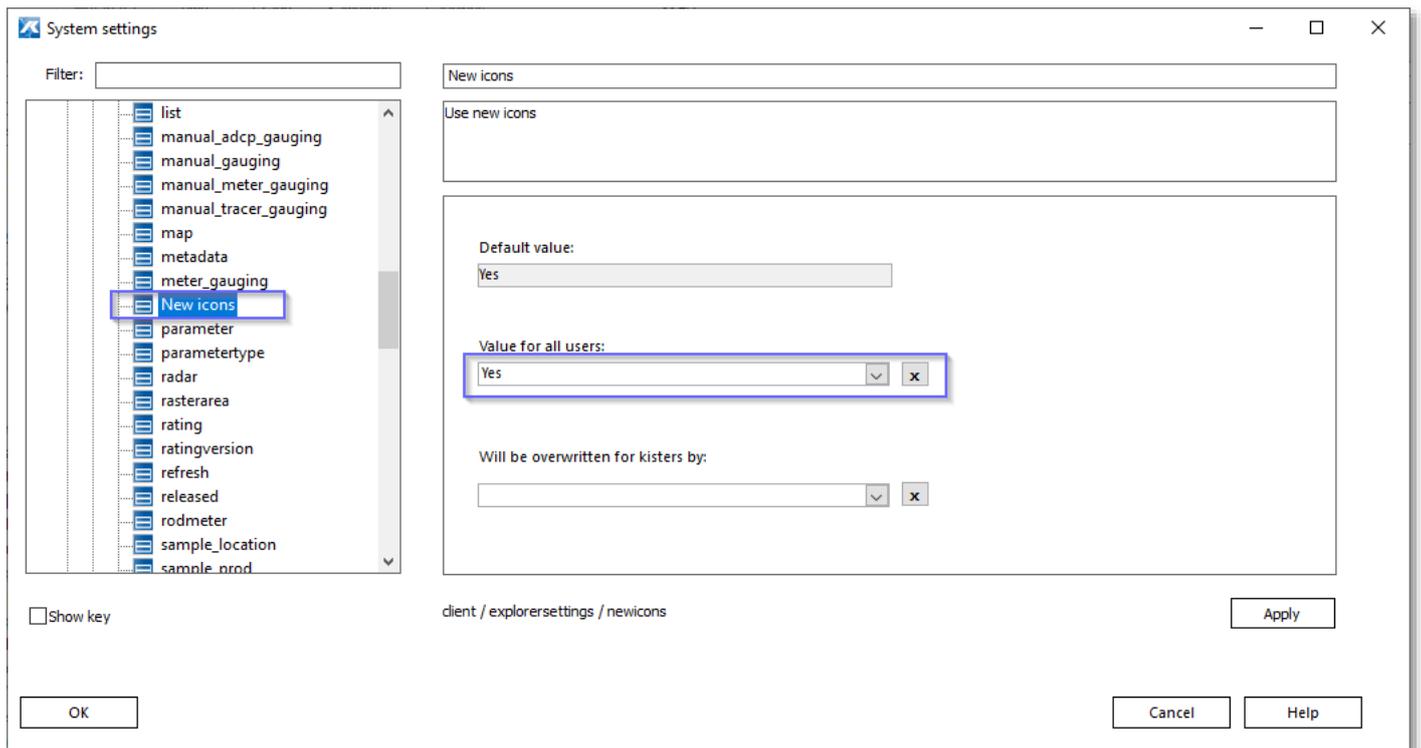
The new toolbar with its snazzy new buttons can be activated via Tools > Settings >>> System settings..., then navigating down through System > Application settings >>> Desktop (or newDesktop):



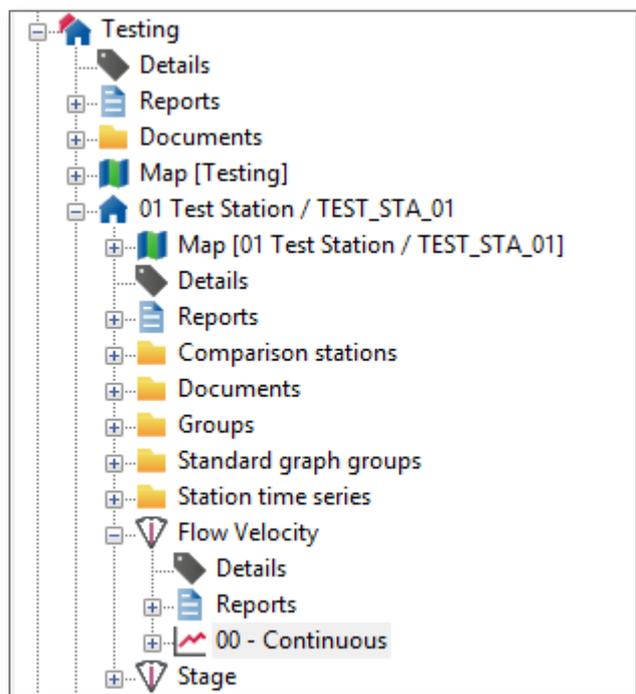
A restart of the application will then update the user/application .ini files and reflect the change:



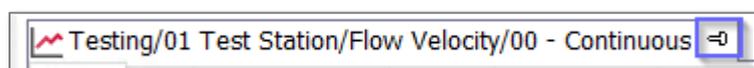
Together with the new toolbar are new icons for objects within WISKI, released prior to the new toolbar last year. These can be activated in the system settings via Local > Explorer settings > Icons >>> New icons:



As before, follow this up with a an application restart to see the new icon designs:



Another fresh GUI feature is the ability to keep tabs with one click, as opposed to doing so via the context menu:

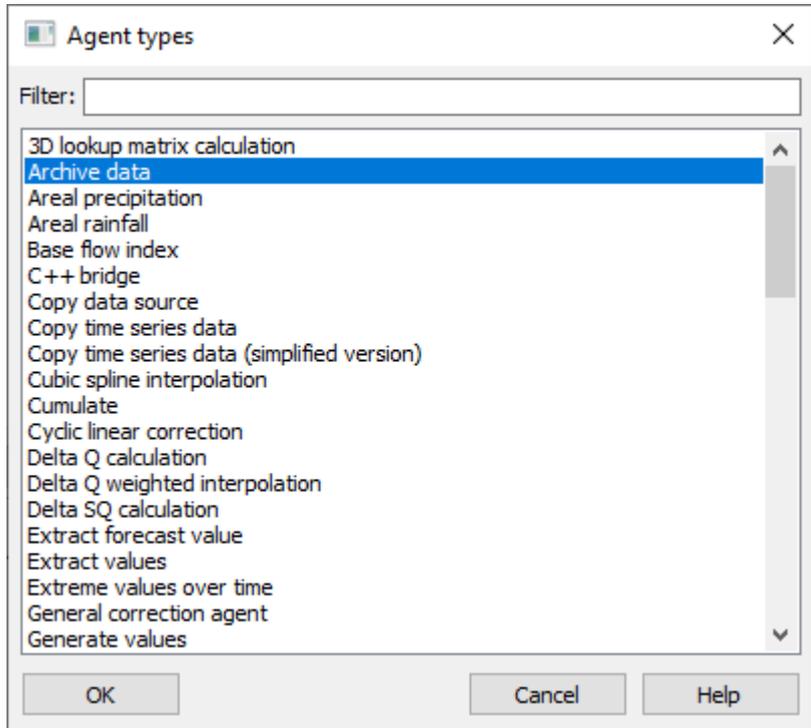


Together with these aesthetic and functional enhancements, development has continued to make the application more accessible to those with visual impairments, resulting in increased navigational capacity without the use of the mouse, and increased integration with screen readers.

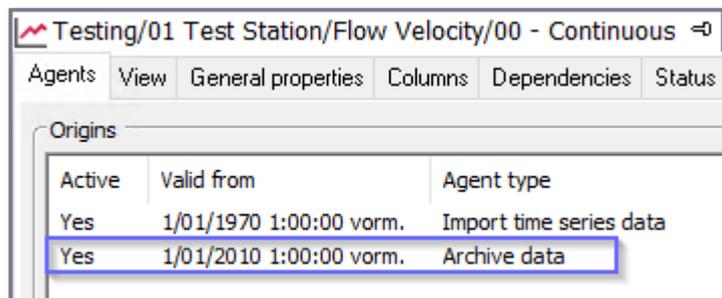
Archive Data Agent

A new agent has been developed which offers an alternative to the release framework to lock time series and prevent them from being edited. This is particularly useful for time series used as reference for other time series which must not be accidentally edited, whether that be manually or through data import.

To apply an archive data agent, from the Agents tab within the time series explorer window, right-click and select Add, or double click within the pane. Select the agent:



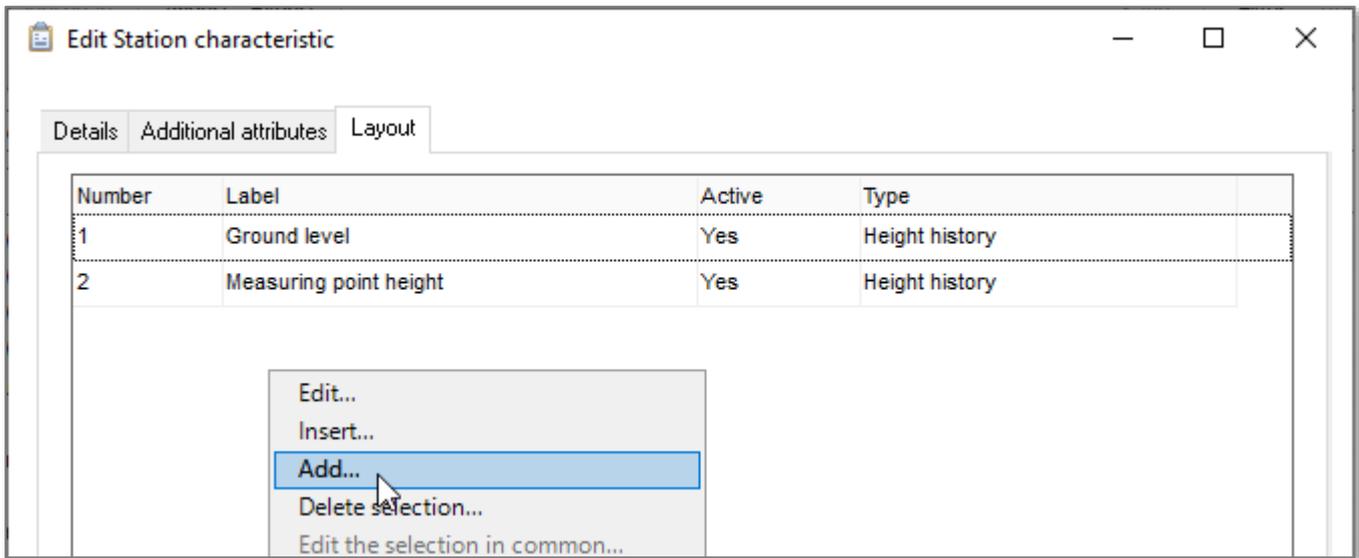
You will then be presented with a simple interface prompting entry of a *Valid from* date, then the setup is complete:



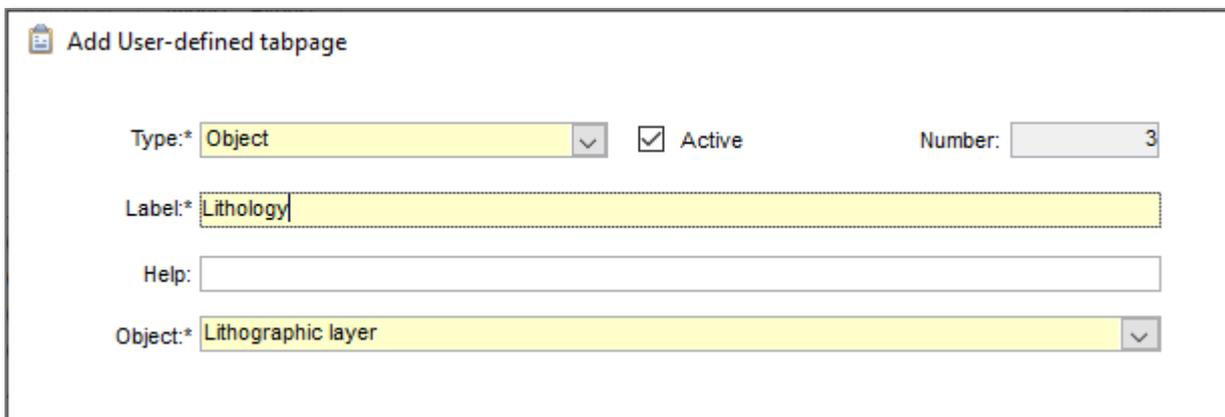
Groundwater Drilling Profiles

Initially introduced in 7.4.11 SR3, this new feature is set up under the Groundwater station characteristic, and is designed to integrate geological/stratigraphic data with existing hydrometric data to provide important contextual information about your station's hydrogeological characteristics. By bringing geological data into the fold, this will enhance the transfer of information between experts in your organisation, improve reporting capabilities and offer a more holistic approach to groundwater data management.

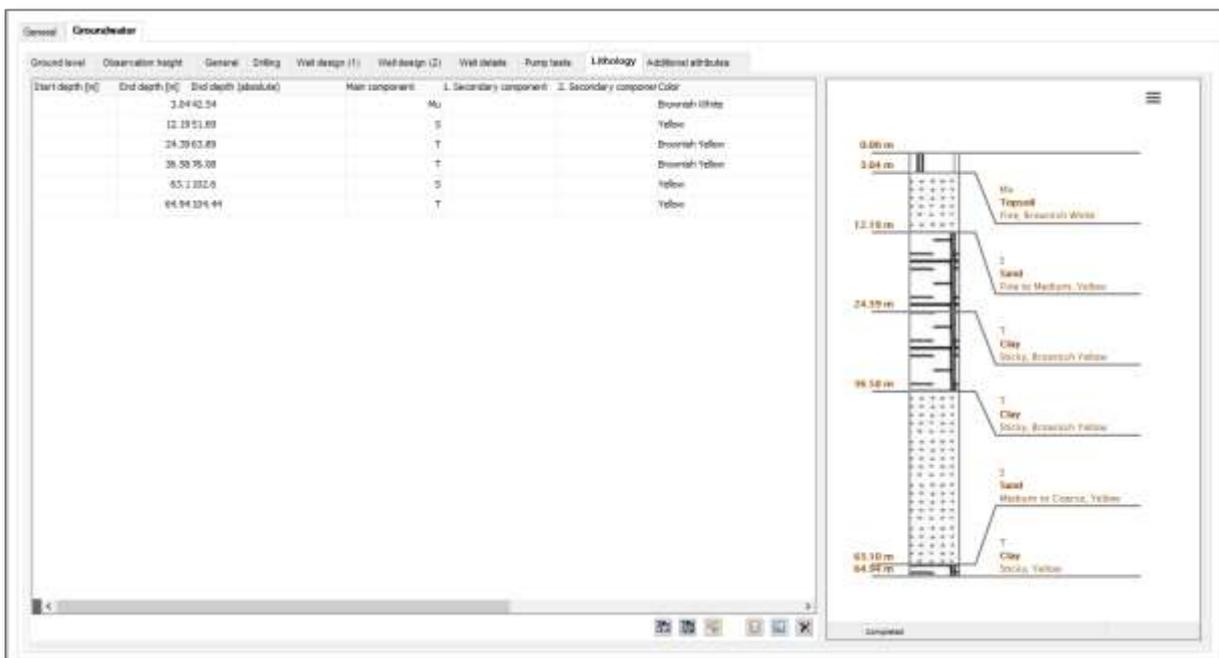
To activate the Lithographic layer object, navigate through System metadata > Management > Types >>> Station characteristic, and open your Groundwater characteristic. Move to the Layout tab, right-click and select Add...:



Select Type = Object, and from the Object field select Lithographic layer, and then apply the tab Label of your choice:



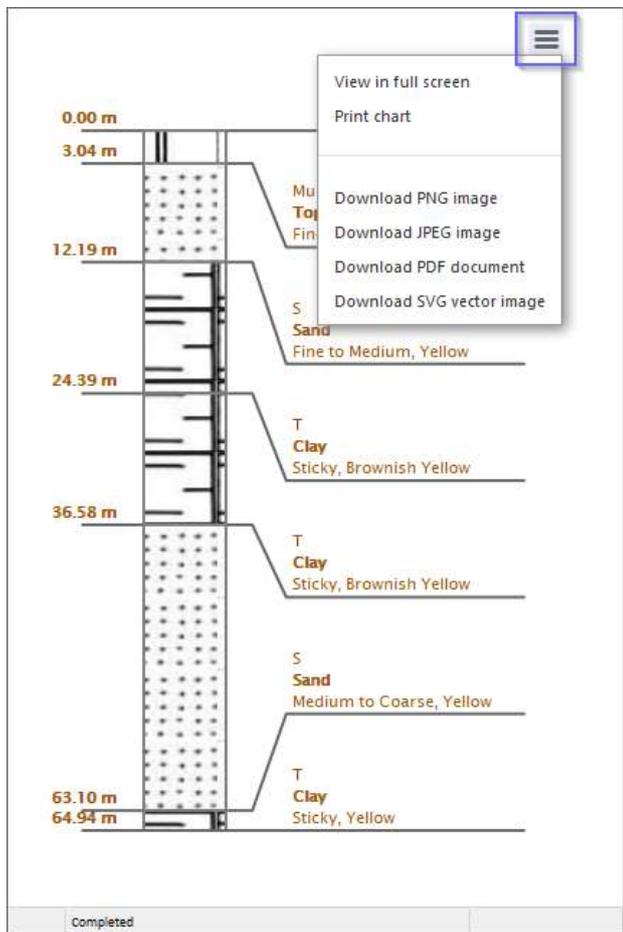
Once set up, shifting to your lithology profile tab from the Groundwater characteristic will give you a summary of each layer present at the station, as well as a visual representation in the right pane. Lithographic layers are also available for import/export as csv files, including all site/station information and other variables as shown below.



To build a profile manually, start by right-clicking inside the Lithology pane and selecting Add...:

The qualitative information that is contained within the lithographic layer is primarily defined in keylists and strings, meaning large lists of stratigraphic/soil data can be imported swiftly for access when manually adding or editing your layers. The naming, symbology and colour of lithographic components are fully customisable to your preference, or alternatively can be derived from national standards.

The soil profile chart can also be saved and exported for reference or reporting purposes via the context menu within the chart itself:



As always, get in touch with us if you're interested in learning more about this feature, or seek information on its setup in your system.

Keycloak – Single Sign on and LDAP/Active Directory support in WISKI

Many of our clients are using Active Directory for access and identity management within their organisation. However, so far, most WISKI installations are not utilising existing central user management infrastructure, and the users are kept separate with their own login credentials governed within WISKI user management. A few of our bigger clients (for example the Bureau of Meteorology) used in the past our native adapter for WISKI to LDAP to integrate Active Directory users. However this native java adapter did not support the standards in authentication such as OpenID Connect, OAuth 2.0, and SAML (Security Assertion Markup Language) which today often does not meet the security standards of state and federal agencies.

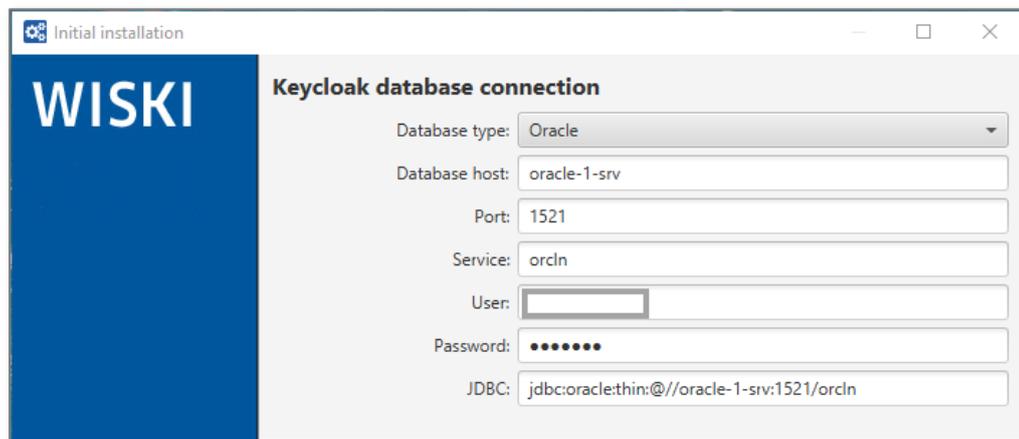
Since the WISKI 7.4.5 stream (most of our clients are using 7.4.7 or higher versions), the OpenID Connect standard was introduced using the internal MitreID provider (part of own WISKI TCA database) which was later extended by integrating the external services of keycloak based on OAuth2 and SAML. These keycloak services allow easy connectivity to LDAP (and secure LDAP) and to Active Directory.

In this context, it is important to know that keycloak is used in WISKI as the Service Provider for Identity Management Providers (Systems) such as Active Directory and not as Identity Provider itself. In the WISKI universe however, keycloak serves as central single sign-on for WISKI clients and the Water Portal applications.

Keycloak is now part of every WISKI distribution and can be deployed using the WISKI Server Manager (WSM). This can be done as initial install or update, as well as installed to existing systems via the Package Management wizard. Once the keycloak zip package is selected the WISKI Server Manager will show the following two additional GUI options:

- Database connection to keycloak schema

Keycloak requires an additional DB schema which needs to be created on the database host. This can be for example as an additional user in the same DB instance (see Oracle example below).

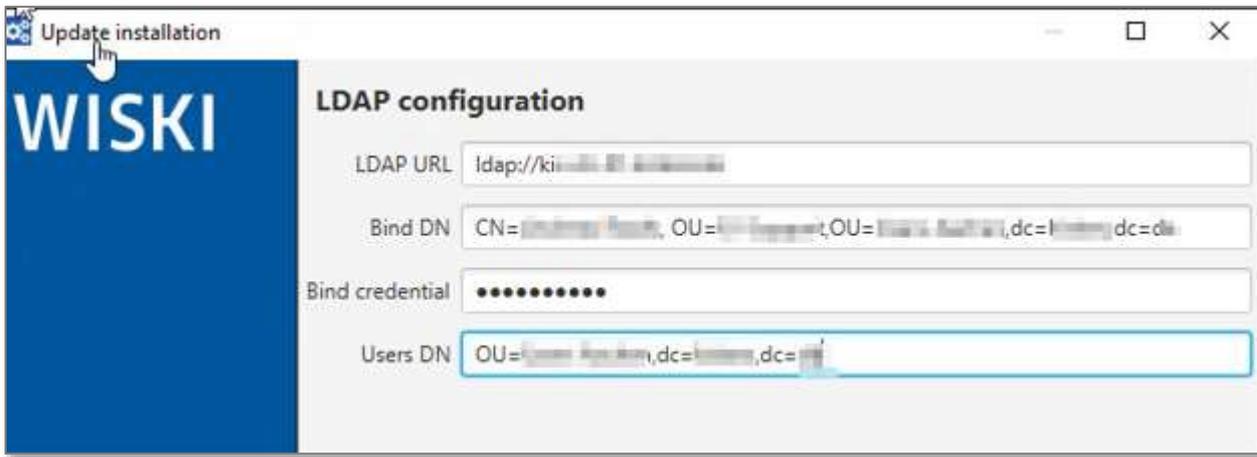


The screenshot shows a window titled "Initial installation" with the WISKI logo on the left. The main area is titled "Keycloak database connection" and contains the following fields:

- Database type: Oracle (dropdown menu)
- Database host: oracle-1-srv
- Port: 1521
- Service: orcln
- User: (empty text field)
- Password: (masked with dots)
- JDBC: jdbc:oracle:thin:@//oracle-1-srv:1521/orcln

- LDAP configuration dialog

The LDAP configuration is customer specific depending on your Active Directory set-up. This configuration needs to be supplied by your Active Directory and LDAP specialist and can be tested in the keycloak GUI (see later below).

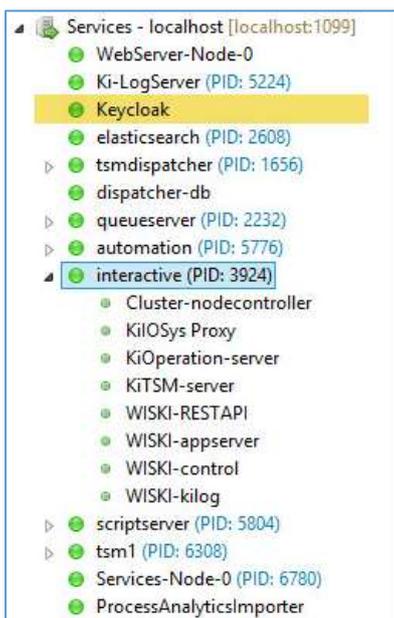


A quick overview of the LDAP elements with description is as follows:

- LDAP: The URL of the machine where the LDAP server is running.
- Bind credentials: The password of the LDAP user.
- Bind DN: The distinguished name of the connection parameters to LDAP including the necessary rights, consisting of common name (CN), organisation unit (OU) and domain controller (DC).
- Users DN: Full Distinguished Name of LDAP tree where your users are. This DN is the parent of LDAP users. It could be for example 'ou=users,dc=example,dc=com', assuming that your typical user will have DN like 'uid=john,ou=users,dc=example,dc=c [...]'; the distinguished name of the current user on LDAP; this indicates how the users are organised on the LDAP server, with the organisation units (OU) and the domain controller(s) (DC).

The WSM keycloak installation routines will install and configure the connection between Active Directory, LDAP, keycloak and WISKI. This includes the set-up of a user federation in keycloak based on the LDAP user, and the client with protocol mapper, which makes the user known to the WISKI dispatcher. The complete authentication uses identity tokens and user credentials are never passed between the applications. The users are created in WISKI by first log-in with a guest role.

Keycloak is part of the WISKI jproccontroller and automatically started (see overview of WISKI services below).



The keycloak service can also be accessed over <https://localhost:7415/auth> (or port 7414 for http). Important is that port 7414 and 7415 allow a connection from the browser to localhost. The administration console of keycloak is configured with a standard admin user (keycloak).

Keycloak can be further configured over the interface. Normally, the keycloak interface is used to check that all components (federation, mapper and clients) are correctly configured as well as for testing and adjusting the connections.

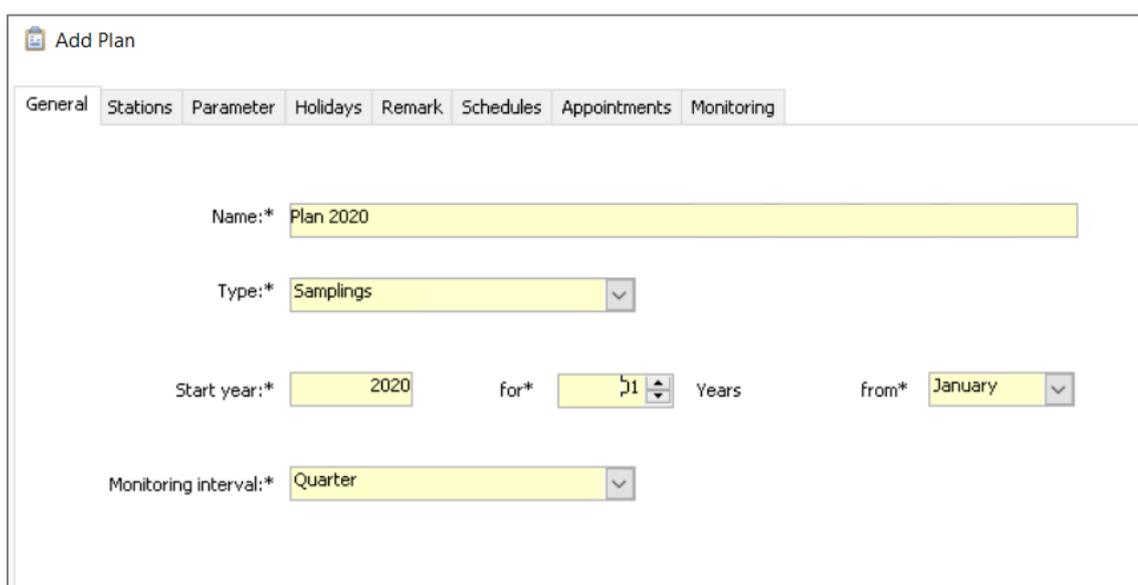
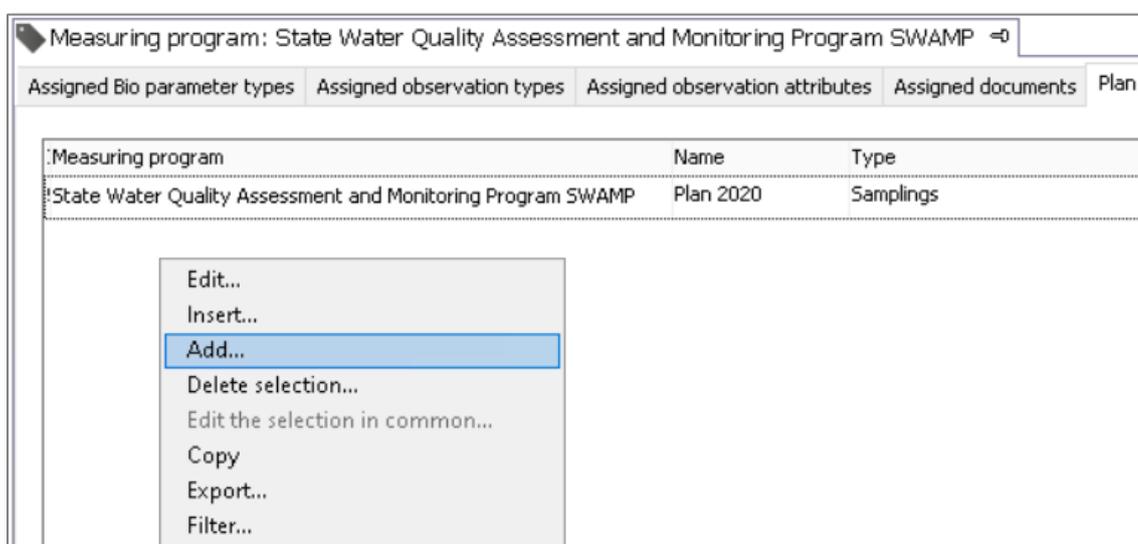
Please let us know when you plan to use keycloak to connect to your Active Directory so we can help to organise the required information before installation of the package.

KiWQM – Sample Planning, Scheduling & Monitoring

Since 7.4.11 SR5 KiWQM now features a fully integrated Sample planning, scheduling and monitoring tool. This feature allows you to: plan future sampling events and testing regimes (on a daily, bi-weekly, weekly, monthly, bi-monthly, yearly basis as well as ad hoc samples), prepare and distribute electronic pre-populated data entry forms as well as monitor your sampling plans for compliance. Below is a brief overview of the new capabilities. Please contact us if you would like further information or a demonstration.

Creation of a Sampling Plan

Select the measuring program and goto the new Plan tab and add your new Plan, specify the name, start year and monitoring interval:



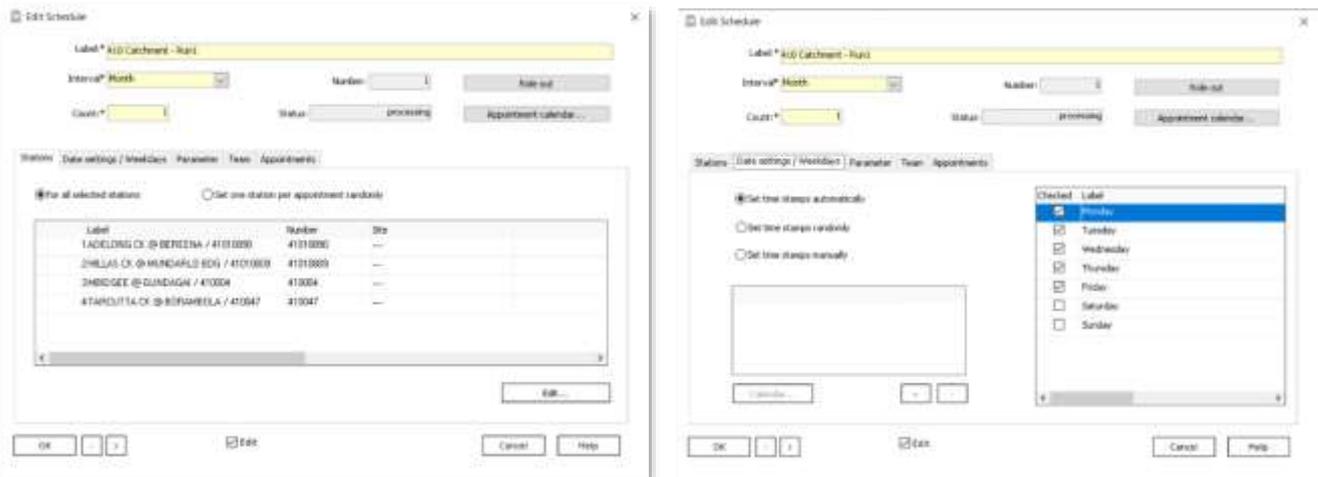
Assign stations to be visited and parameters to be measured. These are restricted to those assigned to the measuring program.

Configure holidays:

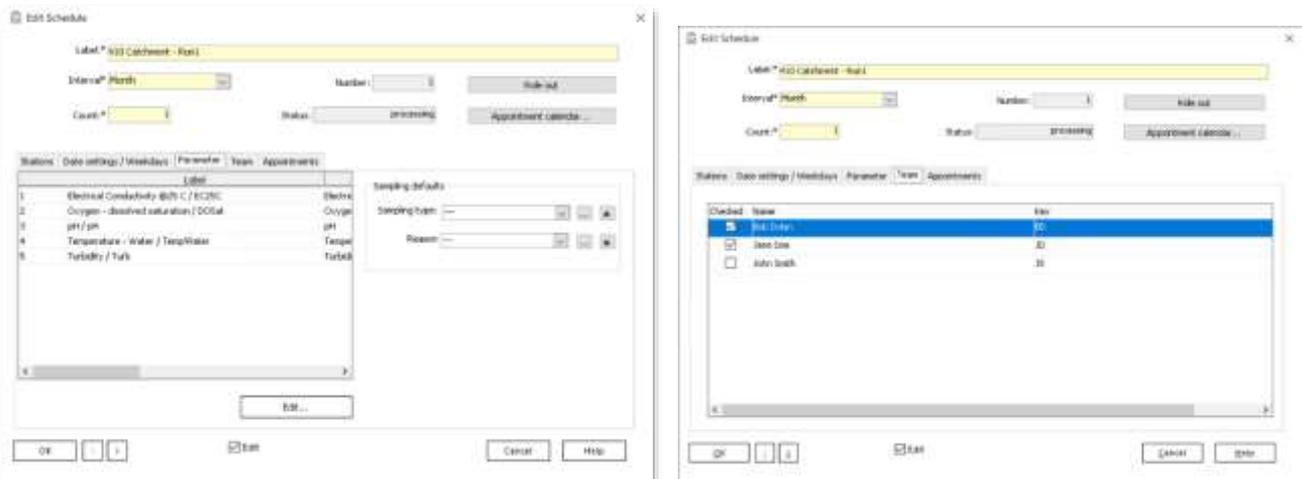


Creation of a Sampling Schedule

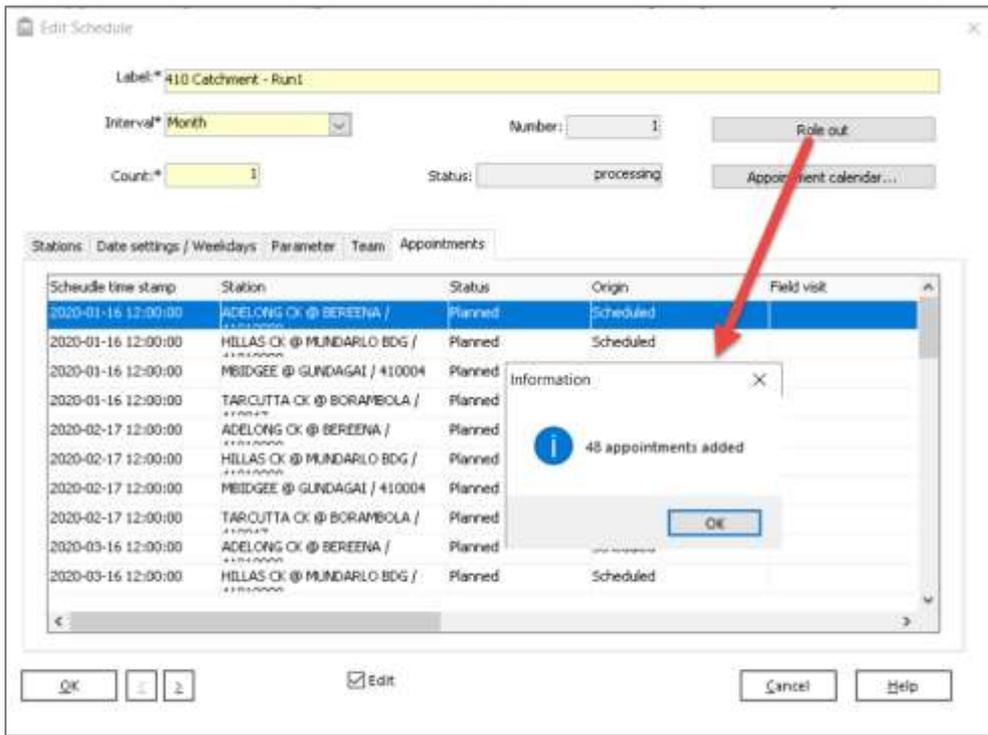
To create a sampling schedule, go to the Schedule tab and add a new schedule, select the Interval type (daily, weekly, monthly, quarterly, annually) and the number of samples within each interval. Stations can be manually or randomly assigned to schedule. The appointment dates can be assigned manually, automatically (applies weekday & holiday restrictions) or randomly. The status indicates if the appointments are not yet created (in progress) or if they have already been applied (done).



Select the parameters to be measured and select the team members that will be assigned to the samplings:

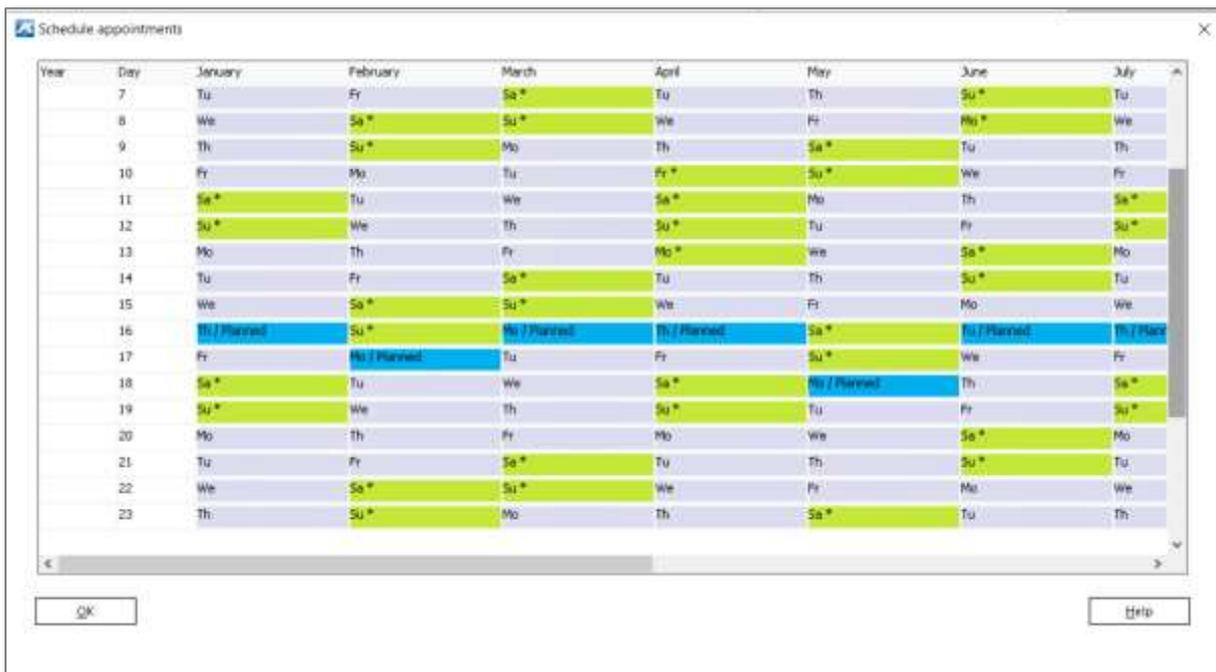


Roll out your schedule to create appointments:

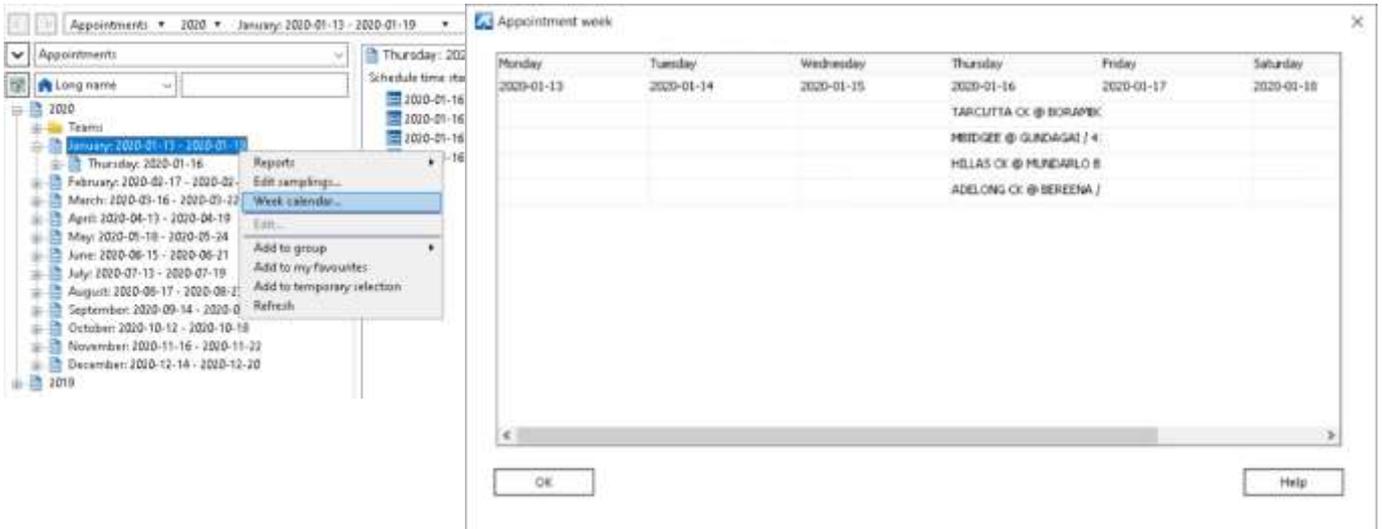


Viewing Appointments

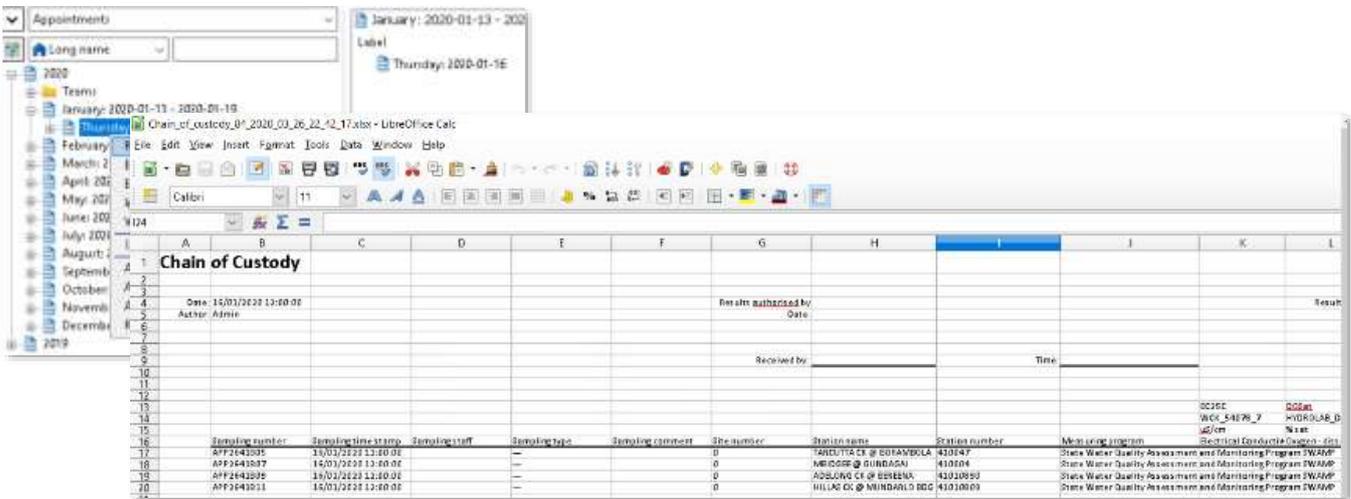
Show your appointments for the rolled out schedule in a calendar from the Appointments calendar button on the Schedule form:



You can also check appointments from the new Appointments explorer view. Here you can also see a calendar view of appointments for the week:

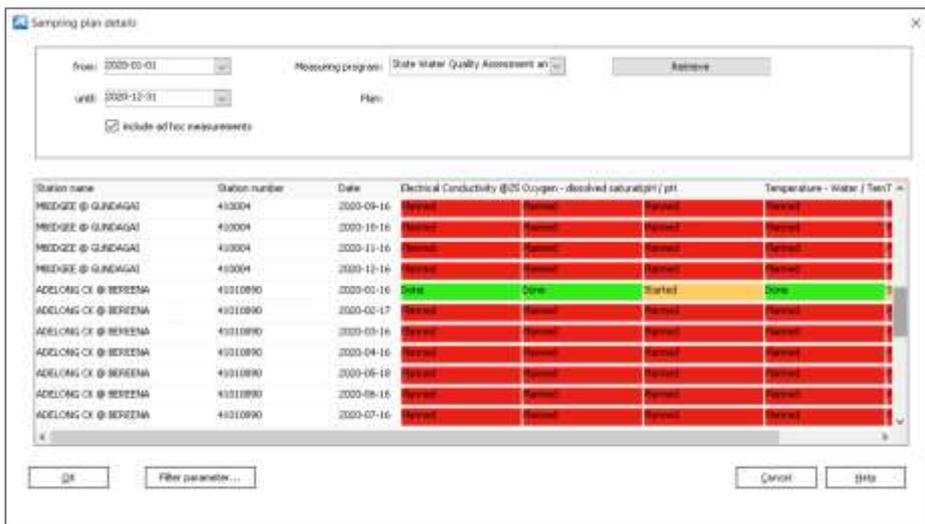


From this view you can also initialise the samplings and create a customised report that can be used as a field or lab work order:



Monitoring Appointments

The progress of scheduled and adhoc samplings can be monitored via the Appointments view by right clicking on the year and selecting appointment details.



The evaluation tab of the plan also provides overviews of the yearly plan, broken down into the Monitoring intervals specified by the plan.

The 'Evaluation' dialog box displays a table with the following columns: Name, Number, Records, Appointment 1, Appointment 2, and Appointment 3. The 'Records' column shows '1 / 3' for the first row. The 'Appointment' columns contain dates: 'Planned / 2020-01-16', 'Planned / 2020-02-17', and 'Planned / 2020-03-18'.

| Name | Number | Records | Appointment 1 | Appointment 2 | Appointment 3 |
|-----------------------------------|--------|---------|----------------------|----------------------|----------------------|
| ABEROROMBIE @ CAMP AP41210123 | | | | | |
| ADELONG CK @ BEREENA 41010890 | | 1 / 3 | Planned / 2020-01-16 | Planned / 2020-02-17 | Planned / 2020-03-18 |
| ADJUNGELLY @ DRIBALAR410038 | | | | | |
| Avon River @ Deards Lane 20810001 | | | | | |
| Avon River @ Pipers Lane20610080 | | | | | |
| BARNARD @ MACKAY 208011 | | | | | |
| BARWON @ BREWARRINA 422002 | | | | | |
| BARWON @ COLLARENEBR422003 | | | | | |
| BARWON @ DANGAR BOG 422001 | | | | | |
| BELL @ NEWIRA 421018 | | | | | |
| BELLINGER @ BELLINGEN 20510051 | | | | | |
| BELUBULA @ CANOWINDR412009 | | | | | |

The 'Sampling plan details' dialog box displays a table with the following columns: Station, Electrical Conductivity @25 Hardness as CaCO3 (calc)pH / pH, and Solids - total dissolved (calc Solids - total disso. The data is as follows:

| Station | Electrical Conductivity @25 Hardness as CaCO3 (calc)pH / pH | Solids - total dissolved (calc Solids - total disso |
|------------------------------|---|---|
| TARCUTTA CK @ BORAMBK | 0 / 0 | 0 / 0 |
| MBIDGEE @ GUNDAGAI / 4 | 0 / 0 | 0 / 0 |
| ADELONG CK @ BEREENA / 1 / 3 | 0 / 0 | 1 / 0 |
| HILLAS CK @ MUNDARLO B | 0 / 0 | 0 / 0 |

WISKI Support Email and Help Desk

Contacts for the WISKI team at KISTERS in Australia:

Vicky, Chris, Markus, David and Callum (web developments) offer specialised support for the KISTERS products WISKI, KiWQM, KiEco, KiDSM, KiALM, Water Portal, WDO and KiWIS in Australia and New Zealand.

The phone number for support is +61 2 6154-5200, and the email address is wiski-support@kisters.com.au.

If you are engaging in a particular dialog with Chris, Vicky, Markus, Callum or David, please cc the support box so a central register of issues can be maintained.

The latest WISKI releases can be found on our download portal at <http://kisters.com.au/downloadswiski.html>, or can be accessed by navigating through to the support page from <http://kisters.com.au>.

To acquire a username and password to access the download portal please contact the KISTERS support team over the phone at (02) 6154 5200 or email at Wiski-Support@kisters.com.au

PreciBal: Weighing Precipitation Gauge

With the PreciBal, HyQuest Solutions launches a cost-efficient precipitation gauge using the highly accurate weighing principle. Determining mass instead of volume significantly increases the accuracy of the measurement.

Furthermore, due to advanced algorithms, the weighing rain gauge measures the water equivalent of all forms of precipitation: rain, snow, drizzle and hail. With PreciBal, HyQuest Solutions completes its line of state-of-the-art pluviometry products providing the most advanced technology available at a reasonable price.

The PreciBal excels with proven and modern technology, durable material and algorithm-based data processing. Various catch diameters are available making the device compliant with most local regulations.

Benefits:

- High accuracy in all weather conditions
- Measurement of liquid, mixed and solid precipitation quantity and intensity
- Load cell technology, edge computing and low power consumption
- Easy maintenance and extended service intervals
- Various output options: Interference free digital SDI12 or RS485, pulse

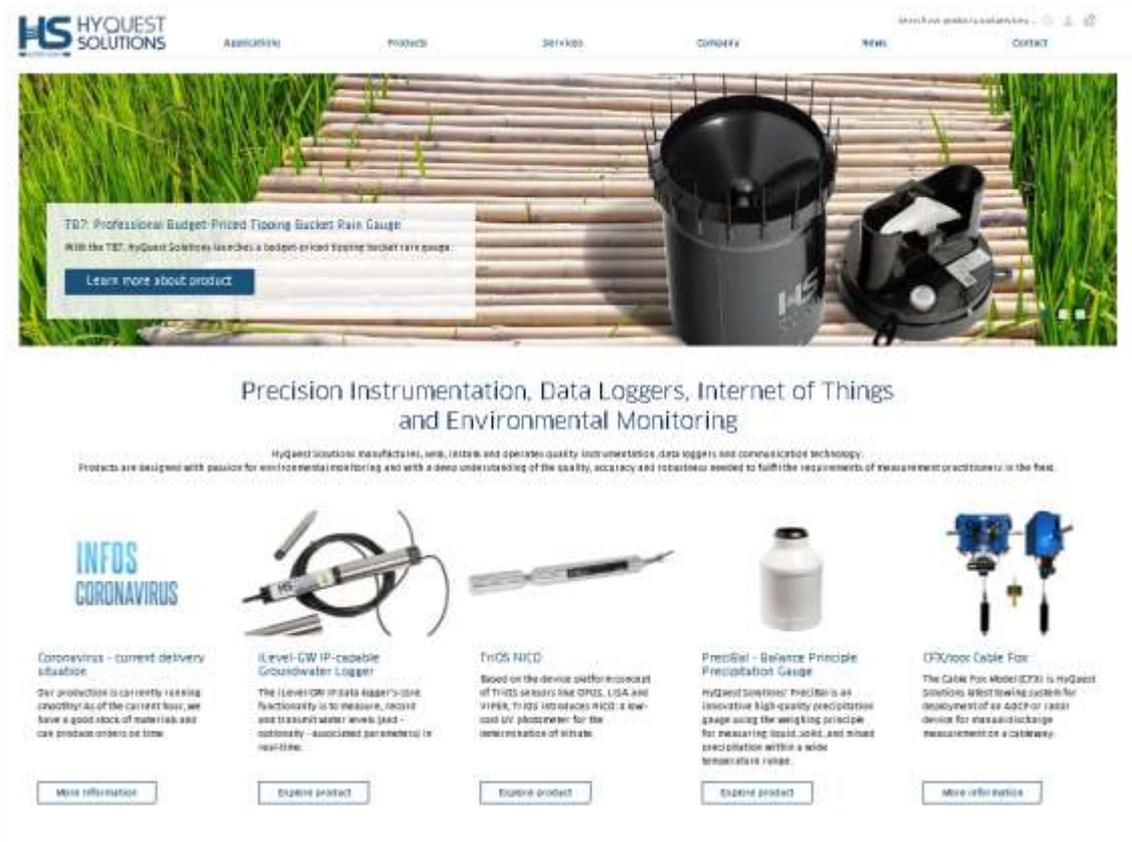
Contact:

sales@hyquestsolutions.com.au



HyQuest Web Site Refresh

HyQuest have refreshed their web site, and you can visit it at www.hyquestsolutions.com for information on the HyQuest range of products and services.



KISTERS Training

It is clear we won't be doing any face to face training for the foreseeable future. However KISTERS Australia has acquired a number of technologies which will enable us to deliver training via the internet, including GoToMeeting, GoToWebinar, Zoom and Skype. Once we gain a better understanding of what's possible we will start offering training courses over the internet. Stay tuned for more information, and please contact us to register your interest in potential future web-based training.

News from KISTERS Overseas

You can keep up with news from other KISTERS offices via the following links:

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

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

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

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, General
Manager
GEOSPHERA Hi-Tech Supplies
85 Via Panoramica
Ercolano (NA) 80056 - ITALY
Web: <http://www.geosphera.com>
Phone: +39 081 777 9541,
Fax: +39 081 739 2596
Mobile: +39 335 601 3998
Email: geosphera@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