ibaPDA-Data-Store-MySQL

Data streaming into MySQL/MariaDB

Manual

Issue 1.0

Measurement Systems for Industry and Energy

www.iba-ag.com

Manufacturer

iba AG

Koenigswarterstr. 44

90762 Fuerth

Germany

Contacts

Main office	+49 911 97282-0
Fax	+49 911 97282-33
Support	+49 911 97282-14
Engineering	+49 911 97282-13
E-mail	iba@iba-ag.com
Web	www.iba-ag.com

The content of this publication has been checked for compliance with the described hardware and software. Nevertheless, discrepancies cannot be ruled out, and we do not provide guarantee for complete conformity. However, the information furnished in this publication is updated regularly. Required corrections are contained in the following regulations or can be downloaded on the Internet.

The current version is available for download on our web site www.iba-ag.com.

Version	Date	Revision - Chapter / Pagee	Author	Version SW
1.0	10-2020	First issue	st	7.2.2

Windows® and SQL Server® are brands and registered trademarks of Microsoft Corporation. Oracle® und MySQL® are registered trademarks of Oracle and/or its partners. MariaDB® is a registered trademark of MariaDB.

Other product and company names mentioned in this manual can be labels or registered trademarks of the corresponding owners.

Content

1 About this manual

This documentation describes the function and application of the data store *ibaPDA-Data-Store-MySQL*.

This documentation is a supplement to the *ibaPDA* manual. Information about all the other characteristics and functions of *ibaPDA* may be found in the *ibaPDA* manual or in the online help.

You will find basic information about data storage in *ibaPDA* in the *ibaPDA* manual part 5.

1.1 Target group and previous knowledge

This documentation addresses qualified professionals, who are familiar with handling electrical and electronic modules as well as communication and measurement technology. A person is regarded as professional if he/she is capable of assessing the work assigned to him/her and recognizing possible risks on the basis of his/her specialist training, knowledge and experience and knowledge of the standard regulations.

This documentation in particular addresses persons, who are concerned with the configuration, test, commissioning or maintenance of the supported database, cloud or cluster storage technology. For the handling of *ibaPDA-Data-Store-MySQL* the following basic knowledge is required and/or useful:

- Windows operating system

- Basic knowledge of *ibaPDA*

- Basic knowledge of databases, cloud or cluster storage technology

1.2   Notations

In this manual, the following notations are used:

Action	Notation
Menu command	Menu *Logic diagram*
Calling the menu command	*Step 1 – Step 2 – Step 3 – Step x* Example: Select the menu *Logic diagram - Add - New function block*.
Keys	<Key name> Example: <Alt>; <F1>
Press the keys simultaneously	<Key name> + <Key name> Example: <Alt> + <Ctrl>
Buttons	<Key name> Example: <OK>; <Cancel>
File names, paths	"Filename", "Path" Example: "Test.doc"

1.3   Used symbols

If safety instructions or other notes are used in this manual, they mean:

Danger!

The non-observance of this safety information may result in an imminent risk of death or severe injury:

- ■ Observe the specified measures.

Warning!

The non-observance of this safety information may result in a potential risk of death or severe injury!

- ■ Observe the specified measures.

Caution!

The non-observance of this safety information may result in a potential risk of injury or material damage!

- ■ Observe the specified measures

Note

A note specifies special requirements or actions to be observed.

Tip

Tip or example as a helpful note or insider tip to make the work a little bit easier.

Other documentation

Reference to additional documentation or further reading.

2 Introduction

Different types of data stores are available in *ibaPDA* for different purposes and methods of data storage. Depending on the licenses registered in the dongle, different types of data stores are available for configuration in the dialog.

This documentation describes the "DB/Cloud timebased data store" type of recording. This recording type writes timebased data to a database, such as SAP HANA, SQL Server, Oracle, MySQL or PostgreSQL.

No measurement files are generated, but the data is written to a table in the database. You define the table structure using the storage profile, see chapter ↗ *Storage profiles*, page 12

Chapter ↗ *Signal selection*, page 15 describes the selection of the signals that are to be recorded.

The data can be continuously recorded or recorded by trigger, see chapter ↗ *Trigger mode*, page 16.

2.1 System requirements

The following system requirements are necessary when using data storage in a MySQL/MariaDB database:

- *ibaPDA* v7.2.0 or higher

- License for *ibaPDA-Data-Store-MySQL*

The licenses are staggered according to the number of signals that should be written in a database. The number of used data stores is unlimited.

Order no.	Product name	Description
30.671050	ibaPDA-Data-Store-MySQL-64	Data streaming into MySQL DB, max. 64 signals
30.671051	ibaPDA-Data-Store-MySQL-256	Data streaming into MySQL DB, max. 256 signals
30.671052	ibaPDA-Data-Store-MySQL-1024	Data streaming into MySQL DB, max. 1024 signals
30.671055	upgrade-ibaPDA-Data-Store-MySQL-64 to 256	License for extension from 64 to 256 signals
30.671056	upgrade-ibaPDA-Data-Store-MySQL-256 to 1024	License for extension from 256 to 1024 signals

Table 1: Available licenses for the data storage in MySQL/MariaDB databases

3   Data store configuration

3.1   Add a data store

The dialog for data storage configuration can be opened in the *Configure – Data storage* main menu or by clicking on the button in the main toolbar.

In order to add a new data store, click on the blue link *Add data store* in the tree structure. You can also right-click on the data store node in the tree structure and choose *Add data store* from the context menu.

Select *DB/Cloud timebased data store* for the recording of timebased data into a database or cloud.
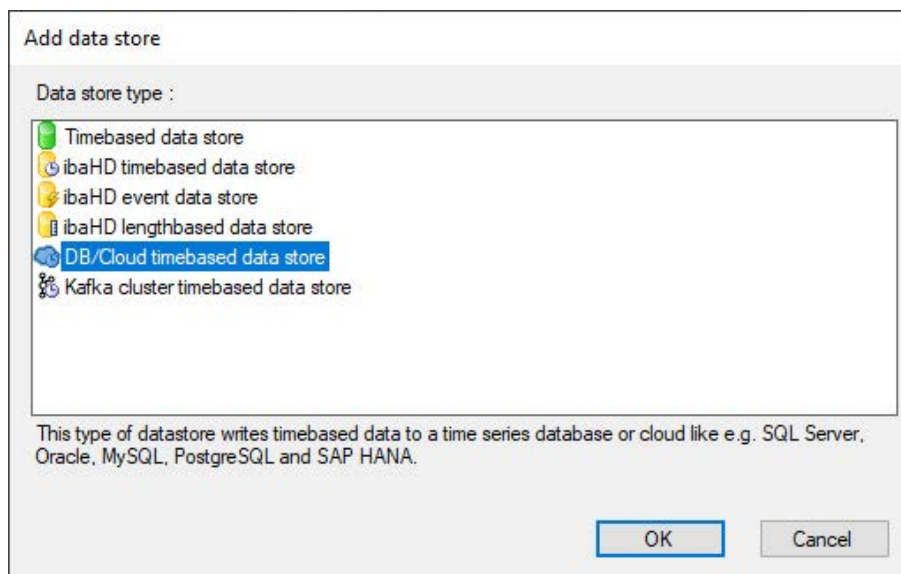

Fig. 1: Add a data store

3.2   Database type MySQL, MariaDB

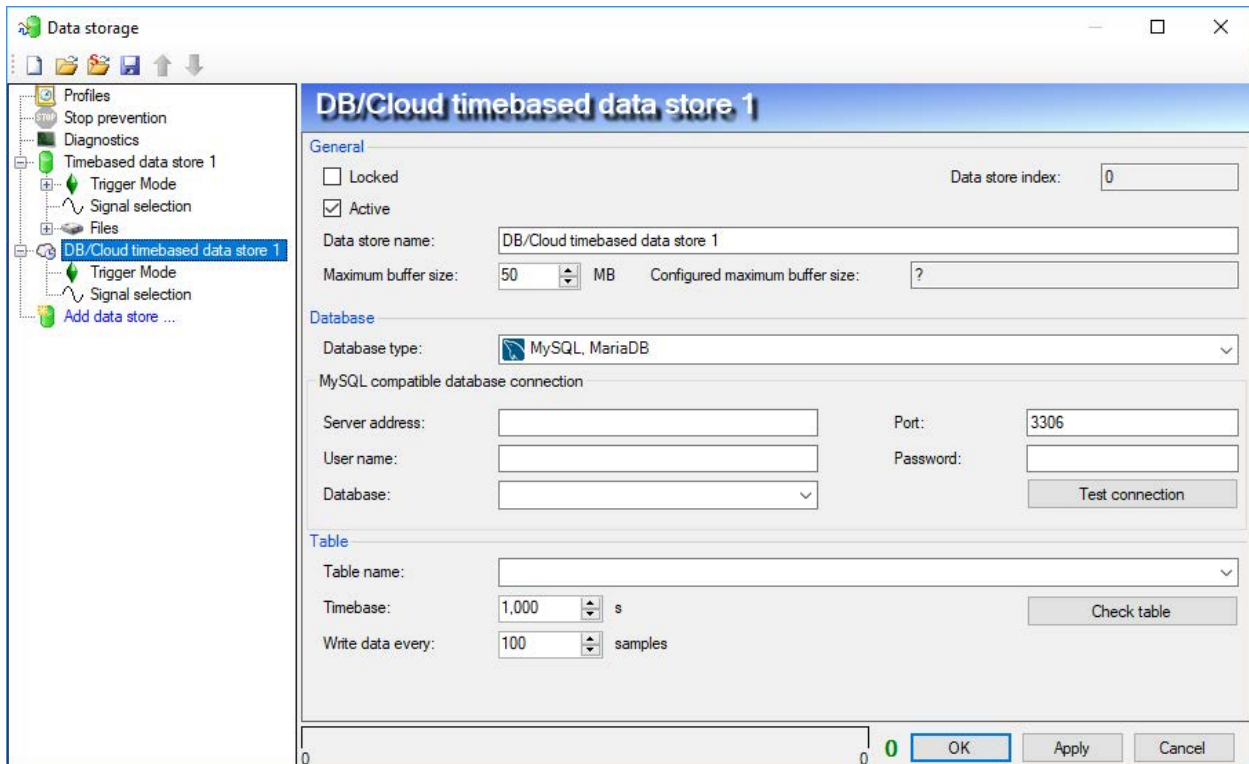First select the database type *MySQL, MariaDB,* and then configure the other necessary settings.

Fig. 2: Configuration database type MySQL, MariaDB

General

Locked
A data store can be locked in order to prevent an accidental or unauthorized change of settings.

Active
A data store must be enabled in order to work. However, you can configure various data stores and disable data stores that are not required.

Data store name
You can enter a name for the data store here.

Maximum buffer size
The maximum buffer size determines how much data is buffered in *ibaPDA* in the event of a temporary connection loss with the database. The buffer is kept in the RAM. When the acquisition is restarted, the buffer memory is deleted.

Configured maximum buffer size
This field indicates the time span that can be temporarily stored in the buffer with the configured settings. Specification in d.hh:mm:ss.

Database

Database type

Select your database type from the drop-down menu, here *MySQL, MariaDB*.

Server address

Name or IP address of the server to be connected to.

Port

TCP port through which the MySQL or a compatible server communicates.

User name/password

Enter user name and password as login. This option must be configured accordingly in the database.

Database

Name of the database to which a connection is to be established. If the server address, port, user name and password are filled in correctly and if the database server is online, the available databases are displayed in the drop-down list.

<Test connection>

Use the <Test connection> button to test the connection to the database. When testing the connection *ibaPDA* tries to retrieve the current list of tables and fills in the drop-down list *Table name* in the *Table* area. If the connection is successful, the version of the database is also displayed.

Table

Table name

Select an existing table from the drop-down menu or enter the name of a new table. The table name can be a full table name consisting of schema name and table name separated by a dot, for example "iba.test". If the table name does not include the schema name then the table will be created in the default schema of the configured user.

Timebase

Enter a timebase for the table. All data in the table are equidistant. All timestamps will be aligned to the timebase.

Write data every x samples

Instead of inserting 1 row at a time into the database it is much more efficient to insert multiple rows at once. This can be controlled via the "Write data every x samples" option.

<Check table>

The <Check table> button can be used to check if the table exists and if all required columns exist and have the correct data type. The required columns depend on the selected signals and their storage profile. Therefore, before checking the table, you should configure the storage profile, see ↗ *Storage profiles*, page 12 and select the signals, see ↗ *Signal selection*, page 15.

If you check the table after configuring the storage profile and selecting the signals, the table will be created automatically in the database by *ibaPDA*.

In case the table already exists in the database and the structure is compliant with the configuration, you will just see a confirmation message.

If a table already exists in the database and the structure is not compliant with the configuration, the following dialog is shown listing the differences:
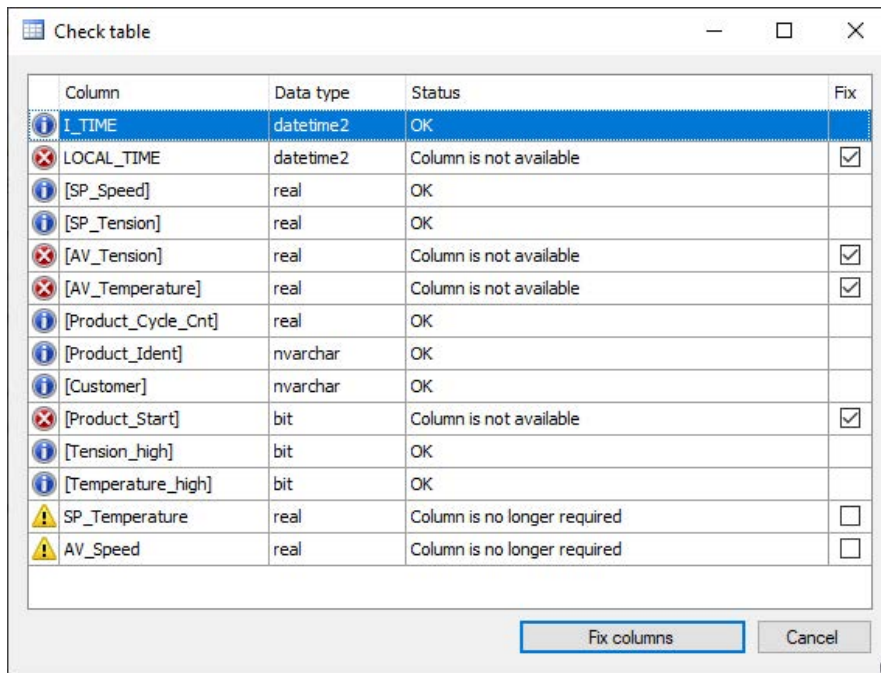

Fig. 3: Check table

Each line corresponds to one column. The first column shows the message type: Info, warning or error. The second column shows the column name, the third column shows the data type and the fourth column shows the current message. In the last column you can select which columns you want to correct. The faulty columns are always selected and cannot be deselected. If you click on <Fix columns> *ibaPDA* tries to correct the selected columns.

In case the changes that have to be done to an existing table according to your configuration are too profound, the existing table has to be deleted and a new one will be created. This will lead to the loss of the data in the existing table.

4 Storage profiles

There is a special storage profile for the *DB/Cloud timebased data store*. In the storage profile you define which value from the signals is written to the table per timebase.

4.1 Add profile

To add a DB profile, select the *Profiles* branch in the tree structure of the data storage configuration dialog. Click the drop-down list icon on the <Add> button in the right pane and select *Time, DB/Cloud* from the drop-down list.
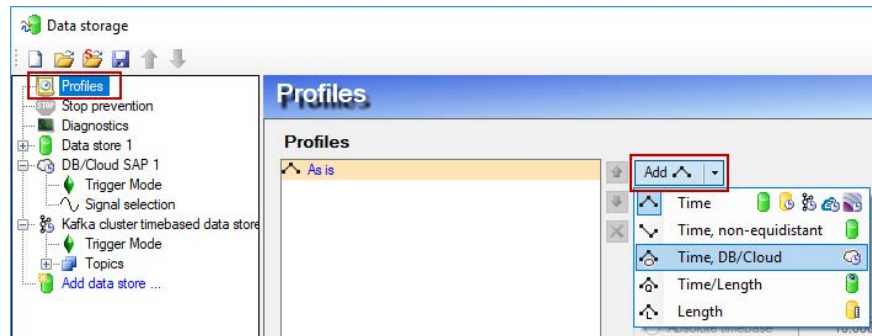

Fig. 4: Add profile
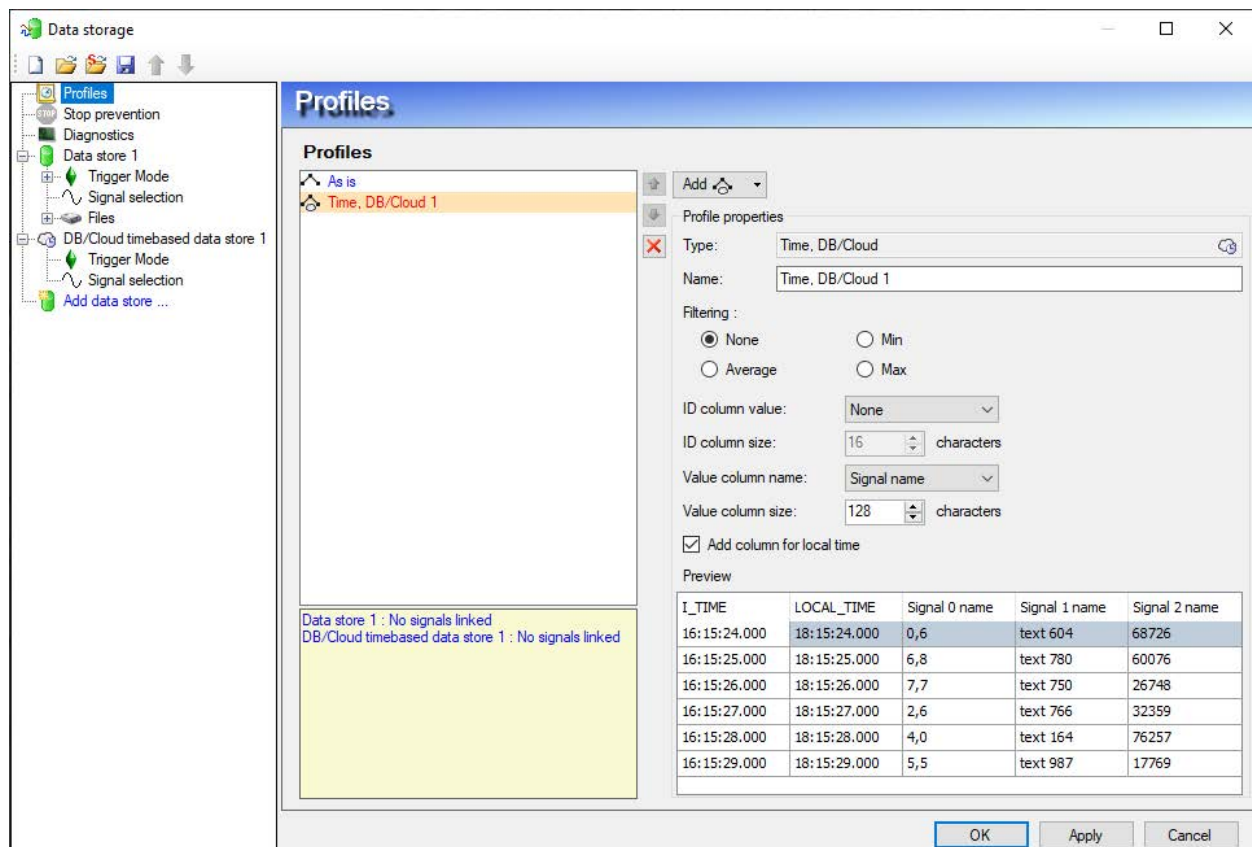
4.2 Profile Time, DB/Cloud

Fig. 5: Profile Time, DB/Cloud

Profile properties

Type
Displays the profile type (information only)

Name
Enter a name for the profile.

Filtering
Select here which value from the signal will be written into the table every timebase. The following values can be selected:

■ *None*: The signal value at the time at which the timebase expires is taken. All other values in the time range are ignored.

■ Min: The smallest signal value within the timebase.

■ Average: The average value of the signal in the timebase.

■ Max: The largest signal value within the timebase.

Configure columns of the table
The table into which data is written using the data store always has a timestamp column with the unchangeable name *I_TIME*. The timestamp is always generated in UTC time.

In addition to the *I_TIME* column, a second time stamp column *LOCAL_TIME* with the local time of the *ibaPDA* server PC can be added optionally. To do this, the option *Add column for local time* must be enabled.

In the table, you can optionally specify an ID column that can be used as a key. The default name is *I_ID.* You can choose whether to use the ID column, and if so, which value should be inserted. The ID column value drop-down menu provides the following options:

■ *None*: No ID column will be used.

■ Fixed: A fixed text will be written into the ID column. This could be used for example when multiple *ibaPDA* systems are writing to the same table as an identification of the *ibaPDA* system. Enter the desired text in the input field to the right.

■ Text signal: The value of a text signal. Select the desired text signal in the selection field to the right.

■ Signal ID: The ID of the signal this profile is applied to. Example:
A_0_1 means analog signal (A) with the signal ID in *ibaPDA* [0:1]
D_0_1 means digital signal (D) with the signal ID in *ibaPDA* [0.1]

■ Signal name: The name of the signal this profile is applied to.

■ Signal comment 1: The first comment of the signal this profile is applied to.

■ Signal comment 2: The second comment of the signal this profile is applied to.

ID column size
Here you can specify a character length for the ID column.

Value column name

The *Value column name* option determines the name of a column in which a signal value is written. You can choose between:

■ Fixed: Fixed column name means, that the data of each signal will be written to a separate row. It is recommended to use the ID column to determine to which signal the data belongs.

■ Signal ID: The colum name is the signal ID. Example:
A_0_1 means analog signal (A) with the signal ID in ibaPDA [0:1].
D_0_1 means digital signal (D) with the signal ID in ibaPDA [0.1].

■ Signal name: The column name is the signal name. Characters that are not allowed in column names in the respective database type are replaced by underscores.

■ Signal comment 1: The column name is derived from the first comment of the signal.

■ Signal comment 2: The column name is derived from the second comment of the signal.

For all settings except "fixed", a row contains the data from multiple signals.

You specify the maximum character length in the *Value column size* field.

The preview shows a general example how a table with the current settings will look like.

Note

Additional information about the storage profiles can be found in the manual *ibaPDA*, part 5.

---

5 Signal selection

To enable signals to be recorded, they must be assigned to a storage profile of type *Time, DB/Cloud.* Select the signal selection node below your *DB/Cloud timebased data store* to open the signal selection dialog.
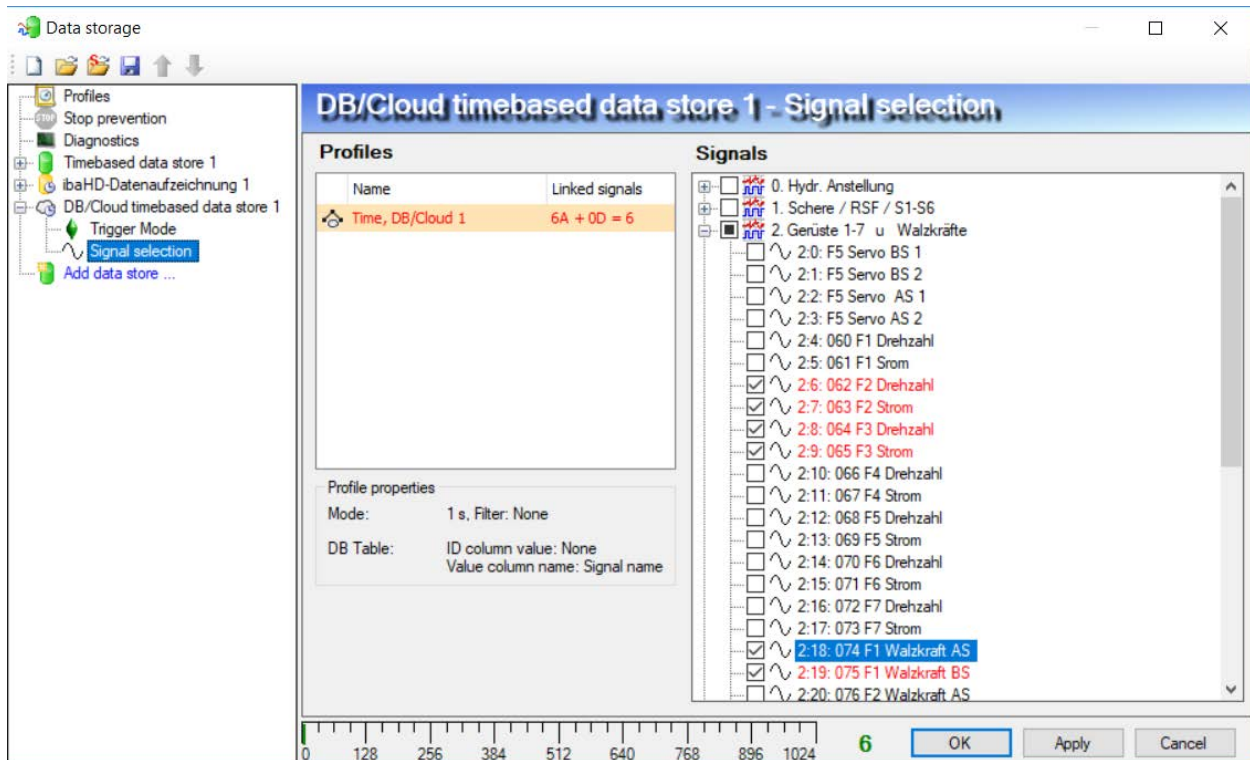

Fig. 6: Signal selection

In the profile list, select the storage profile to which you want to assign certain signals. Set a check mark in the selection fields next to the signals which you would like to assign to this profile.

A signal can only be assigned to one profile per data store.

The *Profile properties* section displays some information about the configured timebase, filtering and column naming of the selected profile.

DB/Cloud data stores are licensed for each database type separately, e. g. SAP HANA, Oracle, SQL Server etc. These licenses are staggered according to the number of signals written to the database. The current number of selected signals in all DB/Cloud data stores of one database type is shown at the bottom of the dialog, similar to the number of configured signals in the I/O Manager.

The licensed number of signals is indicated by the length of the signal strip. In the above example, it is possible to write up to 1,024 signals in several data stores of the same DB type. Currently 6 signals are enabled.

When you have configured all signals you want to write to the database, go back to the main node of the data store. There you can check the table with the <Check table> button.

6   Trigger mode

The description applies to the following types of data stores:

- ▪ Database/cloud timebased
- ▪ Kafka cluster timebased
- ▪ Mindsphere timebased
- ▪ MQTT timebased.

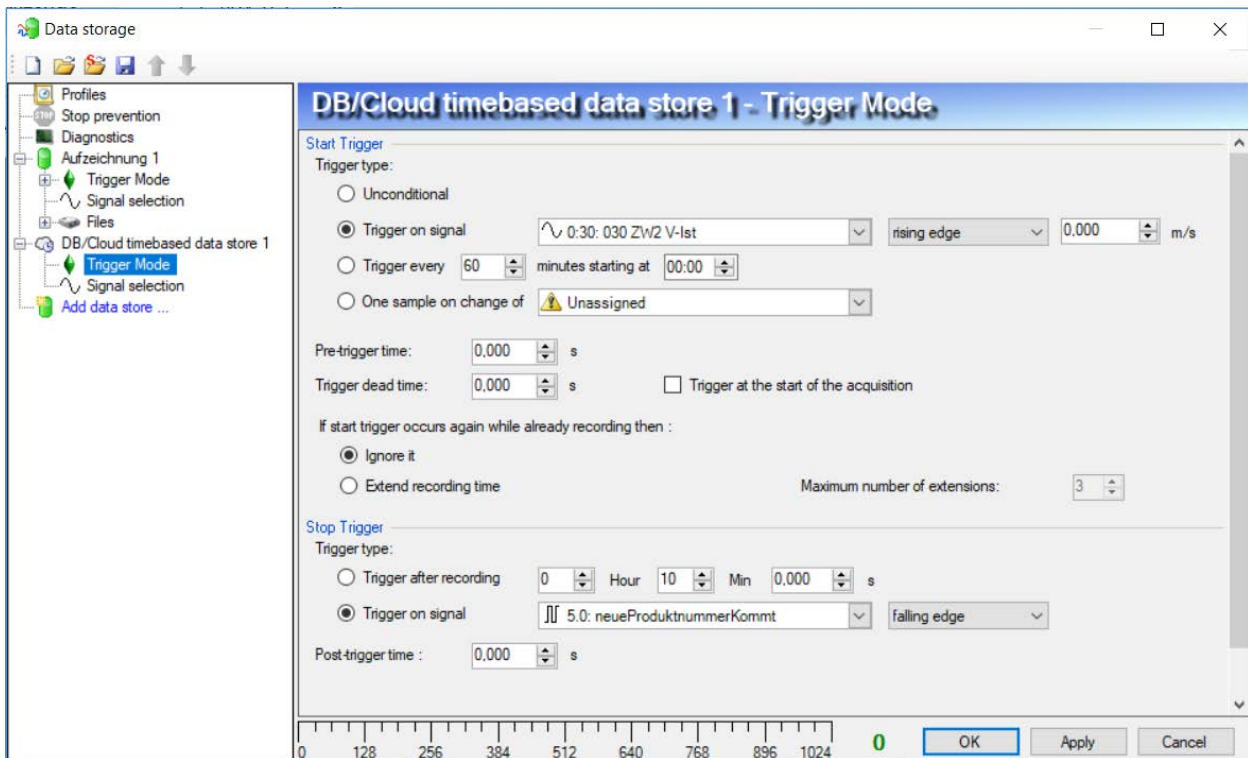In the *Trigger Mode* node, you determine when data is recorded.

Fig. 7: Trigger mode, Database/cloud example

Start trigger

You initially choose whether you would like to continuously record or it should be fired by a trigger.

Unconditional

The data is continuously recorded with this selection. In this case, the recording will start immediately at the start of the measurement or when pressing the "GO" button.

Trigger on signal

If you want the trigger to fire on a measured signal or a virtual signal, you need to check *Trigger on signal* in the option field. In the fields next to this, define the properties of the trigger signal.

- ■ Field 1: Drop-down list for signal selection (available analog and digital signals)

- ■ Field 2: Drop-down list for selecting edges or levels

- ■ Field 3: Drop-down list for selecting the trigger level value given in the specific physical unit (field 3 is only enabled in case of analog trigger signals)

Both analog and digital signals can serve as triggers. The signal to trigger on is to be selected from the drop-down lists (see picture below, field 1). In the drop-down list, you will find the well-known signal tree containing available signals. Select the signal you want to use as trigger signal.
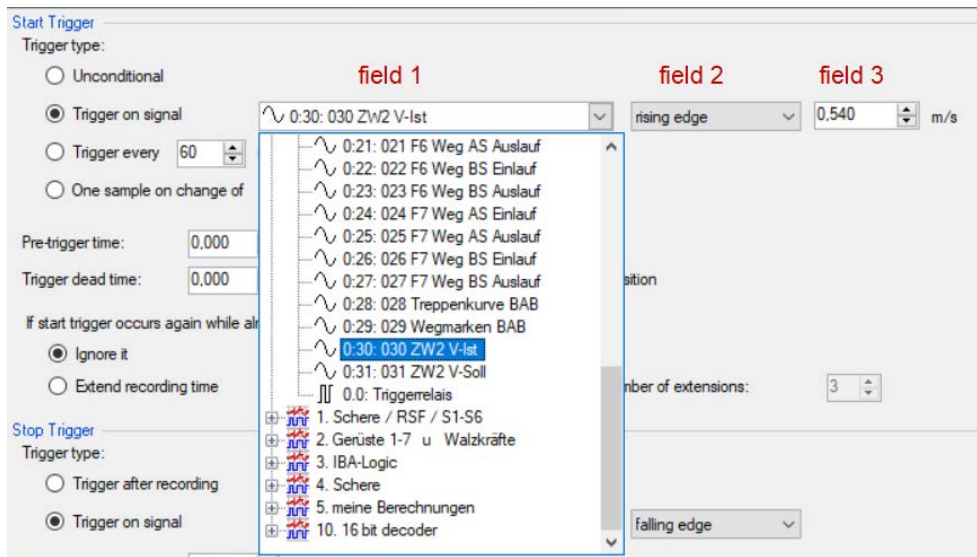

Fig. 8: Configuring "Trigger on signal"

Depending on whether a digital or an analog signal was selected, the fields 2 or 3, respectively, are offered allowing the trigger event to be defined more specifically.

As for analog signals, you can choose between level or edge triggers including a predefined level (field 3).
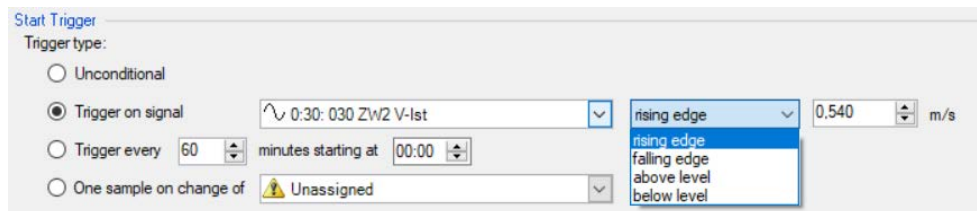

Fig. 9: Configuring "Trigger on signal", analog signal, edge or level

As for digital signals, you can choose between level or edge triggers including the 2 levels logical 0 (FALSE) and logical 1 (TRUE).
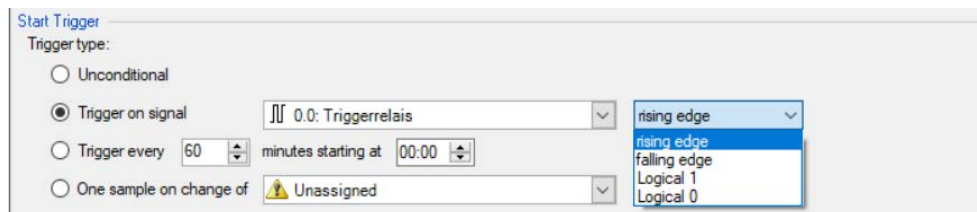

Fig. 10: Configuring "Trigger on signal", digital signal

Trigger every ...
If you want to use a start trigger always at a certain time regularly, you can check the "Trigger every … minutes starting at …" option. Enter the period given in minutes, or select it from the

input field. Value range is from 0 to 1440, which equals one day. Then enter or select the start time for the first trigger. Value range is from 00:00 to 23:59, which equals one day.

One sample on change of
When the value of the selected signal changes, a sample is recorded. The recording will stop after one sample, until the next signal change is detected. A deadtime can be configured to determine a minimum amount of time between samples. Before the deadtime has elapsed, no new sample will be recorded.

Pre-trigger time
You can configure a pre-trigger time and then the recording begins by the pre-trigger time before the trigger event. If the trigger condition is met, the incoming data is added to the data buffered during the pre-trigger time.

Trigger dead time
This property is available for the start triggers "Trigger on signal", "Trigger every ..." and "One sample on change of". The trigger dead time determines the time of suppressing subsequent triggers after a trigger occurred.

If the dead time, for instance, is set to 5 seconds, all other triggers are ignored for the duration of 5 seconds after the first trigger occurrence.

Trigger at the start of the acquisition
If you want the recording to start immediately at acquisition start or as soon as you apply a new data storage configuration, you also need to check the *Trigger on acquisition start* option. If you do not enable this option, the recording first starts once the trigger is fired.

If start trigger occurs again while file is already recording, then:
You can determine here what should happen if a new start trigger occurs while a recording is already running.

- Ignore it:
  Selecting this option will cause the system to ignore any new start trigger during a running recording for as long as the stop trigger occurs

- Extend recording time:
  If this option is enabled, it extends the duration of the running recording upon occurrence of another start trigger during an ongoing recording. This occurs as often as set in the "Maximum number of extensions on single file" field. If the max. number of extensions is reached, all subsequent start triggers will be ignored. Of course, the recording is stopped immediately by any stop trigger.

Stop trigger

The settings for the stop trigger are made in the same way as those for the start trigger. Here, both analog and digital signals can also be used as triggers.

Trigger after recording of x hours x minutes x seconds
Here you can configure a time span according to which the recording is ended - after the occurrence of the start trigger.

Trigger on signal

See explanation for start trigger above.

Post trigger time

You can configure a post trigger time and then the recording ends by the post trigger time after the stop trigger event.

7   Diagnostics

7.1   Data storage status

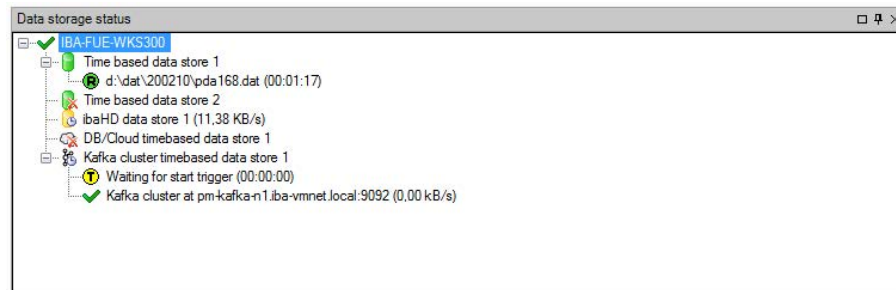The data storage status window shows the current status of the data stores.

Fig. 11: Example of data storage status window

All defined data stores and their respective status are displayed here, depending on the data store, with server address, acquisition duration, write speed, etc.

The icon in front of the name indicates the current status of the storage:

Wait for the start trigger (only for triggered recording)

Recording in progress

Post-trigger phase; stop trigger occurred, but acquisition continues until the post-trigger time is over

Disabled or faulty data store is indicated by a red cross in the data store icon.

Right-clicking on this node allows you to manually send a start or stop trigger.

7.2   Diagnostics of data stores

The *Diagnostics* node in the data storage tree offers information about the system load by the data stores. The measurement must be running.
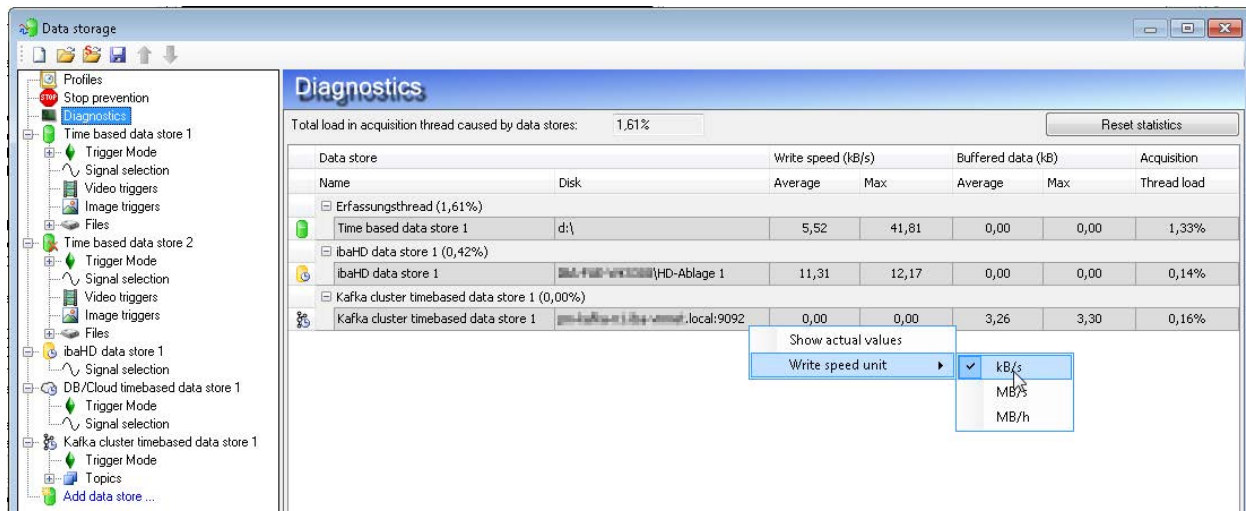

Fig. 12: Diagnostics of data stores

The performance values of all data stores are shown in the table. There is one row per data store. The rows are grouped according to the threads that write the data.

In each group row is the name of the thread and (in brackets) its share of the load. The load average is displayed by default. But, you can switch between the average and actual value using the context menu.

The *Disk* column indicates the respective target to which the data is written, for example a hard disk partition, the address of the database, the address of the Kafka cluster, etc. The *Write speed* indicates how fast the data is written. The *Memory buffer (kB)* columns indicate how much data is buffered in *ibaPDA*.

The *Acquisition Thread load* column indicates various information depending on the data stores. For timebased data stores, the *Acquisition Thread load* column indicates the amount of time needed for the run length encoding and writing to a disk or in a Kafka cluster. For database/ cloud, MQTT and MindSphere data stores, the column indicates the load caused by the analysis of the triggers and creation of the row data.

For HD data stores, the partial processing time will be displayed, that is used for the creation of the data to be written on the HD server. These values already contain the run length encoding for time-based stores, event trigger calculation for event-based stores and the calculation of the length-based data for length-based stores.

Additional information about diagnostics can be found in the *ibaPDA* manual, part 5.

8 Support and contact

Support

Phone: +49 911 97282-14

Fax: +49 911 97282-33

Email: support@iba-ag.com

Note

| i | If you require support, indicate the serial number (iba-S/N) of the product or the license number. |

Contact

Headquarters

iba AG
Koenigswarterstrasse 44
90762 Fuerth
Germany

Phone: +49 911 97282-0

Fax: +49 911 97282-33

Email: iba@iba-ag.com

Mailing address

iba AG
Postbox 1828
D-90708 Fuerth, Germany

Delivery address

iba AG
Gebhardtstrasse 10
90762 Fuerth, Germany

Regional and Worldwide

For contact data of your regional iba office or representative
please refer to our web site

www.iba-ag.com.