



ibaPDA-Data-Store-SAP-HANA

Data streaming into SAP HANA

Manual
Issue 1.3

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

Unless explicitly stated to the contrary, it is not permitted to pass on or copy this document, nor to make use of its contents or disclose its contents. Infringements are liable for compensation.

© iba AG 2021, All rights reserved.

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 / Page	Author	Version SW
1.3	11-2022	GUI new, SAP HANA 64 bit Client for v8.1.0	st, mm	8.1.0

Windows® and SQL Server® are brands and registered trademarks of Microsoft Corporation. SAP HANA® is a registered trademark of SAP.

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

Content

1 About this manual..... 4

1.1 Target group and previous knowledge 4

1.2 Notations 5

1.3 Used symbols..... 6

2 Introduction 7

2.1 System requirements..... 7

3 Data store configuration..... 8

3.1 Add a data store 8

3.2 Database type SAP HANA 9

3.3 Buffer 12

4 Storage profiles 15

4.1 Add profile 15

4.2 Profile Time, DB/Cloud 15

5 Signal selection 18

6 Trigger mode..... 19

7 Diagnostics..... 23

7.1 Data storage status 23

7.2 Diagnostics of data stores..... 24

8 Appendix 25

8.1 Notes on SAP HANA client installation 25

9 Support and contact..... 29

1 About this manual

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

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-SAP-HANA* 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 <i>Logic diagram</i>
Calling the menu command	<i>Step 1 – Step 2 – Step 3 – Step x</i> Example: Select the menu <i>Logic diagram – Add – New function block</i> .
Keys	<Key name> Example: <Alt>; <F1>
Press the keys simultaneously	<Key name> + <Key name> Example: <Alt> + <Ctrl>
Buttons	<Key name> Example: <OK>; <Cancel>
Filenames, paths	Filename, Path Example: Test.docx

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 15.

Chapter ↗ *Signal selection*, page 18 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 19.

2.1 System requirements

The following system requirements are necessary when using data storage in a SAP HANA database:

- *ibaPDA* v8.1.0 or higher
- License for *ibaPDA-Data-Store-SAP-HANA*
- SAP HANA client installation (64 bit)

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
30.670141	ibaPDA-Data-Store-SAP-HANA-64	Data streaming to SAP HANA; 64 signals
30.670142	ibaPDA-Data-Store-SAP-HANA-256	Data streaming to SAP HANA; 256 signals
30.670143	ibaPDA-Data-Store-SAP-HANA-1024	Data streaming to SAP HANA; 1024 signals
30.670151	upgrade-ibaPDA-Data-Store-SAP-HANA-64 to 256	License for extension from 64 to 256 signals
30.670152	upgrade-ibaPDA-Data-Store-SAP-HANA-256 to 1024	License for extension from 256 to 1024 signals

Table 1: Available licenses for the data storage in SAP HANA databases

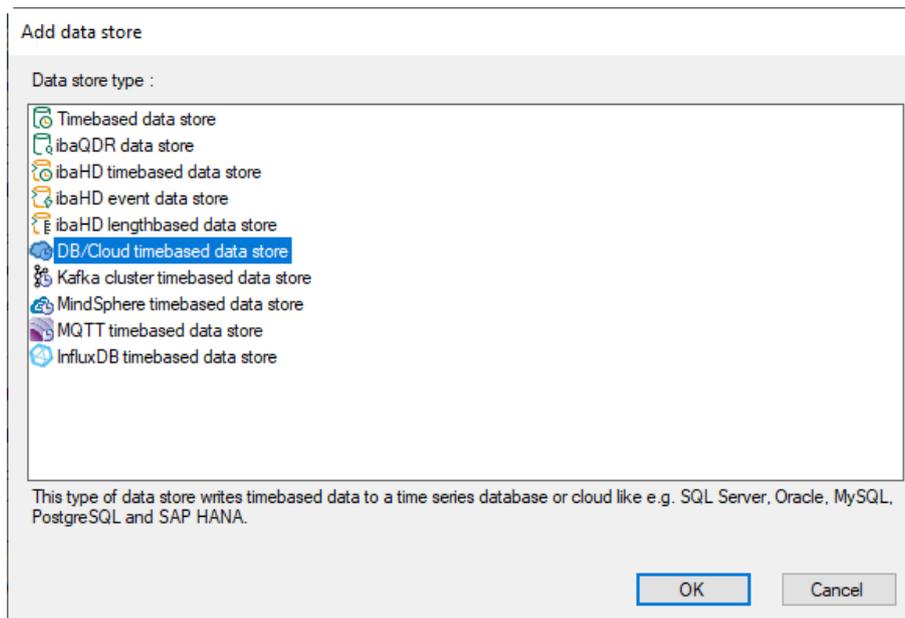
3 Data store configuration

3.1 Add a data store

Open the dialog for data storage configuration in the *Configure – Data storage* main menu or by clicking on 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.



3.2 Database type SAP HANA

First select the database type *SAP HANA* and then configure the other necessary settings.

The screenshot shows the 'Data storage' configuration window for 'DB/Cloud timebased data store 1'. The window is divided into several sections:

- General:**
 - Locked
 - Active
 - Name: DB/Cloud timebased data store 1
 - Data store index: 0
- Database:**
 - Database type: SAP HANA
 - SAP Hana database connection:
 - Server address: hanaServer:39013
 - User name: system
 - HANA client version: 2.4.171.0
 - Database: systemdb
 - Password: [masked]
 - Test connection button
- Table:**
 - Table name: IBA.TESTTABELLE
 - Timebase: 1,000 s
 - Write data every: 100 samples
 - Check table button

At the bottom of the window, there is a progress bar and buttons for OK, Apply, and Cancel.

General

Locked

You can lock a data store in order to prevent an accidental or unauthorized change of settings.

Active

Activate the data storage in order to record data. However, you can configure various data stores and disable data stores that are not required.

Data store index

Unique index of all existing DB/Cloud data stores. You need to reference this index e.g. in the virtual function `DataStoreInfoDB/Cloud()` for generating diagnostic data for a specific DB/Cloud data store.

Data store name

You can enter a name for the data store here.

Database

Database type

Select your database type from the drop-down menu, here *SAP HANA*.

Server address, Database

Enter the address of the server and the name of the database you want to connect to.

User name/Password

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

HANA client version

The version of the SAP HANA client is displayed here (information only).

Note



Prerequisite for using the data store is the installation of the SAP HANA Client (64 bit) including the ADO.NET Data Provider.

For instructions on how to install the client and how to verify the correct registration of the data providers, see chapter [Notes on SAP HANA client installation](#), page 25

<Test connection>

Use the <Test connection> button to test the connection. When testing the connection, *ibaPDA* attempts to call up the current list of available tables and fill out the *Table name* drop-down list in the *Table* area. If the connection is successful, the version of the database is also output.

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.

Note that you can only specify the time base when creating the table because it is a system parameter of the table. The time basis can be changed later only by deleting and recreating the table, which causes the data in the table to be lost.

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>

Via the <Check table> button you can 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 15 and select the signals, see [Signal selection](#), page 18.

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, a confirmation message appears.

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:

Column	Data type	Status	Fix
I_TIME	datetime2	OK	<input type="checkbox"/>
LOCAL_TIME	datetime2	Column is not available	<input checked="" type="checkbox"/>
[SP_Speed]	real	OK	<input type="checkbox"/>
[SP_Tension]	real	OK	<input type="checkbox"/>
[AV_Tension]	real	Column is not available	<input checked="" type="checkbox"/>
[AV_Temperature]	real	Column is not available	<input checked="" type="checkbox"/>
[Product_Cycle_Cnt]	real	OK	<input type="checkbox"/>
[Product_Ident]	nvarchar	OK	<input type="checkbox"/>
[Customer]	nvarchar	OK	<input type="checkbox"/>
[Product_Start]	bit	Column is not available	<input checked="" type="checkbox"/>
[Tension_high]	bit	OK	<input type="checkbox"/>
[Temperature_high]	bit	OK	<input type="checkbox"/>
SP_Temperature	real	Column is no longer required	<input type="checkbox"/>
AV_Speed	real	Column is no longer required	<input type="checkbox"/>

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.
- 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. They cannot be deselected because a correction is mandatory. 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.

3.3 Buffer

The data storage uses a memory buffer and additionally a file buffer that can be enabled optionally.

The description applies to all types of data stores that transfer data to external systems and where temporary accessibility and available bandwidth issues may occur, such as:

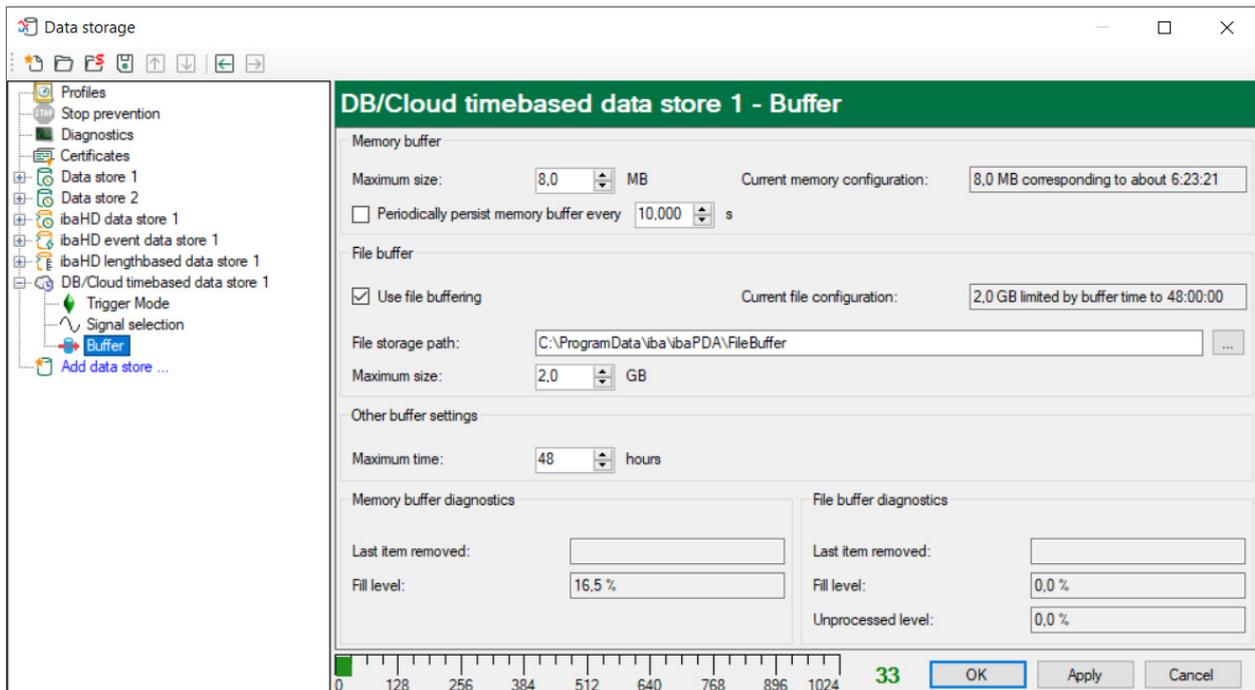
- ibaHD timebased/eventbased/lengthbased
- DB/Cloud timebased
- Kafka cluster timebased
- MQTT timebased
- MindSphere timebased
- InfluxDB timebased

Data to be sent to the target system always passes through the internal *ibaPDA* memory buffer. If the connection to the target system exists, the data is sent there from the memory buffer immediately. If the connection is lost, or the data cannot be sent out fast enough, the data remain in the memory buffer. The memory buffer is located in the RAM of the *ibaPDA* computer and is therefore limited and volatile. If, for example, the acquisition is restarted, the buffered data will be lost. If the memory buffer grows beyond the configured size during ongoing acquisition, the oldest values are deleted and thus lost.

To improve this, a file buffer can additionally be enabled, which can buffer much larger amounts of data. The data is stored in files in a directory in a local drive of the *ibaPDA* server. When the file buffer is enabled, data is transferred from the overflowing memory buffer to the file buffer. If the acquisition is finished or restarted (e.g. by applying a modified IO configuration), data that may be in the memory buffer at this time is also transferred to the file buffer.

After reconnecting to the target system, the oldest data is always transferred first. Newer values are added to the buffer in the meantime. If there is still buffered data in the file buffer when the acquisition is started, it will be handled and processed in the same way. The data is saved in the format that was configured in the data store at the time of buffering and it is also sent in this format when the connection is established again.

You configure the buffering in the *Buffer* node of the respective data store. The figure below shows the buffer configuration using the example of DB/Cloud data store.



Memory buffer

The memory buffer is always enabled. It cannot be deactivated, since data to be transmitted always passes through the buffer before being forwarded to the target system.

Maximum size

Enter here the maximum total size for items buffered in memory. If the maximum size is exceeded, there are 2 options:

- When file buffering is disabled, the oldest item in memory is deleted (and is lost forever).
- When file buffering is enabled, the oldest part of the buffer memory is moved to a buffer file.

Periodically persist memory buffer every ... s

This option can be enabled only if file buffering is enabled. If the option is enabled, the entire memory buffer is periodically swapped to a buffer file.

Enter a duration after which the memory buffer is periodically stored. It must be between 10 and 600 s.

With this option you can ensure that as little data as possible is lost in case of a system failure.

Current memory configuration

Display of the approximate time period that can be temporarily stored in the memory buffer with the configured settings. Specified in d.hh:mm:ss.

File buffer

Use file buffering

By default, the file buffer is not used. Here you can enable file buffering.

Current file configuration

Display of the approximate time period that can be temporarily stored in the file buffer with the configured settings. Specified in d.hh:mm:ss.

File storage path

In the *File storage path* field you can select a location for the files. You can enter the directory directly into the text field, or select it via the browse button <...>. The configured file directory must be located on a local hard disk of the *ibaPDA* server computer.

The same file directory can be used for several data stores, because the buffer files of a data store have a unique name. Files from different data stores can thus be distinguished by their name.

Maximum size

You can configure the maximum total size of the buffer files of a data store. The buffer files themselves have the file extension *.buf*, the index file for managing the buffer files has the extension *.info*. The maximum size is the total size of all these files. If the maximum buffer size is exceeded, the oldest buffer file is deleted.

Other buffer settings

Maximum time

Stored data older than the maximum time will not be transferred to the target system. Files older than the maximum time can be deleted. You can enter a value between 1 and 1000 hours.

Memory buffer diagnostics / File buffer diagnostics

Last item removed

Indicates when the last item was taken from this part of the buffer.

Fill level

The fill level indicates what percentage of the buffer size is currently filled with buffered data.

Unprocessed level

Items transferred to the target system are not deleted immediately in the file buffer. Only when a buffer file is completely read, it is deleted. Therefore, it is possible that only a part of a buffer file contains data that has not yet been transferred. The fill level refers to the existing buffer files, while the "unprocessed level" indicates the percentage of data in the file buffer that has not yet been transferred.

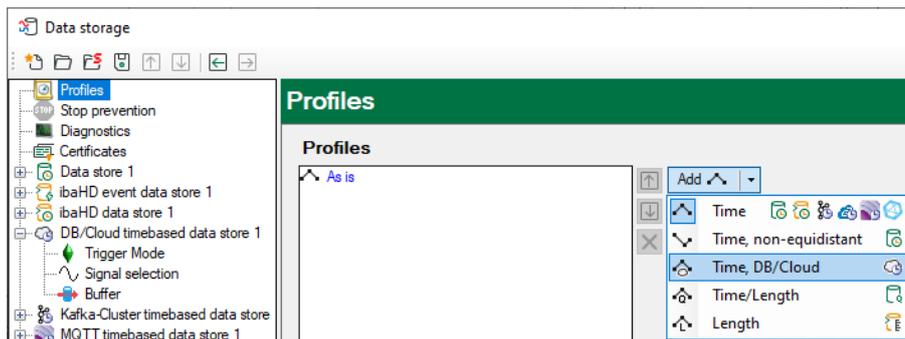
4 Storage profiles

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.

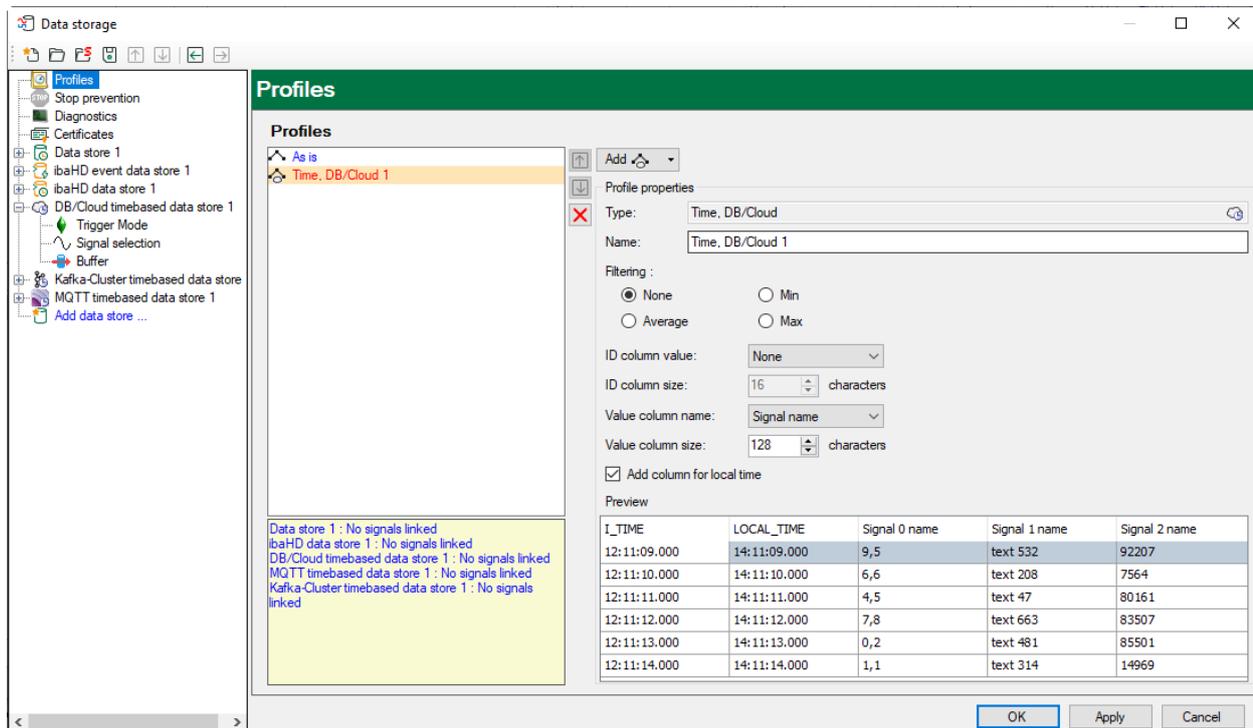
4.1 Add profile

To add a DB profile, select the *Profiles* node in the tree structure of the data storage configuration dialog.

Open the drop-down list on the right under <Add> and select *Time, DB/Cloud*.



4.2 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 computer 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 column 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.

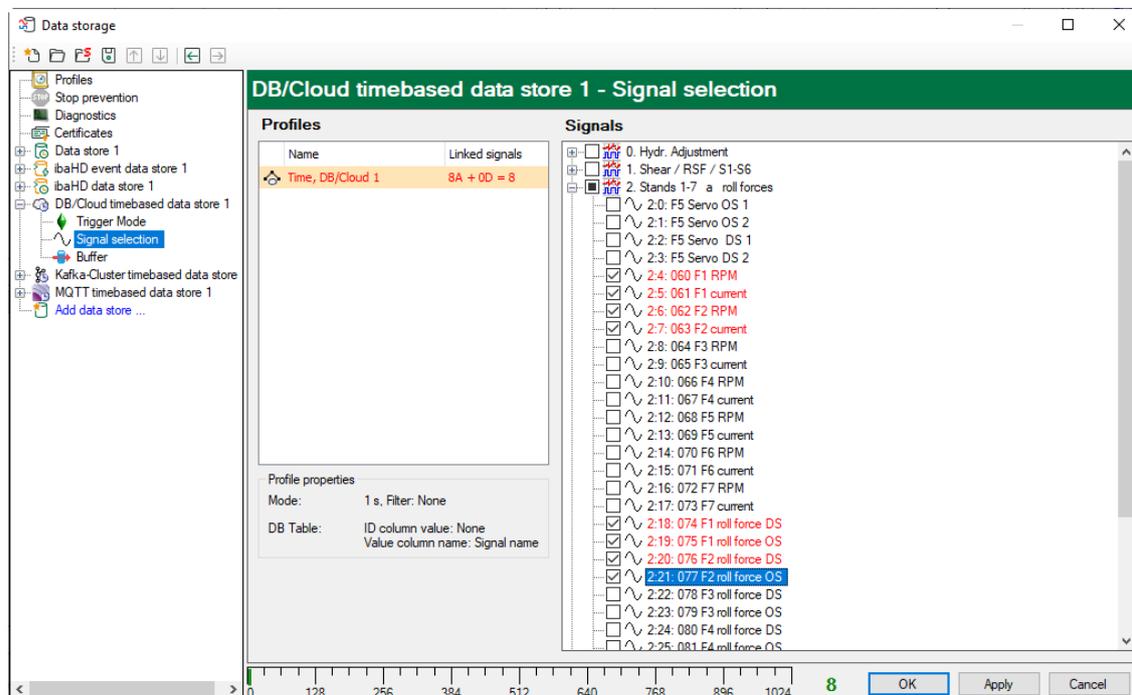
The preview shows a general example how a table with the current settings will look like.

Note

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

5 Signal selection

To record signals, you have to assign the signals to a storage profile of the type *Time, DB/Cloud*. Select the signal selection node below your *DB/Cloud timebased data store* to open the signal selection dialog.



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.

You can find the current number of selected signals in all DB/Cloud data stores of one database type at the bottom of the dialog, similar to the number of configured signals in the I/O Manager. The length of the bar corresponds to the licensed number of signals.

In the above example, it is possible to write up to 1024 signals in several data stores of the same DB type. Currently 8 signals are enabled.

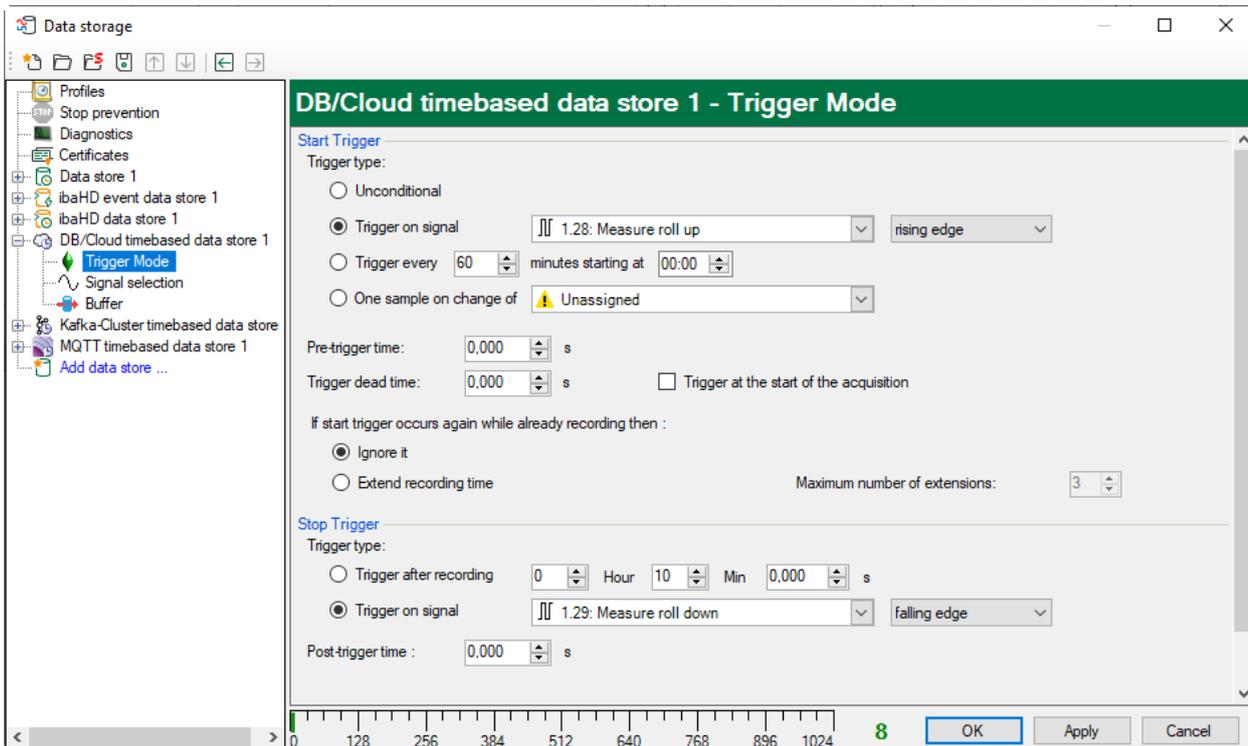
When you have configured all signals you want to write to the database, go back to the main node of the data store. There you can check the table with the <Check table> button.

6 Trigger mode

The description applies to the following types of data stores:

- DB/Cloud timebased
- Kafka cluster timebased
- MQTT timebased
- MindSphere timebased
- InfluxDB timebased.

In the *Trigger Mode* node, you determine when data is recorded, here using the example of *DB/Cloud timebased data store*.



Start trigger

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

Unconditional

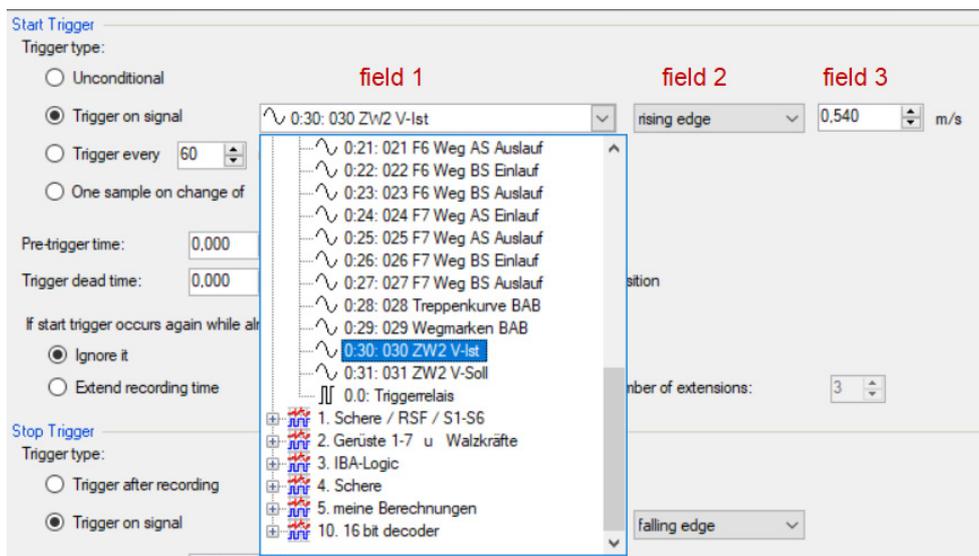
The data is continuously recorded with this selection. In this case, the recording will start immediately at the start of the measurement or when pressing the "GO" button.

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.

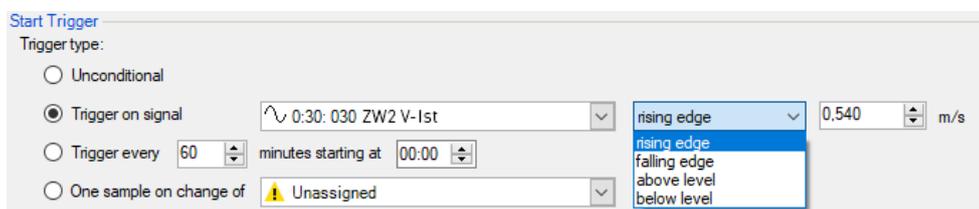
- Field 1: Drop-down list for signal selection (available analog and digital signals)
- Field 2: Drop-down list for selecting edges or levels
- Field 3: Drop-down list for selecting the trigger level value given in the specific physical unit (field 3 is only enabled in case of analog trigger signals)

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.

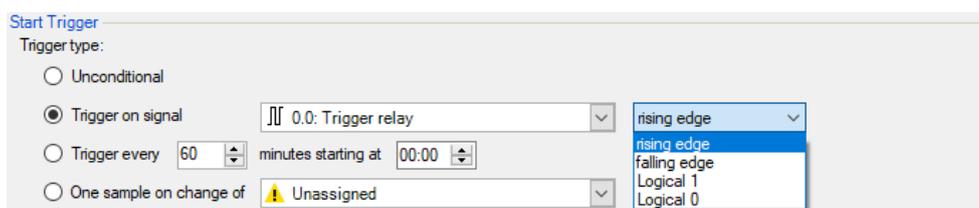


Depending on whether a digital or an analog signal was selected, the fields 2 or 3, respectively, are offered allowing the trigger event to be defined more specifically.

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



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



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:
If you select this option, the system ignores any new start trigger during a running recording for as long as the stop trigger occurs
- Extend recording time:
If you select this option, the system 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" 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.

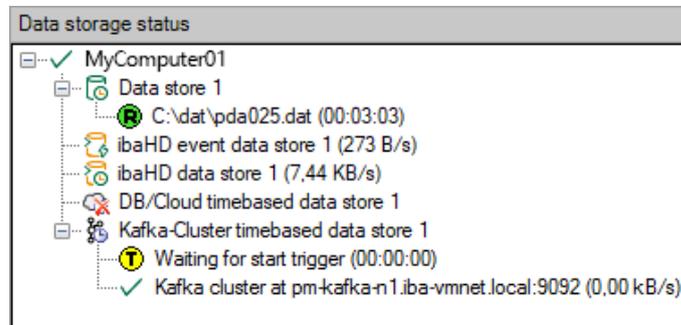
Post trigger time

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.



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.

The screenshot shows the 'Data storage' application window with the 'Diagnostics' tab selected. The window title is 'Data storage'. On the left, there is a tree view showing the data storage hierarchy, including 'Profiles', 'Stop prevention', 'Diagnostics', 'Certificates', and several data stores like 'Data store 1', 'ibaHD event data store 1', 'ibaHD data store 1', 'DB/Cloud timebased data store 1', 'Kafka-Cluster timebased data store', and 'MQTT timebased data store 1'. The main area displays a table of performance metrics for various data stores. The table has columns for 'Data store Name', 'Destination', 'Write speed (KB/s) Average and Max', 'Memory buffer (KB) Average and Max', 'File buffer (MB) Average and Max', and 'Acquisition Thread load'. A 'Total load in acquisition thread caused by data stores:' field shows 0,36% with a 'Reset statistics' button. A context menu is open over the table, showing 'Show actual values' and 'Write speed unit' with options 'KB/s', 'MB/s', and 'MB/h'. At the bottom, there are 'OK', 'Apply', and 'Cancel' buttons.

Data store Name	Destination	Write speed (KB/s)		Memory buffer (KB)		File buffer (MB)		Acquisition Thread load
		Average	Max	Average	Max	Average	Max	
Erfassungsthread (0,36%)								
Data store 1	C:\	9,02	53,64	0,00	0,00			0,19%
ibaHD event data store 1 (0,92%)								
ibaHD event data store 1	MyComputer01\HD_TRIG	0,27	0,39	0,00	0,01			0,02%
ibaHD data store 1 (0,21%)								
ibaHD data store 1	MyComputer01\HD_Time2	7,01	9,28	0,00	1,59			0,04%
DB/Cloud timebased data store 1 (0,00%)								
DB/Cloud timebased dat...	hanaServer:39013	0,00	0,00	0,00	0,00			0,03%
Kafka-Cluster timebased data store 1 (0,00%)								
Kafka-Cluster timebased...	localhost:9092	0,00	0,00	268,83	296,07			0,02%
MQTT timebased data store 1 (0,00%)								
MQTT timebased data st...	localhost	0,00	0,00	333,16	355,22			0,05%

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) the utilization (100% means the thread is overloaded). The load average is displayed by default. But, you can switch between the average and actual value using the context menu.

The *Destination* 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 columns *File buffer (MB)* indicate how much data is buffered in the file buffer.

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. For DB/Cloud, MQTT, Kafka Cluster, InfluxDB 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 timebased stores, event trigger calculation for event stores and the calculation of the length-based data for lengthbased stores.

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

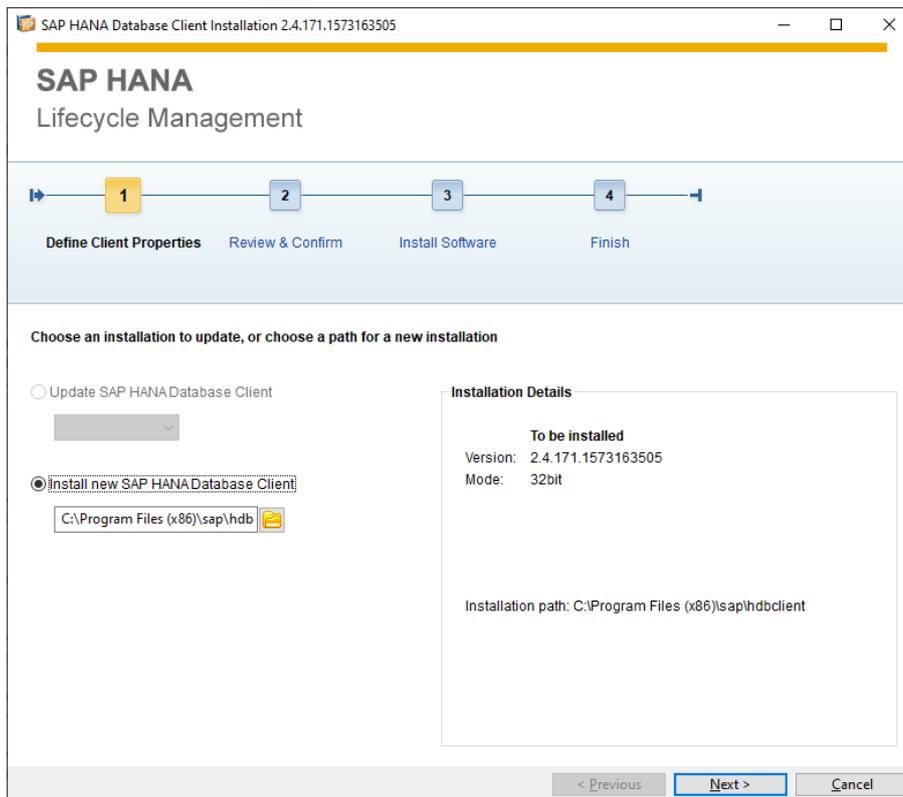
8 Appendix

8.1 Notes on SAP HANA client installation

Writing data from *ibaPDA* to a SAP HANA database via the SQL interface or DB/Cloud requires the installation of the SAP HANA client (64 bit).¹⁾

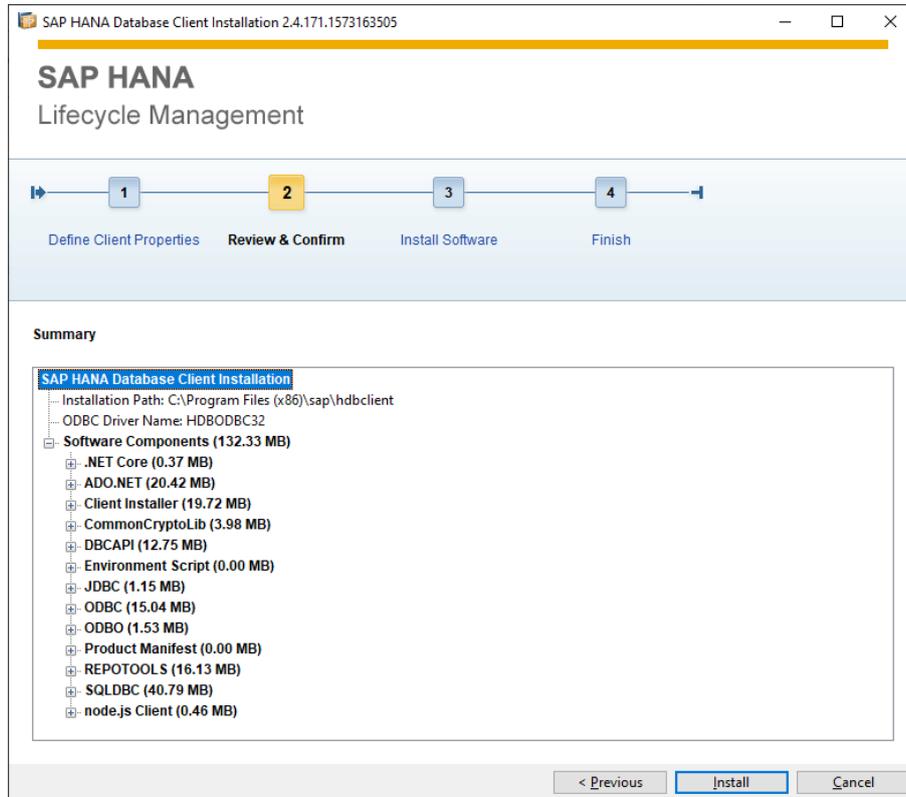
The installation of the SAP HANA client is described below step by step.

1. The installer with graphical user interface is started by double clicking on [hdbsetup.exe](#).
2. Select the option *Install new SAP HANA Database Client*.
Specify the installation directory and continue with <Next>.

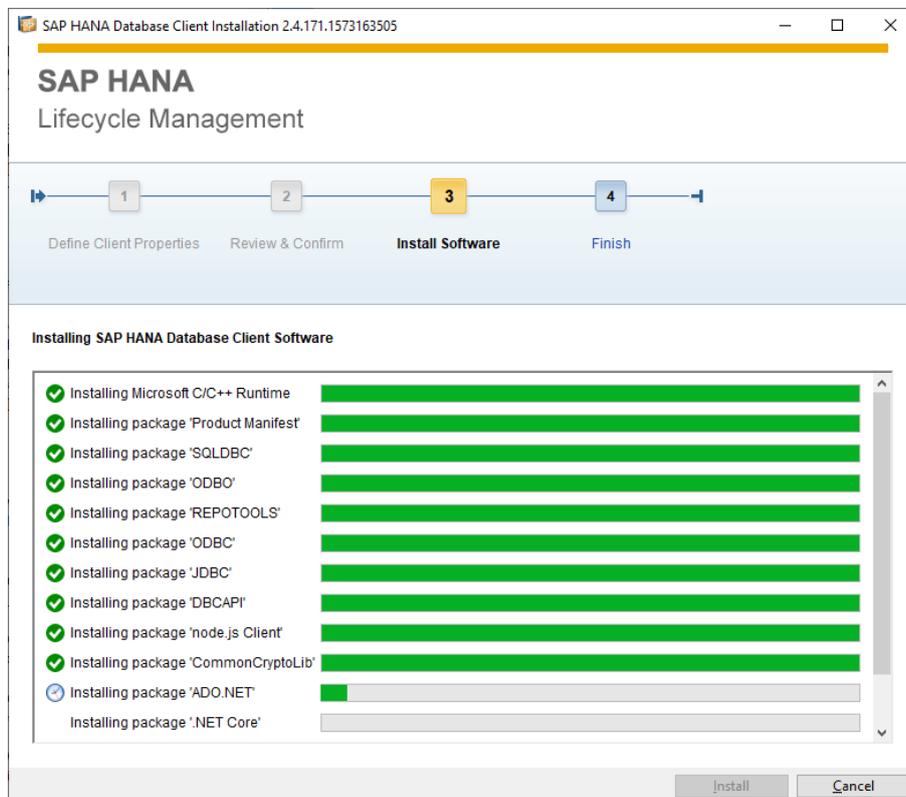


1) The SAP HANA client 64 bit runs with all ibaPDA versions including v8.1 32 bit.

3. A manual selection of the different driver components is not possible. All components are always installed. It is important that the ADO.NET component appears in the list. Confirm the settings for the installation and continue with <Install>.



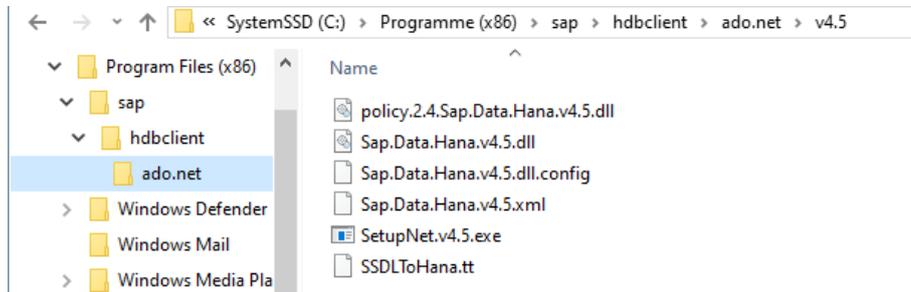
→ The installation is executed.



4. Confirm the successful installation with <Finish>.

→ After successful installation of the SAP HANA Client you find the ADO.NET Data Provider components in the following path by default:

`C:\Program Files (x86)\sap\hdbclient\ado.net\v4.5`



5. The data providers must also be registered in the `machine.config` configuration file of the computer. This is normally done automatically by the SAP HANA Client Installer. For verification open the configuration e.g. with the notepad text editor. You find the file here:

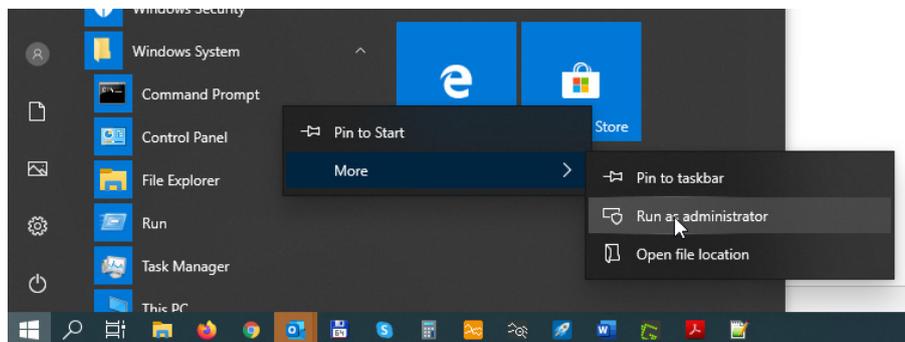
`%windir%\Microsoft.NET\Framework\v4.0.30319\Config`

In the file there should be an entry "SAP HANA Data" under the node "<DbProviderFactories>".

```
<system.data>
<DbProviderFactories>
  <add name="SAP HANA Data Provider" invariant="Sap.Data.Hana" description=".NET Framework Data Provider for SAP HANA" type="Sap.Data.Hana.HanaFactory, Sap.Data.Hana.v4.5, Version=2.4.171.0, Culture=neutral, PublicKeyToken=0326b8ea63db4bc4" />
</DbProviderFactories>
</system.data>
```

6. If the "SAP HANA Data Provider" registration is not available, you can execute it manually.

- To do this, open a command prompt with the option *Run as administrator*.



- Change to the installation directory of SAP HANA Data Providers:
`C:\Program Files (x86)\sap\hdbclient\ado.net\v4.5\`
- Execute the command "SetupNet.v4.5.exe /i".

→ The SAP HANA Data Providers is then registered in the `machine.config` configuration file.

```
Administrator: Eingabeaufforderung
Microsoft Windows [Version 10.0.18362.959]
(c) 2019 Microsoft Corporation. Alle Rechte vorbehalten.

C:\Windows\system32>cd C:\Program Files (x86)\sap\hdbclient\ado.net\v4.5
C:\Program Files (x86)\sap\hdbclient\ado.net\v4.5>SetupNet.v4.5.exe /i

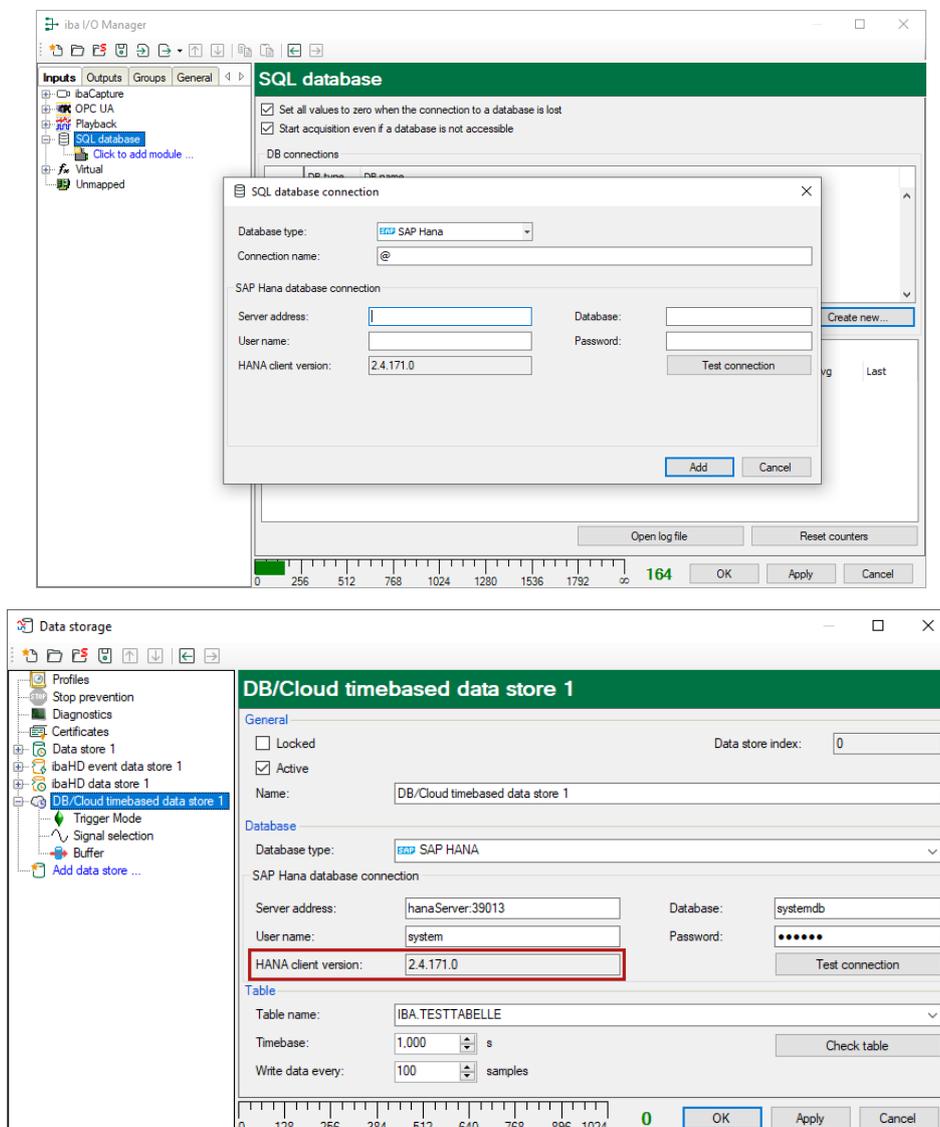
HANA ADO.NET Setup Utility Version 2.4.171.0
Copyright © 2013 SAP SE or an SAP affiliate company. All rights reserved.

Succeeded

C:\Program Files (x86)\sap\hdbclient\ado.net\v4.5>
```

7. Restart the *ibaPDA* service.

→ In the configuration of the database connections in the *ibaPDA* client (SQL interface and DB/Cloud data store) a HANA client version number is now displayed:



9 Support and contact

Support

Phone: +49 911 97282-14
Fax: +49 911 97282-33
Email: support@iba-ag.com

Note



If you need support for software products, please state the license number or the CodeMeter container number (WIBU dongle). For hardware products, please have the serial number of the device ready.

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.