



# ibaPDA-Data-Store-SQL-Server

## Data streaming into SQL Server

Manual  
Issue 1.3

Measurement Systems for Industry and Energy  
[www.iba-ag.com](http://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](mailto:iba@iba-ag.com)  
Web [www.iba-ag.com](http://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 2022, 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](http://www.iba-ag.com).

| Version | Date    | Revision - Chapter / Page | Author | Version SW |
|---------|---------|---------------------------|--------|------------|
| 1.3     | 08-2022 | Configuration, port       | st     | 8.0.0      |

Windows® and SQL Server® are brands and registered trademarks of Microsoft Corporation.

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 SQL Server..... 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 Support and contact.....25**

# 1 About this manual

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

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-SQL-Server* 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><br>Example:<br>Select the menu <i>Logic diagram – Add – New function block</i> . |
| Keys                          | <Key name><br>Example: <Alt>; <F1>  |
| Press the keys simultaneously | <Key name> + <Key name><br>Example: <Alt> + <Ctrl>  |
| Buttons                       | <Key name><br>Example: <OK>; <Cancel>   |
| Filenames, paths              | <a href="#">Filename, Path</a><br>Example: <a href="#">Test.docx</a>  |

## 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 SQL Server database:

- *ibaPDA* v7.3.0 or higher
- License for *ibaPDA-Data-Store-SQL-Server*

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  |
|-----------|--|--|
| 30.671030 | ibaPDA-Data-Store-SQL-Server-64                  | Data streaming into SQL Server database, max. 64 signals   |
| 30.671031 | ibaPDA-Data-Store-SQL-Server-256                 | Data streaming into SQL Server database, max. 256 signals  |
| 30.671032 | ibaPDA-Data-Store-SQL-Server-1024                | Data streaming into SQL Server database, max. 1024 signals |
| 30.671035 | upgrade-ibaPDA-Data-Store-SQL-Server-64 to 256   | License for extension from 64 to 256 signals               |
| 30.671036 | upgrade-ibaPDA-Data-Store-SQL-Server-256 to 1024 | License for extension from 256 to 1024 signals             |

Table 1: Available licenses for the data storage in SQL Server databases

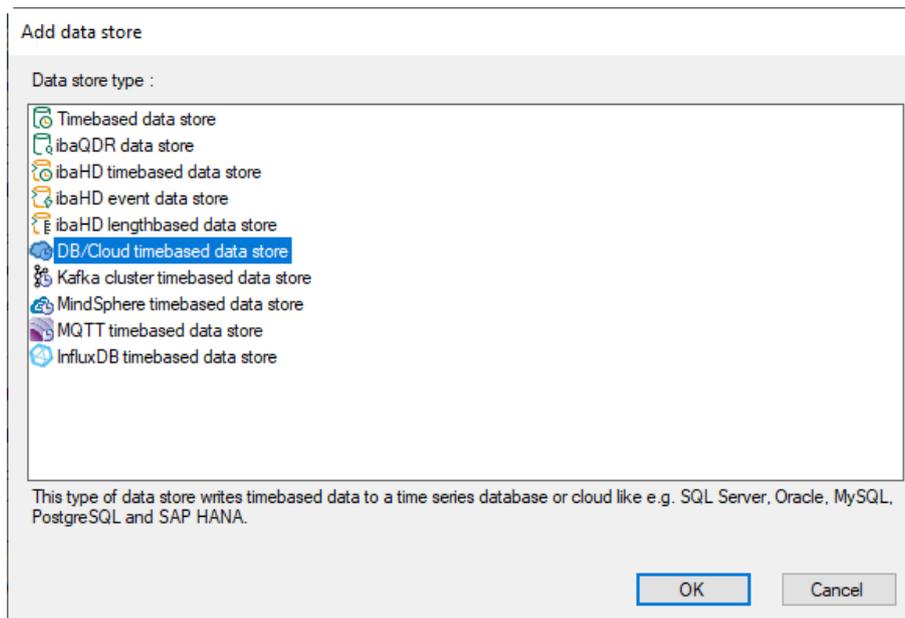
## 3 Data store configuration

### 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.



## 3.2 Database type SQL Server

First select the database type *SQL Server* and then configure the other necessary settings.



### 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.

### Database

#### Database type

Select your database type from the drop-down menu, here SQL Server.

#### Server address

Enter the IP address, host name, or a combination of host name and instance name of the SQL Server database server. You can enter the server addresses or select them from the drop-down list. The drop-down list shows the database instances that can be reached from the *ibaPDA* server. Only the registered SQL Server instances on the computer where *ibaPDA* server is installed are shown.

A port can also be specified with the server address. The port is separated from the server name with a comma. If no port is specified for the server address, the default port 1433 is used.

Example: "192.168.0.1,1433"

## Database

Enter the database here. If the Server address, authentication (see below) and required user name and password are entered, the drop-down list shows all available databases for the connection and you may select one of them.

## Authentication

Different options for authentication are available:

- Windows authentication

The user account the *ibaPDA* server is running on is used to connect to the SQL Server instance. This is usually the system account. The user can also be changed to another user with administrator rights.

---

### Note



When using the local system account, the corresponding login NT-AUTHORITY\SYSTEM in SQL Server must be given additional permissions that are not set up by default. Required are the rights to connect to the database, to create and modify a table, and to write to a table. Clarify in advance with your DB administrator whether the extension of rights is possible or whether a different login or authentication should be selected.

---

- SQL Server authentication

Authentication that is stored within the database. A user name and password are required to use this authentication method.

- Azure Active Directory – Universal with MFA:

Multi-factor authentication, as configured on Azure. More information can be found here: <https://docs.microsoft.com/en-us/azure/sql-database/sql-database-aad-authentication-configure?tabs=azure-powershell>

- Active Directory – Password or Integrated

These authentication methods can be used for members of a user group in Active Directory that has access to the database. A trust relationship must be established between the user and its device and the domain. If you select *Active Directory - Password*, user name and password must be entered.

Depending on the selection you have to enter the user name or a password.

### <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 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, 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:

| Column              | Data type | Status                       | Fix                                 |
|---------------------|-----------|------------------------------|-------------------------------------|
| [T_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 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.

### 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/event/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 s 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 / 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* 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.



### 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 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 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 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.



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 all types of data stores that transfer data to external systems, such as:

- ibaHD time/event/lengthbased
- 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:  
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.



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 hierarchy of data stores. 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 context menu is open over the 'Write speed' column, showing options to 'Show actual values' and 'Write speed unit' (KB/s, MB/s, MB/h). The total load in acquisition thread caused by data stores is shown as 0,36%.

| Data store Name                              | Destination             | Write speed (KB/s) Average | Write speed (KB/s) Max | Memory buffer (KB) Average | Memory buffer (KB) Max | File buffer (MB) Average | File buffer (MB) Max | Acquisition Thread load |
|--|-------------------------|----------------------------|------------------------|----------------------------|------------------------|--------------------------|----------------------|-------------------------|
| 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                     | IBA-FUE-WKS366\HD_TRIG  | 0,27                       | 0,39                   | 0,00                       | 0,01                   |                          |                      | 0,02%                   |
| ibaHD data store 1 (0,21%)                   |                         |                            |                        |                            |                        |                          |                      |                         |
| ibaHD data store 1                           | IBA-FUE-WKS366\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) 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 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 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](http://www.iba-ag.com)**.