
Agile Data

unknown

Jun 07, 2026

CONTENTS

1 Overview	3
1.1 Simple to learn	4
1.2 Not a traditional ORM	4
1.3 Concern Separation	4
1.3.1 Class: Field	4
1.3.2 Class: Model	5
1.3.3 Class: Persistence	5
1.4 Code Layers	5
1.4.1 Domain-Model Code	6
1.4.2 Persistence-specific code	6
1.4.3 Generic Persistence-code	7
1.5 Persistence Scaling	7
2 Quickstart	9
2.1 Requirements	9
2.2 Core Concepts	9
2.2.1 Persistence Domain vs Business Domain	10
2.2.2 Class vs In-Line definition	10
2.2.3 Model State	11
2.3 Getting Started	13
2.3.1 Adding Fields	14
2.3.2 Table Joins	14
2.3.3 Understanding Persistence	15
2.4 References between Models	16
2.4.1 One to Many	16
2.4.2 Many to Many	16
2.4.3 One to One	17
2.4.4 Implementation of References	17
2.5 Actions	17
2.5.1 Aggregation actions	17
2.5.2 Field-reference actions	18
2.5.3 Advanced Use of Actions	18
2.6 Expressions	19
2.7 Conclusion	19
3 Introduction to Architectural Design	21
3.1 The Domain Layer Scope	22
3.1.1 The Danger of Raw Queries	22
3.1.2 Purity levels of Domain code	22
3.2 Domain Logic	23

3.2.1	Domain Models	23
3.2.2	Domain Model Methods	23
3.2.3	Domain Model Fields	24
3.2.4	Domain Model Relationship	25
3.3	Persistence backed Domain Logic	25
3.3.1	ID Field	26
3.4	Persistence-specific Code	26
3.4.1	Domain Model Expressions	26
3.4.2	Persistence Hooks	27
3.5	DataSet Declaration	27
3.6	Domain Conditions	28
3.7	Related DataSets	28
3.7.1	Domain Model Actions	29
3.7.2	Unique Features of Persistence Layer	30
4	Model	31
4.1	Understanding Model	32
4.1.1	Model object = Data Set	32
4.1.2	Model object = meta information	32
4.1.3	Domain vs Persistence	33
4.1.4	Good naming for a Model	33
4.2	Initialization	34
4.2.1	Fields	34
4.2.2	Actions	36
4.2.3	Hooks	37
4.2.4	Inheritance	38
4.3	Associating Model with Database	38
4.4	Populating Data	39
4.5	Working with selective fields	40
4.6	Setting and Getting active record data	40
4.7	Title Field, ID Field and Model Caption	42
4.7.1	ID Field	42
4.7.2	Title Field	42
4.7.3	Model Caption	43
4.8	Setting limit and sort order	43
5	Typecasting	45
5.1	Value types	46
5.1.1	Undefined type	46
5.1.2	Type of IDs	46
5.1.3	Supported types	46
5.1.4	Types and UI	47
5.1.5	Array and Object types	47
6	Loading and Saving (Persistence)	49
6.1	Associating with Persistence	49
6.1.1	Inserting Record with a specific ID	50
6.2	Type Converting	51
6.2.1	Strict Types an Normalization	51
6.2.2	Typecasting	52
6.2.3	Validation	52
6.2.4	Multi-column fields	53
6.2.5	Dates and Time	54
6.2.6	Customizations	54

6.3	Duplicating and Replacing Records	54
6.3.1	Create copy of existing record	54
6.3.2	Duplicate then save under a new ID	55
6.4	Working with Multiple DataSets	55
6.4.1	Cloning versus New Instance	55
6.4.2	Looking for duplicates	55
6.4.3	Archiving Records	56
6.5	Working with Multiple Persistencies	56
6.5.1	Creating Cache with Memcache	57
6.5.2	Using Read / Write Replicas	58
6.5.3	Archive Copies into different persistence	59
6.5.4	Store a specific record	59
6.6	Actions	59
6.6.1	Action Types	60
6.6.2	SQL Actions	60
6.6.3	SQL Actions on Linked Records	61
6.6.4	Action Matrix	61
7	Fetching results	63
7.1	Iterate through model data	63
7.1.1	Raw Data Fetching	63
7.1.2	Fetching data through action	64
8	Field	65
8.1	Purpose of Field	65
8.1.1	Field Type	65
8.1.2	Basic Properties	66
8.1.3	UI Presentation	68
9	Conditions and DataSet	69
9.1	Basic Usage	69
9.1.1	Operations	70
9.1.2	Multiple Conditions	70
9.1.3	Adding OR Conditions	70
9.1.4	Defining your classes	71
9.2	Vendor-dependent logic	71
9.2.1	Field Matching	71
9.2.2	Expression Matching	71
9.2.3	SQL Expression Matching	72
9.2.4	Custom Parameters in Expressions	72
9.2.5	Expression as first argument	73
9.3	Advanced Usage	73
9.3.1	Model Scope	73
9.3.2	Conditions on Referenced Models	75
10	SQL Extensions	77
10.1	Default Model Classes	77
10.1.1	SQL Field	77
10.1.2	SQL Reference	78
10.1.3	Expressions	79
10.2	Transactions	79
10.3	Custom Expressions	80
10.4	Actions	80
10.4.1	Action: select	81
10.4.2	Action: count	81

10.4.3	Action: field	81
10.4.4	Action: fx	81
10.5	Stored Procedures	82
10.5.1	Compatibility Warning	82
10.5.2	as a Model method	82
10.5.3	as a Model Field	84
10.5.4	as an Action	84
10.5.5	as a Temporary Table	85
10.5.6	as a Model Source	86
11	Static Persistence	87
11.1	Usage	87
11.1.1	Saving Records	88
12	References	89
12.1	Persistence	90
12.2	Safety and Performance	90
12.3	hasMany Reference	91
12.3.1	Dealing with many-to-many references	91
12.3.2	Dealing with NON-ID fields	91
12.3.3	Concatenating Fields	92
12.3.4	Add Aggregate Fields	92
12.3.5	Available Aggregation Functions	93
12.3.6	Aggregate Expressions	93
12.4	hasMany / refLink	93
12.5	hasOne reference	94
12.5.1	Traversing loaded model	94
12.5.2	Traversing DataSet	94
12.5.3	Importing Fields	95
12.5.4	Importing hasOne Title	96
12.6	User-defined Reference	96
12.7	Reference Discovery	97
12.8	Deep traversal	97
12.9	Reference Aliases	98
12.10	Various ways to specify options	99
12.11	References with New Records	99
12.12	Reference Classes	100
13	Expressions	101
13.1	Defining Expression	101
13.2	No-table Model Expression	102
13.3	Expression Callback	102
13.4	Model Reloading after Save	103
14	Model from multiple joined tables	105
14.1	Join Basics	105
14.1.1	Strong and Weak joins	106
14.1.2	Join relationship definitions	106
14.1.3	Method Proxying	107
14.1.4	Create and Delete behavior	107
14.1.5	Implementation Detail	108
14.2	SQL-specific joins	108
14.2.1	Implementation Details	108
14.2.2	Specifying complex ON logic	108

15 Model Aggregates	111
15.1 Grouping	111
16 Hooks	113
16.1 Model Operation Hooks	113
16.1.1 Example with beforeSave	114
16.1.2 Arguments	114
16.1.3 Interrupting	114
16.1.4 Insert/Update Hooks	115
16.1.5 beforeSave, afterSave Hook	115
16.1.6 Loading, Deleting	115
16.1.7 Hook execution sequence	116
16.1.8 How to prevent actions	116
16.1.9 onRollback Hook	116
16.2 Persistence Hooks	117
16.2.1 Persistence\Sql	117
16.3 Other Hooks:	118
17 Advanced Topics	119
17.1 SubTypes	119
17.1.1 Best practice for specifying relation type	119
17.1.2 Type substitution on loading	120
17.2 Audit Fields	120
17.3 Soft Delete	122
17.3.1 Soft Delete that overrides default delete()	124
17.4 Creating Unique Field	125
17.5 Using WITH cursors	126
17.6 Creating Many to Many relationship	126
17.6.1 1. Create Intermediate Entity - InvoicePayment	126
17.6.2 2. Update Invoice and Payment model	127
17.6.3 3. How to use	127
17.7 Creating Related Entity Lookup	128
17.8 Inserting Hierarchical Data	130
17.9 Related Record Conditioning	131
17.10 Narrowing Down Existing References	131
18 Loading and Saving CSV Files	133
18.1 Setting Up	133
18.2 Exporting and Importing data from CSV	133
19 Agile DSQL Documentation	135
19.1 Overview	135
19.1.1 Goals of DSQL	135
19.1.2 DSQL by example	136
19.1.3 DSQL is Part of Agile Toolkit	136
19.2 Quickstart	136
19.2.1 Basic Concepts	136
19.2.2 Getting Started	137
19.2.3 Creating Objects and PDO	138
19.2.4 Query Building	139
19.2.5 Query Mode	139
19.2.6 Fetching Result	139
19.3 Connection	140
19.4 Expressions	141
19.4.1 Properties, Arguments, Parameters	142

19.4.2	Creating Expression	143
19.4.3	Expression Template	143
19.4.4	Nested expressions	143
19.4.5	Rendering	144
19.4.6	Executing Expressions	144
19.4.7	Magic an Debug Methods	145
19.4.8	Escaping Methods	145
19.4.9	Other Properties	146
19.5	Queries	146
19.5.1	Method invocation principles	147
19.5.2	Query Modes	147
19.5.3	Chaining	148
19.5.4	Using query as expression	148
19.5.5	Modifying Select Query	149
19.5.6	Insert and Replace query	157
19.5.7	Update Query	158
19.5.8	Delete Query	158
19.5.9	Dropping attributes	158
19.5.10	Other Methods	159
19.5.11	Properties	160
19.6	Results	160
19.7	Transactions	161
19.8	Advanced Topics	161
19.8.1	Advanced Connections	162
19.8.2	Extending Query Class	163
19.8.3	Manual Query Execution	165
19.8.4	Exception Class	165
19.9	Vendor support and Extensions	165
19.9.1	3rd party vendor support	166
20	Indices and tables	167
	PHP Namespace Index	169
	Index	171

Contents:

OVERVIEW

Agile Data is a unique SQL/NoSQL access library that promotes correct Business Logic design in your PHP application and implements database access in a flexible and scalable way.

Agile Data

PHP Business Logic Framework



1.1 Simple to learn

We have designed Agile Data to be very friendly for those who started programming recently and teach them correct patterns through clever architectural design.

Agile Data carries the spirit of PHP language in general and gives developer ability to make choices. The framework can be beneficial even in small applications, but the true power of Agile Data is realized when it's paired with Agile UI or Agile API projects. (<https://github.com/atk4/ui>, <https://github.com/atk4/api>).

1.2 Not a traditional ORM

Agile Data implementation has several significant differences to a traditional ORM (Hibernate / Doctrine style). I will discuss those in more detail further in documentation, however it's important to note the reason of not following ORM pattern:

- More suitable for mapping remote databases
- Give developer control over generated queries
- Better support for Persistence-specific features (e.g. SQL expressions)
- True many-to-many deep traversal and avoiding (explicit eager pre-loading)
- Better aggregation abstraction

To find out more, how Agile Data compares to other PHP data mappers and ORM frameworks, see <https://medium.com/@romaninsh/objectively-comparing-orm-dal-libraries-e4f095de80b5>

1.3 Concern Separation

Design of Agile Data follows principle of “concern separation”, but all of the basic functionality is divided into 3 major areas:

- Fields (or Columns)
- DataSets (or Rows)
- Databases (or Persistencies)

Each of the above corresponds to a PHP class, which may use composition principle to hide implementation details.

By design, you will be able to mix and match any any Field with any Database to work with your DataSets.

If you have worked with other ORMs, read the following sections to avoid confusion:

1.3.1 Class: Field

- Represent logical data column (e.g. “date_of_birth”)
- Stores column meta-information (e.g. ['type' => 'date', 'caption' => 'Birth Date'])
- Handles value normalization
- Documentation: *Field*

Note: Meta-information may be a persistence detail, (*Field::\$actual*) or presentation detail (*Field::\$ui*). Field class does not interpret the value, it only stores it.

1.3.2 Class: Model

- Represent logical Data Set (e.g. Active Users)
- Stores data location and criteria
- Stores list of Fields
- Stores individual row
- Handle operations over single or all records from Data Set
- Documentation: *Model*

Note: Model object is defined in such a way to contain enough information to fully provide all information for generic UI, or generic API, and generic persistence implementations.

Note: Unlike ORMs Model instances are never created during iterating. Also, in most cases, you never instantiate multiple instances of a model class.

1.3.3 Class: Persistence

- Represent external data storage (e.g. MySQL database)
- Stores connection information
- Translate single or multi-record operations into vendor-specific language
- Type-casts standard data types into vendor-specific format
- Documentation: Persistence

1.4 Code Layers

How is code:

```
select name from user where id = 1
```

is different to the code?:

```
$user->load(1)->get('name');
```

While both achieve similar things, the SQL-like code is what we call “persistence-specific” code. The second example is “domain model” code. The job of Agile Data is to map “domain model” code into “persistence-specific” code.

The design and features of Agile Data allow you to perform wider range of operations, be more expressive and efficient while remaining in “domain model”.

In normal application, all the database operations can be expressed in domain model without any degradation in performance due to large data volume or higher database latency.

It's typical for a web application that uses Agile Data in “domain model” to execute no more than 3-4 requests per page even for highly complex data pages (such as dashboards) and without use of stored procedures.

Next I'll show you how code from different “code layers” may look like:

1.4.1 Domain-Model Code

Code is unaware of physical location of your data or which persistence are you using:

```
$user = new User($db);

$user = $user->load(20); // load specific user record into PHP
echo $user->get('name') . ': '; // access field values

$gross = $user->ref('Invoice')
    ->addCondition('status', 'due')
    ->ref('Lines')
    ->action('sum', 'gross')
    ->getOne();

// get sum of all gross fields for due invoices
```

Another important aspect of Domain-model code is that fields such as `gross` or `name` can be either a physical values in the database or can be mapped to expressions (such as `vat + net`).

A typical method of your model class will be written in “domain-model” code.

Note: the actual execution and number of queries may vary based on capabilities of persistence. The above example executes a total of 2 queries if used with SQL database.

1.4.2 Persistence-specific code

This is a type of code which may change if you decide to switch from one persistence to another. For example, this is how you would define `gross` field for SQL:

```
$model->addExpression('gross', ['expr' => '[net] + [vat]']);
```

If your persistence does not support expressions (e.g. you are using Redis or MongoDB), you would need to define the field differently:

```
$model->addField('gross');
$model->onHook(Model::HOOK_BEFORE_SAVE, function (Model $entity) {
    $entity->set('gross', $entity->get('net') + $entity->get('vat'));
});
```

When you use persistence-specific code, you must be aware that it will not map into persistencies that does not support features you have used.

In most cases that is OK as if you prefer to stay with same database type, for instance, the above expression will still be usable with any SQL vendor, but if you want it to work with NoSQL, then your solution might be:

```

if ($model->hasMethod('addExpression')) {
    // method is injected by Persistence
    $model->addExpression('gross', ['expr' => '[net] + [vat]']);
} else {
    // persistence does not support expressions
    $model->addField('gross');
    $model->onHook(Model::HOOK_BEFORE_SAVE, function (Model $entity) {
        $entity->set('gross', $entity->get('net') + $entity->get('vat'));
    });
}

```

1.4.3 Generic Persistence-code

A final type of code is also persistence-specific, but it is agnostic to your data-model. The example would be implementation of aggregation with “GROUP BY” feature in SQL.

<https://github.com/atk4/report/blob/develop/src/GroupModel.php>

This code is specific to SQL databases, but can be used with any Model, so in order to use grouping with Agile Data, your code would be:

```

$aggregate = new AggregateModel(new Sale($db));
$aggregate->setGroupBy(['contractor_to', 'type'], [ // groups by 2 columns
    'c' => ['expr' => 'count(*)'], // defines aggregate formulas for fields
    'qty' => ['expr' => 'sum([])'], // [] refers back to qty
    'total' => ['expr' => 'sum([amount])'], // can specify any field here
]);

```

1.5 Persistence Scaling

Although in most cases you would be executing operation against SQL persistence, Agile Data makes it very easy to use models with a simpler persistencies.

For example, consider you want to output a “table” to the user using HTML by using Agile UI:

```

$table = \Atk4\Ui\Table::addTo($app);

$table->setModel(new User($db));

echo $table->render();

```

Class `\Atk4\Ui\Table` here is designed to work with persistencies and models - it will populate columns of correct type, fetch data, calculate totals if needed. But what if you have your data inside an array? You can use `Persistence\Static_` for that:

```

$table = \Atk4\Ui\Table::addTo($app);

$table->setModel(new User(new Persistence\Static_(
    ['name' => 'John', 'is_admin' => false, 'salary' => 34_400.0],
    ['name' => 'Peter', 'is_admin' => false, 'salary' => 42_720.0],
)));

```

(continues on next page)

(continued from previous page)

```
echo $table->render();
```

Even if you don't have a model, you can use Static persistence with Generic model class to display VAT breakdown table:

```
$table = \Atk4\Ui\Table::addTo($app);

$table->setModel(new Model(new Persistence\Static_(
    ['VAT_rate' => '12.0%', 'VAT' => 36.0, 'Net' => 300.0],
    ['VAT_rate' => '10.0%', 'VAT' => 52.0, 'Net' => 520.0],
]));

echo $table->render();
```

It can be made even simpler:

```
$table = \Atk4\Ui\Table::addTo($app);

$table->setModel(new Model(new Persistence\Static_(
    'John',
    'Peter',
]));

echo $table->render();
```

Agile UI even offers a wrapper for static persistence:

```
$table = \Atk4\Ui\Table::addTo($app);

$table->setSource(['John', 'Peter']);

echo $table->render();
```

QUICKSTART

Agile Data Framework is built around some unique concepts. Your knowledge of other ORM, ActiveRecord and QueryBuilder tools could be helpful, but you should carefully go through the basics if you want to know how to use Agile Data efficiently.

The distinctive goal for Agile Data is ability to “execute” complex operations on the database server directly, such as aggregation, sub-queries, joins and unions but only if the database server supports those operations.

Developer would normally create a declaration like this:

```
$user->hasMany('Order', ['model' => [Order::class]])  
->addField('total', ['aggregate' => 'sum']);
```

It is up to Agile Data to decide what's the most efficient way to implement the aggregation. Currently only SQL persistence is capable of constructing aggregate sub-query.

2.1 Requirements

If you wish to try out some examples in this guide, you will need the following:

- PHP 7.4 or above
- any of supported database - Sqlite, MySQL/MariaDB, PostgreSQL, MSSQL or Oracle

2.2 Core Concepts

- Business Model (see *Model*)

You define business logic inside your own classes that extend *Model*. Each class you create represent one business entity.

Model has 3 major characteristic: Business Logic definition, DataSet mapping and Active Record.

See: *Model*

- Persistence (see *Loading and Saving (Persistence)*)

Object representing a connection to database. Linking your Business Model to a persistence allows you to load/save individual records as well as execute multi-record operations (Actions)

For developer, persistence should be a secondary concern, after all it is possible to switch from one persistence to another and compensate for the feature differences without major refactoring.

- DataSet (see *Conditions and DataSet*)

A set of physical records stored on your database server that correspond to the Business Model.

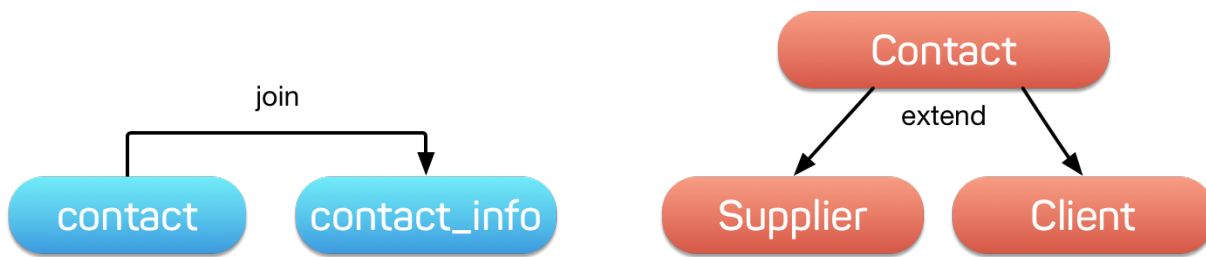
- Active Record (see *Setting and Getting active record data*)

Model can load individual record from DataSet, work with it and save it back into DataSet. While the record is loaded, we call it an Active Record.

- Action (see *Actions*)

Operation that Model performs on all of DataSet records without loading them individually. Actions have 3 main purposes: data aggregation, referencing and multi-record operations.

2.2.1 Persistence Domain vs Business Domain



It is very important to understand that there are two “domains” when it comes to your data. If you have used ORM, ActiveRecord or QueryBuilders, you will be thinking in terms of “Persistence Domain”. That means that you think in terms of “tables”, “fields”, “foreign keys” and “group by” operations.

In larger application developers does not necessarily have to know the details of your database structure. In fact - structure can often change and code that depend on specific field names or types can break.

More importantly, if you decide to store some data in different database either for caching (memcache), unique features (full-text search) or to handle large amounts of data (BigData) you suddenly have to carefully consider that in your application.

Business Domain is a layer that is designed to hide all the logic of data storage and focus on representing your business model in great detail. In other words - Business Logic is an API you and the rest of your developer team can use without concerning about data storage.

Agile Data has a rich set of features to define how Business Domain maps into Persistence Domain. It also allows you to perform most actions with only knowledge of Business Domain, keeping the rest of your application independent from your database choice, structure or patterns.

2.2.2 Class vs In-Line definition

Business model entity in Agile Data is represented through PHP object. While it is advisable to create each entity in its own class, you do not have to do so.

It might be handy to use in-line definition of a model. Try the following inside console:

```
$m = new Model($db, 'contact_info');  
$m->addField('address_1');  
$m->addField('address_2');  
$m->addCondition('address_1', '!=', null);  
$m = $m->loadAny();
```

(continues on next page)

(continued from previous page)

```
$m->get();
$m->executeCountQuery(); // same as ((int) $m->action('count')->getOne())
```

Next, exit and create file `src/Model_ContactInfo.php`:

```
<?php
class Model_ContactInfo extends Model
{
    public $table = 'contact_info';

    protected function init(): void
    {
        parent::init();

        $this->addField('address_1');
        $this->addField('address_2');
        $this->addCondition('address_1', '!=', null);
    }
}
```

Save, exit and run console again. You can now type this:

```
$m = new Model_ContactInfo($db);
$m = $m->loadAny();
$m->get();
```

Note: Should the “addCondition” be located inside model definition or inside your inline code? To answer this question - think - would `Model_ContactInfo` have application without the condition? If yes then either use `addCondition` in-line or create 2 classes.

2.2.3 Model State

When you create a new model object, you can change its state to perform various operations on your data. The state can be broken down into the following categories:

Persistence

When you create instance of a model (`new Model()`) you need to specify `Persistence` as a parameter. If you don't you can still use the model, but it won't be able to `Model::load()` or `Model::save()` data.

Once model is associated with one persistence, you cannot re-associate it. Method `Model::init()` will be executed only after persistence is known, so that method may make some decisions based on chosen persistence. If you need to store model inside a different persistence, this is achieved by creating another instance of the same class and copying data over. You must however remember that any fields that you have added in-line will not be recreated.

DataSet (Conditions)

Model object may have one or several conditions applied. Conditions will limit which records model can load (make active) and save. Once the condition is added, it cannot be removed for safety reasons.

Suppose you have a method that converts DataSet into JSON. Ability to add conditions is your way to specify which records to operate on:

```
public function myexport(\Atk4\Data\Model $m, array $fields = null)
{
    return json_encode($m->export($fields));
}

$m = new Model_User($db);
$m->addCondition('country_id', '2');

myexport($m, ['id', 'username', 'country_id']);
```

If you want to temporarily add conditions, then you can either clone the model or use `Model::loadBy`.

Active Record

Active Record is a third essential piece of information that your model stores. You can load / unload records like this:

```
$m = new Model_User($db);
$m = $m->loadAny();

$m->get(); // inside console, this will show you what's inside your model

$m->set('email', 'test@example.com');
$m->save();
```

You can call `$m->isLoading()` to see if there is active record and `$m->getId()` will store the ID of active record. You can also un-load the record with `$m->unload()`.

By default no records are loaded and if you modify some field and attempt to save unloaded model, it will create a new record.

Model may use some default values in order to make sure that your record will be saved inside DataSet:

```
$m = new Model_User($db);
$m->addCondition('country_id', 2);
$m->set('username', 'peter');
$m->save();

$m->get(); // will show country_id as 2
$m->set('country_id', 3);
$m->save(); // will generate exception because model you try to save doesn't match
↳ conditions set
```

Other Parameters

Apart from the main 3 pieces of “state” your Model holds there can also be some other parameters such as:

- order
- limit
- onlyFields

You can also define your own parameters like this:

```
$m = new Model_User($db, ['audit' => false]);
$m->audit
```

This can be used internally for all sorts of decisions for model behavior.

2.3 Getting Started

It’s time to create the first Model. Open `src/Model_User.php` which should look like this:

```
<?php
class Model_User extends Model
{
    public $table = 'user';

    protected function init(): void
    {
        parent::init();

        $this->addField('username');
        $this->addField('email');

        $j = $this->join('contact_info', 'contact_info_id');
        $j->addField('address_1');
        $j->addField('address_2');
        $j->addField('address_3');
        $j->hasOne('country_id', ['model' => [Model_Country::class]]);
    }
}
```

Extend either the base Model class or one of your existing classes (like Model_Client). Define \$table property unless it is already defined by parent class. All the properties defined inside your model class are considered “default” you can re-define them when you create model instances:

```
$m = new Model_User($db, 'user2'); // will use a different table
$m = new Model_User($db, ['table' => 'user2']); // same
```

Note: If you’re trying those lines, you will also have to create this new table inside your MySQL database:

```
create table user2 as select * from user
```

As I mentioned - `Model::init` is called when model is associated with persistence. You could create model and associate it with persistence later:

```
$m = new Model_User();  
  
$m->setPersistence($db); // calls $m->invokeInit()
```

You cannot add conditions just yet, although you can pass in some of the defaults:

```
$m = new Model_User(null, ['table' => 'user2']);  
  
$m->setPersistence($db); // will use table user2
```

2.3.1 Adding Fields

Methods `Model::addField()` and `Model::addFields()` can declare model fields. You need to declare them before you are able to use. You might think that some SQL reverse-engineering could be good at this point, but this would mimic your business logic after your presentation logic, while the whole point of Agile Data is to separate them, so you should, at least initially, avoid using generators.

In practice, `Model::addField()` creates a new 'Field' object and then links it up to your model. This object is used to store some information about your field, but it also participates in some field-related activity.

2.3.2 Table Joins

Similarly, `Model::join()` creates a Join object and stores it in `$j`. The Join object defines a relationship between the master `Model::$table` and some other table inside persistence domain. It makes sure relationship is maintained when objects are saved / loaded:

```
$j = $this->join('contact_info', 'contact_info_id');  
$j->addField('address_1');  
$j->addField('address_2');
```

That means that your business model will contain 'address_1' and 'address_2' fields, but when it comes to storing those values, they will be sent into a different database table and the records will be automatically linked.

Lets once again load up the console for some exercises:

```
$m = new Model_User($db);  
  
$m = $m->loadBy('username', 'john');  
$m->get();
```

At this point you'll see that address has also been loaded for the user. Agile Data makes management of related records transparent. In fact you can introduce additional joins depending on class. See classes `Model_Invoice` and `Model_Payment` that join table `document` with either `payment` or `invoice`.

As you load or save models you should see actual queries in the console, that should give you some idea what kind of information is sent to the database.

Adding Fields, Joins, Expressions and References creates more objects and 'adds' them into Model (to better understand how Model can behave like a container for these objects, see [documentation on Agile Core Containers](#)). This archi-

texture of Agile Data allows database persistence to implement different logic that will properly manipulate features of that specific database engine.

2.3.3 Understanding Persistence

To make things simple, console has already created persistence inside variable `$db`. Load up `console.php` in your editor to look at how persistence is set up:

```
$app->db = Persistence::connect($dsn, $user, $pass);
```

The `$dsn` can also be using the PEAR-style DSN format, such as: “mysql://user:pass@db/host”, in which case you do not need to specify `$user` and `$pass`.

For some persistence classes, you should use constructor directly:

```
$array = [
    1 => ['name' => 'John'],
    2 => ['name' => 'Peter'],
];

$db = new Persistence\Array_($array);
$m = new Model($db);
$m->addField('name');
$m = $m->load(2);
echo $m->get('name'); // Peter
```

There are several Persistence classes that deal with different data sources. Lets load up our console and try out a different persistence:

```
$a = [
    'user' => [],
    'contact_info' => [],
];

$ar = new Persistence\Array_($a);
$m = new Model_User($ar);
$m->set('username', 'test');
$m->set('address_1', 'street');

$m->save();

var_dump($a); // shows you stored data
```

This time our `Model_User` logic has worked pretty well with Array-only persistence logic.

Note: Persisting into Array or MongoDB are not fully functional as of 1.0 version. We plan to expand this functionality soon, see our development [roadmap](#).

2.4 References between Models

Your application normally uses multiple business entities and they can be related to each-other.

Warning: Do not mix-up business model references with database relations (foreign keys).

References are defined by calling `Model::hasOne()` or `Model::hasMany()`. You always specify destination model and you can optionally specify which fields are used for conditioning.

2.4.1 One to Many

Launch up console again and let's create reference between 'User' and 'System'. As per our database design - one user can have multiple 'system' records:

```
$m = new Model_User($db);
$m->hasMany('System', ['model' => [Model_System::class]]);
```

Next you can load a specific user and traverse into System model:

```
$m = $m->loadBy('username', 'john');
$s = $m->ref('System');
```

Unlike most ORM and ActiveRecord implementations today - instead of returning array of objects, `Model::ref()` actually returns another Model to you, however it will add one extra Condition. This type of reference traversal is called "Active Record to DataSet" or One to Many.

Your Active Record was user john and after traversal you get a model with DataSet corresponding to all Systems that belong to user john. You can use the following to see number of records in DataSet or export DataSet:

```
$s->isLoading();
$s->executeCountQuery();
$s->export();
$s->action('count')->getDebugQuery();
```

2.4.2 Many to Many

Agile Data also supports another type of traversal - 'DataSet to DataSet' or Many to Many:

```
$c = $m->ref('System')->ref('Client');
```

This will create a Model_Client instance with a DataSet corresponding to all the Clients that are contained in all of the Systems that belong to user john. You can examine the this model further:

```
$c->isLoading();
$c->executeCountQuery();
$c->export();
$c->action('count')->getDebugQuery();
```

By looking at the code - both MtM and OtM references are defined with 'hasMany'. The only difference is the loaded() state of the source model.

Calling `ref()->ref()` is also called Deep Traversal.

2.4.3 One to One

The third and final reference traversal type is “Active Record to Active Record”:

```
$cc = $m->ref('country_id');
```

This results in an instance of `Model_Country` with Active Record set to the country of user john:

```
$cc->isLoaded();
$cc->getId();
$cc->get();
```

2.4.4 Implementation of References

When reference is added using `Model::hasOne()` or `Model::hasMany()`, the new object is created and added into Model of class `Reference\HasMany` or `Reference\HasOne` (or `Reference\HasOneSql` in case you use SQL database). The object itself is quite simple and you can fetch it from the model if you keep the return value of `hasOne()` / `hasMany()` or call `Model::getReference()` with the same identifier later on. You can also use `Model::hasReference()` to check if reference exists in model.

Calling `Model::ref()` will proxy into the `ref()` method of reference object which will in turn figure out what to do.

Additionally you can call `Model::addField()` on the reference model that will bring one or several fields from related model into your current model.

Finally this reference object contains method `Reference::getModel()` which will produce a (possibly) fresh copy of related entity and will either adjust it's `DataSet` or set the active record.

2.5 Actions

Since NoSQL databases will always have some specific features, Agile Data uses the concept of ‘action’ to map into vendor-specific operations.

2.5.1 Aggregation actions

SQL implements methods such as `sum()`, `count()` or `max()` that can offer you some basic aggregation without grouping. This type of aggregation provides some specific value from a data-set. SQL persistence implements some of the operations:

```
$m = new Model_Invoice($db);
$m->executeCountQuery();
$m->action('fx', ['sum', 'total'])->getOne();
$m->action('fx', ['max', 'shipping'])->getOne();
```

Aggregation actions can be used in Expressions with `hasMany` references and they can be brought into the original model as fields:

```
$m = new Model_Client($db);
$m->getReference('Invoice')->addField('max_delivery', ['aggregate' => 'max', 'field' =>
  ↳ 'shipping']);
```

(continues on next page)

(continued from previous page)

```
$m->getReference('Payment')->addField('total_paid', ['aggregate' => 'sum', 'field' =>
↪ 'amount']);
$m->export(['name', 'max_delivery', 'total_paid']);
```

The above code is more concise and can be used together with reference declaration, although this is how it works:

```
$m = new Model_Client($db);
$m->addExpression('max_delivery', ['expr' => $m->refLink('Invoice')->action('fx', ['max',
↪ 'shipping'])]);
$m->addExpression('total_paid', ['expr' => $m->refLink('Payment')->action('fx', ['sum',
↪ 'amount'])]);
$m->export(['name', 'max_delivery', 'total_paid']);
```

In this example calling `refLink` is similar to traversing reference but instead of calculating DataSet based on Active Record or DataSet it references the actual field, making it ideal for placing into sub-query which SQL action is using. So when calling like above, `action()` will produce expression for calculating max/sum for the specific record of Client and those calculation are used inside an `Expression()`.

Expression is a special type of read-only Field that uses sub-query or a more complex SQL expression instead of a physical field. (See *Expressions* and *References*)

2.5.2 Field-reference actions

Field referencing allows you to fetch a specific field from related model:

```
$m = new Model_Country($db);
$m->action('field', ['name'])->get();
$m->action('field', ['name'])->getDebugQuery();
```

This is useful with `hasMany` references:

```
$m = new Model_User($db);
$m->getReference('country_id')->addField('country', 'name');
$m = $m->loadAny();
$m->get(); // look for 'country' field
```

`hasMany::addField()` again is a short-cut for creating expression, which you can also build manually:

```
$m->addExpression('country', $m->refLink('country_id')->action('field', ['name']));
```

2.5.3 Advanced Use of Actions

Actions prove to be very useful in various situations. For instance, if you are looking to add a new user:

```
$m = new Model_User($db);
$m->set('username', 'peter');
$m->set('address_1', 'street 49');
$m->set('country', 'UK');
$m->save();
```

Normally this would not work, because `country` is read-only expression, however if you wish to avoid creating an intermediate select to determine ID for 'UK', you could do this:

```

$m = new Model_User($db);
$m->set('username', 'peter');
$m->set('address_1', 'street 49');
$m->set('country_id', (new Model_Country($db))->addCondition('name', 'UK')->action('field
↪', ['id']));
$m->save();

```

This way it will not execute any code, but instead it will provide expression that will then be used to lookup ID of 'UK' when inserting data into SQL table.

2.6 Expressions

Expressions that are defined based on Actions (such as aggregate or field-reference) will continue to work even without SQL (although might be more performance-expensive), however if you're stuck with SQL you can use free-form pattern-based expressions:

```

$m = new Model_Client($db);
$m->getReference('Invoice')->addField('total_purchase', ['aggregate' => 'sum', 'field' =>
↪ 'total']);
$m->getReference('Payment')->addField('total_paid', ['aggregate' => 'sum', 'field' =>
↪ 'amount']);

$m->addExpression('balance', ['expr' => '[total_purchase] + [total_paid]']);
$m->export(['name', 'balance']);

```

2.7 Conclusion

You should now be familiar with the basics of Agile Data. To find more information on specific topics, use the rest of the documentation.

Agile Data is designed in an extensive pattern - by adding more objects inside Model a new functionality can be introduced. The described functionality is never a limitation and 3rd party code or you can add features that Agile Data authors are not even considered.

INTRODUCTION TO ARCHITECTURAL DESIGN

Layering is one of the most common techniques that software designers use to break apart a complicated software system. A modern application would have three primary layers:

- Presentation - Display of information (HTML generation, UI, API or CLI interface)
- Domain - Logic that is the real point of the system
- Data Source - Communication with databases, messaging systems, transaction managers, other packages

A persistence mechanism is a way how you save the data from some kind of in-memory model to the database. Apart from data-bases modern system also use REST services or interact with caches or files to load/store data.

Due to implementation specifics of the various data sources, making a “universal” persistence logic that can store Domain objects efficiently is not a trivial task. Various frameworks implement “Active Record”, “ORM” and “Query Builder” patterns in attempts to improve data access.

The common problems when trying to simplify mapping of domain logic include:

- Performance
 - Traversing references where you deal with millions of related records
 - Executing multi-row database operation
- Reduced features
 - Inability to use vendor-specific features such as SQL expression syntax
 - Derive calculations from multi-row sub-selects
 - Tweak persistence-related operations
- Abstraction
 - Domain objects are often restricted by database schema
 - Difficult to use Domain objects without database connection (e.g. in Unit Tests)

Agile Data implements a fresh concepts that separates your Domain from persistence cleanly yet manages to solve problems mentioned above.

The concepts implemented by Agile Data framework may require some getting used to (especially if you used some traditional ORMs or Active Record implementations before).

Once you learn the concept behind Agile Data, you’ll be able to write “Domain objects” of your application with ease through a readable code and without impact on your application performance or feature restrictions.

3.1 The Domain Layer Scope

Agile Data is a framework that will allow you to define your Domain objects and will map them into database of your choice.

You can use Agile Data with SQL (PDO-compatible) vendors, NoSQL (MongoDB) or memory Arrays. Support for other database vendors can be added through add-ons.

3.1.1 The Danger of Raw Queries

If you still think that writing SQL queries is the most efficient way to work with database, you are probably not considering other disadvantages of this approach:

- Parameters you specify to a query need to be escaped
- Complex queries are more difficult to write and debug
- Various parts of your application may want to change query (soft-delete add-on?)
- Optimization in your database may impact your Domain logic and even presentation
- Changing your database vendor or storing object data in cache is harder
- Difficult to maintain code

There are more problems such as difficulty in unit-testing your Domain object code.

3.1.2 Purity levels of Domain code

Agile Data focuses on creating “patterns” that can live in “Domain” layer. There are three levels of code “purity”:

- Implement patterns for working with for Domain objects.
- Implement patterns for “persistence-backed” Domain objects.
- Implement extensions for “persisting”

Some of your code will focus on working with Domain object without any concern about “persistence”. A good example is “Validation”. When you Validate your Domain object you just need to check field values, you would not even care where data came from.

Most of your code, however, will assume existence of SOME “persistence”, but will not rely on anything specific. Calculating total amount of your shopping basked price is such an operation. Basket items are stored somewhere - array, SQL or NoSQL and all you need is to calculate sum(amount). You don’t even know how “amount” field is called in the database.

While most of relational mapping solutions would load all basket items, Agile Data performs same operations inside database if possible.

Finally - some of your code may rely of some specific database vendor features. Example would be defining an expression using “IF (expr, val1, val2)” expression for some field of Domain model or using stored procedure as the source instead of table.

Agile Data offers you ability to move as much code as possible to the level with highest “purity”, but even if you have to write chunk of SQL code, you can do it without compromising cross-vendor compatibility.

3.2 Domain Logic

When dealing with Domain logic, you work with a single object.

When we start developing a new application, we first decide on the Model structure. Think what models your application will use and how they are related. Do not think in terms of “tables”, but rather think in terms of “objects” and properties of those objects.

All of those model properties are “declared”.

3.2.1 Domain Models

Congratulations, you have just designed a model layer of your application. Remember that it had nothing to do with your database structure, right?

- Client
- Order
- Admin

A code to declare a model:

```
class Model_User extends \Atk4\Data\Model
{
}

class Model_Client extends Model_User
{
}

class Model_Admin extends Model_User
{
}

class Model_Order extends \Atk4\Data\Model
{
}
```

3.2.2 Domain Model Methods

Next we need to write down various “functions” your application would have to perform and attribute those to individual models. At the same time think about object inheritance.

- User
 - sendPasswordReminder()
- Client (extends User)
 - register()
 - checkout()
- Admin (extends User)
 - showAuditLog()

- Order

Code:

```
class Model_Client extends Model_User
{
    public function sendPasswordReminder()
    {
        mail($this->get('email'), 'Your password is: ' . $this->get('password'));
    }
}
```

At this stage you should not think about “saving” your entries. Think of your objects as if they would forever exist in your memory. Also don’t bother with basic actions such as adding new order or deleting order.

3.2.3 Domain Model Fields

Our next step is to define object fields (or properties). Remember that inheritance is at play here so you can take advantage of OOP:

- User
 - name, is_vip, email, password, password_change_date
- Client
 - phone
- Admin
 - permission_level
- Order
 - description, amount, is_paid

Those fields are not just mere “properties”, but have more “meta” information behind them and that’s why we call them “fields” and not “properties”. A typical field contain information about field name, caption, type, validation rules, persistence rules, presentation rules and more. Meta information is optional and it can be used by automated processes (such as presentation or persistence).

For instance, `is_paid` has a `type('boolean')` which means it will be stored as 1/0 in MySQL, but will use true/false in MongoDB. It will be displayed as a checkbox. Those decisions are made by the framework and will simplify your life, however if you want to do things differently, you will still be able to override default behavior.

Code to declare fields:

```
class Model_Order extends \Atk4\Data\Model
{
    protected function init(): void
    {
        parent::init();

        $this->addField('description');
        $this->addField('amount')->type('atk4_money');
        $this->addField('is_paid')->type('boolean');
    }
}
```

Code to access field values:

```
$order->set('amount', 1200.2);
```

3.2.4 Domain Model Relationship

Next - references. Think how those objects relate to each-other. Think in terms of “specific object” and not database relations. Client has many Orders. Order has one Client.

- User
 - hasMany(Client)
- Client
 - hasOne(User)

There are no “many-to-many” relationship in Domain Model because relationships work from a specific record, but more on that later.

Code (add inside `init()`):

```
class Model_Client extends Model_User
{
    protected function init(): void
    {
        parent::init();

        $this->hasMany('Order', ['model' => [Model_Order::class]]);
    }
}

class Model_Order extends \Atk4\Data\Model
{
    protected function init(): void
    {
        parent::init();

        $this->hasOne('Client', ['model' => [Model_Client::class]]);

        // addField declarations
    }
}
```

3.3 Persistence backed Domain Logic

Once we establish that Model object and set its persistence layer, we can start accessing it. Here is the code:

```
$order = new Model_Order();
// $order is not linked with persistence

$order = new Model_Order();
$order->setPersistence($db); // same as $order = new Model_Order($db)
// $order is associated with specific persistence layer $db
```

3.3.1 ID Field

Each object is stored with some unique identifier, so you can load and store object if you know it's ID:

```
$order = $order->load(20);  
$order->set('amount', 1200.2);  
$order->save();
```

3.4 Persistence-specific Code

Finally, some code may rely on specific features of your persistence layer.

3.4.1 Domain Model Expressions

A final addition to our Domain Model are expressions. Those are the “formulas” where the value cannot be changed directly, but is actually derived from other values.

- User
 - is_password_expired
- Client
 - amount_due, total_order_amount

Here field `is_password_expired` is the type of expression that is based on the field `password_change_date` and system date. In other words the value of this expression will be different depending on parameter outside of your app.

Field `amount_due` is a sum of amount for all Orders by specific User for which condition “is_paid=false” is met. `total_order_amount` is similar, however there is no condition on the order.

With all of the above we have finished our “Domain Model” declaration. We haven't done any assumptions on where and how data is stored, which vendor we are using or how we can ensure that expressions will operate.

This is, however, a good point for you to write the initial batch of the code.

Code:

```
class Model_User extends \Atk4\Data\Model  
{  
    protected function init(): void  
    {  
        parent::init();  
  
        $this->addField('password');  
        $this->addField('password_change_date');  
  
        $this->addExpression('is_password_expired', [  
            'expr' => '[password_change_date] < (NOW() - INTERVAL 1 MONTH)',  
            'type' => 'boolean',  
        ]);  
    }  
}
```

3.4.2 Persistence Hooks

Hooks can help you perform operations when object is being persisted:

```
class Model_User extends \Atk4\Data\Model
{
    protected function init(): void
    {
        parent::init();

        // add fields here

        $this->onHookShort(Model::HOOK_BEFORE_SAVE, function () {
            if ($this->isDirty('password')) {
                $this->set('password', encrypt_password($this->get('password')));
                $this->set('password_change_date', $this->expr('now()'));
            }
        });
    }
}
```

3.5 DataSet Declaration

So far we have only looked at a single record - one User or one Order. In practice our application must operate with multiple records.

DataSet is an object that represents collection of Domain model records that are persisted:

```
$order = new Model_Order($db);
$order = $order->load(10);
```

In scenario above we loaded a specific record. Agile Data creates a separate object/entity when loading by cloning the original object/model.

You can iterate over the DataSet:

```
$sum = 0;
foreach (new Model_Order($db) as $order) {
    $sum += $order->get('amount');
}
```

At this point, I'll jump ahead a bit and will show you an alternative code:

```
$sum = (new Model_Order($db))->fx0(['sum', 'amount']->getOne());
```

It will have same effect as the code above, but will perform operation of adding up all order amounts inside the database and save you a lot of CPU cycles.

3.6 Domain Conditions

If your database has 3 clients - 'Joe', 'Bill', and 'Steve' then the DataSet of "Client" has 3 records.

DataSet concept lives in "Domain Logic" therefore you can use it safely without worrying that you will introduce unnecessary bindings into persistence and break single-purpose principle of your objects:

```
foreach ($clients as $client) {
    // echo $client->get('name') . "\n";
}
```

The above is a Domain Model code. It will iterate through the DataSet of "Clients" and output 3 names. You can also "narrow down" your DataSet by adding a restriction:

```
$sum = 0;
foreach ((new Model_Order($db))->addCondition('is_paid', true) as $order) {
    $sum += $order->get('amount');
}
```

And again it's much more effective to do this on database side:

```
$sum = (new Model_Order($db))
->addCondition('is_paid', true)
->fx0(['sum', 'amount'])
->getOne();
```

3.7 Related DataSets

Next, let's look on the orders of specific user. How would you load orders of a specific user. Depending on your past experience you might think about "querying" Order table with condition on user_id. We can't do that, because "query", "table" and "user_id" are persistence details and we must keep them outside of business logic. Other ORM solution give you something like this:

```
$arrayOfOrders = $user->orders();
```

Unfortunately this has practical performance implications and scalability constraints. What if your user is having millions of orders? Even with lazy-loading, you will be operating with million "id" records.

Agile Data implements traversal as a simple operation that converts one DataSet into another:

```
$userModel->addCondition('is_vip', true);
$vipOrders = $userModel->ref('Order');

$sum = $vipOrders->fx0(['sum', 'amount'])->getOne();
```

The implementation of ref() is pretty powerful - \$userModel can address 3 users in the database and only 2 of those users are VIP. Typical ORM would require you to fetch all VIP records and then perform additional queries to find their orders.

Agile Data, however, perform traversal without accessing database at all. After ref() is executed, you have a new DataSet with a condition based on user sub-query. The actual implementation may be different depending on vendor, but Agile Data will prefer not to fetch list of "user_id"s without need.

3.7.1 Domain Model Actions

Persistence layer in Agile Data uses intelligent mapping of your Domain Logic into DatabaseVendor-specific operations.

To continue my example from above, I'll use a query method to calculate number of orders placed by VIP clients:

```
$vipOrderCount = $vipOrders->fx(['count'])->getOne();
```

This code will attempt to execute a single-query only, however the ability to optimize your request relies on the capabilities of database vendor. The actual database operation(s) might look like this on SQL database:

```
select count(*) from `order` where user_id in
  (select id from user where type = "user" and is_vip = 1)
```

While with MongoDB, the query could be different:

```
$ids = collections.client.find({'is_vip': true}).field('id');
return collections.order.find({'user_id': $ids}).count();
```

Finally the code above will work even if you use a simple Array as a data source:

```
$db = new \Atk4\Data\Persistence\Array_(
  'client' => [
    [
      'name' => 'Joe',
      'email' => 'joe@yahoo.com',
      'Orders' => [
        ['amount' => 10],
        ['amount' => 20],
      ],
    ],
    [
      'name' => 'Bill',
      'email' => 'bill@yahoo.com',
      'Orders' => [
        ['amount' => 35],
      ],
    ],
  ],
);
```

So getting back to the operation above, lets look at it in more details:

```
$vipOrderCount = $vipOrders->fx(['count'])->getOne();
```

While “\$vipOrders” is actually a DataSet, executing count() will cross you over into persistence layer. However this method is returning a new object, which is then executed when you call getOne(). For SQL persistencies it returns \Atk4\Data\Persistence\Sql\Query object, for example.

Even though for a brief moment you had your hands on a “database-vendor specific” object, you have immediately converted Action into an actual value. As result your code is universal and is not persistence-specific. In Agile Data we permit code like that in our Domain Model and we call it “Domain Model Action”.

Let me define this properly: Domain Model Action is an operation that can be executed in your Domain Model layer which assumes existence of SOME Persistence for your model, but not a specific one.

As long as your Domain Model is restricted to generic Domain Model Actions, it will not violate SRP (Single Responsibility Principle)

3.7.2 Unique Features of Persistence Layer

More often than not, your application is designed and built with a specific persistence layer in mind. If you are using SQL database, you want to

Before we talk “databases”, we must outline a few challenges:

- our business model described above should work with various database vendors
- we should be able to perform basic Unit tests on our domain logic
- single vs multiple records
- ..add more..

MODEL

class Atk4\Data\Model

Probably the most significant class in ATK Data - Model - acts as a Parent for all your entity classes:

```
class User extends \Atk4\Data\Model
```

You must define individual classes for all your business entities. Other frameworks may rely on XML or annotations, in ATK everything is defined inside your “Model” class through pure PHP (See Initialization below)

Once you create instance of your model class, it can be recycled. With a single object you can load/unload individual records (See Single record Operations below):

```
$m = new User($db);  
  
$entity = $m->load(3);  
$entity->set('note', 'just updating');  
$entity->save();  
  
$m = $m->load(8);  
...
```

and even perform operations on multiple records (See Persistence Actions below).

When data is loaded from associated Persistence, it is automatically converted into a native PHP type (such as DateTime object) through a process called Typecasting. Various rules apply when you set value for model fields (Normalization) or when data is stored into database that does support a field type (Serialization)

Furthermore, because you define Models as a class, it is very easy to introduce your own extensions which may include Hooks and Actions.

There are many advanced topics that ATK Data covers, such as References, Joins, Aggregation, SQL actions, Unions, Deep Traversal and Containment.

The design is also very extensible allowing you to introduce new Field types, Join strategies, Reference patterns, Action types.

I suggest you to read the next section to make sure you fully understand the Model and its role in ATK Data.

4.1 Understanding Model

Please understand that Model in ATK Data is unlike models in other data frameworks. The Model class can be seen as a “gateway” between your code and many other features of ATK Data.

For example - you may define fields and relations for the model:

```
$model->addField('age', ['type' => 'integer']);  
$model->hasMany('Children', ['model' => [Person::class]]);
```

Methods `addField` and `hasMany` will ultimately create and link model with a corresponding `Field` object and `Reference` object. Those classes contain the logic, but in 95% of the use-cases, you will not have to dive deep into them.

4.1.1 Model object = Data Set

From the moment when you create instance of your model class, it represents a `DataSet` - set of records that share some common traits:

```
$allUsers = new User($db); // John, Susan, Leo, Bill, Karen
```

Certain operations may “shrink” this set, such as adding conditions:

```
$maleUsers = $allUsers->addCondition('gender', 'M');  
  
send_email_to_users($maleUsers);
```

This essentially filters your users without fetching them from the database server. In my example, when I pass `$maleUsers` to the method, no records are loaded yet from the database. It is up to the implementation of `send_email_to_users` to load or iterate records or perhaps approach the data-set differently, e.g. execute multi-record operation.

Note that most operations on Model are mutating (meaning that in example above `$allUsers` will also be filtered and in fact, `$allUsers` and `$maleUsers` will reference same object. Use `clone` if you do not wish to affect `$allUsers`).

4.1.2 Model object = meta information

By design, Model object does not have direct knowledge of higher level objects or specific implementations. Still - Model will be a good place to deposit some meta-information:

```
$model->addField('age', ['ui' => ['caption' => 'Put your age here']]);
```

Model and Field class will simply store the “ui” property which may (or may not) be used by ATK UI component or some add-on.

4.1.3 Domain vs Persistence

When you declare a model Field you can also store some persistence-related meta-information:

```
// override how your persistence formats date field
$model->addField('date_of_birth', ['type' => 'date', 'persistence' => ['format' => 'Ymd
→']]);

// declare field which is not saved
$model->addField('secret', ['neverPersist' => true]);

// relocate into a different field
$model->addField('old_field', ['actual' => 'new_field']);

// or even into a different table
$model->join('new_table')
->addField('extra_field');
```

Model also has a property `$table`, which indicate name of default table/collection/file to be used by persistence. (Name of property is decided to avoid beginner confusion)

4.1.4 Good naming for a Model

Some parts of this documentation were created years ago and may use class notation: `Model_User`. We actually recommend you to use namespaces instead:

```
namespace yourapp\Model;

use \Atk4\Data\Model;

class User extends Model
{
    protected function init(): void
    {
        parent::init();

        $this->addField('name');

        $this->hasMany('Invoices', ['model' => [Invoice::class]]);
    }
}
```

PHP does not have a “class” type, so `Invoice::class` will translate into a string “`yourapp\Model\Invoice`” and is a most efficient way to specify related class name.

You way also use `new Invoice()` there but be sure not to specify any argument, unless you intend to use cross-persistence referencing (this is further explained in Advanced section)

4.2 Initialization

Atk4\Data\Model::init()

Method init() will automatically be called when your Model is associated with Persistence object. It is commonly used to declare fields, conditions, relations, hooks and more:

```
class Model_User extends Atk4\Data\Model
{
    protected function init(): void
    {
        parent::init();

        $this->addField('name');
        $this->addField('surname');
    }
}
```

You may safely rely on \$this->getPersistence() result to make choices:

```
if ($this->getPersistence() instanceof \Atk4\Data\Persistence\Sql) {
    // calculating on SQL server is more efficient!!
    $this->addExpression('total', ['expr' => '[amount] + [vat]']);
} else {
    $this->addCalculatedField('total', ['expr' => function (self $m) {
        return $m->get('amount') + $m->get('vat');
    }, 'type' => 'float']);
}
```

To invoke code from init() methods of ALL models (for example soft-delete logic), you use Persistence's "afterAdd" hook. This will not affect ALL models but just models which are associated with said persistence:

```
$db->onHook(Persistence::HOOK_AFTER_ADD, function (Persistence $persistence, Model
→$model) use ($acl) {
    $fields = $model->getFields();

    $acl->disableRestrictedFields($fields);
});

$invoice = new Invoice($db);
```

4.2.1 Fields

Each model field is represented by a Field object:

```
$model->addField('name');

var_dump($model->getField('name'));
```

Field handles many very valuable operations which would otherwise fall on the shoulders of developer (Read more here [Field](#))

```
Atk4\Data\Model::addField($name, $seed)
```

Creates a new field object inside your model (by default the class is 'Field'). The fields are implemented on top of Containers from Agile Core.

Second argument to addField() will contain a seed for the Field class:

```
$this->addField('surname', ['default' => 'Smith']);
```

You may also specify your own Field implementation:

```
$this->addField('amount_and_currency', [MyAmountCurrencyField::class]);
```

Read more about *Field*

```
Atk4\Data\Model::addFields(array $fields, $seed = [])
```

Creates multiple field objects in one method call. See multiple syntax examples:

```
$m->addFields(['name'], ['default' => 'anonymous']);

$m->addFields([
    'last_name',
    'login' => ['default' => 'unknown'],
    'salary' => ['type' => 'atk4_money', CustomField::class, 'default' => 100],
    ['tax', CustomField::class, 'type' => 'atk4_money', 'default' => 20],
    'vat' => new CustomField(['type' => 'atk4_money', 'default' => 15]),
]);
```

Read-only Fields

Although you may make any field read-only:

```
$this->addField('name', ['readOnly' => true]);
```

There are two methods for adding dynamically calculated fields.

```
Atk4\Data\Model::addExpression($name, $seed)
```

Defines a field as server-side expression (e.g. SQL):

```
$this->addExpression('total', ['expr' => '[amount] + [vat]']);
```

The above code is executed on the server (SQL) and can be very powerful. You must make sure that expression is valid for current `$this->getPersistence()`:

```
$product->addExpression('discount', ['expr' => $this->refLink('category_id')->fieldQuery(
    ↪ 'default_discount')]);
// expression as a sub-select from referenced model (Category) imported as a read-only
↪ field
// of $product model

$product->addExpression('total', ['expr' => 'if ([is_discounted], ([amount] +
    ↪ [vat])*[discount], [amount] + [vat])']);
// new "total" field now contains complex logic, which is executed in SQL
```

(continues on next page)

(continued from previous page)

```
$product->addCondition('total', '<', 10);
// filter products that cost less than 10.0 (including discount)
```

For the times when you are not working with SQL persistence, you can calculate field in PHP.

```
Atk4\Data\Model::addCalculatedField($name[, 'expr' => $callback ])
```

Creates new field object inside your model. Field value will be automatically calculated by your callback method right after individual record is loaded by the model:

```
$this->addField('term', ['caption' => 'Repayment term in months', 'default' => 36]);
$this->addField('rate', ['caption' => 'APR %', 'default' => 5]);

$this->addCalculatedField('interest', ['expr' => function (self $m) {
    return $m->calculateInterest();
}, 'type' => 'float']);
```

Important: always use argument `$m` instead of `$this` inside your callbacks. If model is to be cloned, the code relying on `$this` would reference original model, but the code using `$m` will properly address the model which triggered the callback.

This can also be useful for calculating relative times:

```
class MyModel extends Model
{
    use HumanTiming; // see https://stackoverflow.com/questions/2915864/php-how-to-find-
    ↪ the-time-elapsed-since-a-date-time

    protected function init(): void
    {
        parent::init();

        $this->addCalculatedField('event_ts_human_friendly', ['expr' => function (self
    ↪ $m) {
            return $this->humanTiming($m->get('event_ts'));
        }]);
    }
}
```

4.2.2 Actions

Another common thing to define inside `Model::init()` would be a user invocable actions:

```
class User extends Model
{
    protected function init(): void
    {
        parent::init();
    }
}
```

(continues on next page)

(continued from previous page)

```

        $this->addField('name');
        $this->addField('email');
        $this->addField('password');

        $this->addUserAction('send_new_password');
    }

    public function send_new_password()
    {
        // .. code here

        $this->save(['password' => .. ]);

        return 'generated and sent password to ' . $m->get('name');
    }
}

```

With a method alone, you can generate and send passwords:

```

$user = $user->load(3);
$user->send_new_password();

```

but using `$this->addUserAction()` exposes that method to the ATK UI widgets, so if your admin is using Crud, a new button will be available allowing passwords to be generated and sent to the users:

```

Crud::addTo($app)->setModel(new User($app->db));

```

Read more about `Model\UserAction`

4.2.3 Hooks

Hooks (behaviors) can allow you to define callbacks which would trigger when data is loaded, saved, deleted etc. Hooks are typically defined in `Model::init()` but will be executed accordingly.

There are countless uses for hooks and even more opportunities to use hook by all sorts of extensions.

Validation

Validation is an extensive topic, but the simplest use-case would be through a hook:

```

$this->addField('name');

$this->onHookShort(Model::HOOK_VALIDATE, function () {
    if ($this->get('name') === 'C#') {
        return ['name' => 'No sharp objects are allowed'];
    }
});

```

Now if you attempt to save object, you will receive `ValidationException`:

```
$model->set('name', 'Swift');
$model->saveAndUnload(); // all good

$model->set('name', 'C#');
$model->saveAndUnload(); // exception here
```

Other Uses

Other uses for model hooks are explained in *Hooks*

4.2.4 Inheritance

ATK Data models are really good for structuring hierarchically. Here is example:

```
class VipUser extends User
{
    protected function init(): void
    {
        parent::init();

        $this->addCondition('purchases', '>', 1000);

        $this->addUserAction('send_gift');
    }

    public function send_gift()
    {
        ...
    }
}
```

This introduces a new business object, which is a sub-set of User. The new class will inherit all the fields, methods and actions of “User” class but will introduce one new action - `send_gift`.

4.3 Associating Model with Database

After talking extensively about model definition, lets discuss how model is associated with persistence. In the most basic form, model is associated with persistence like this:

```
$m = new User($db);
```

If model was created without persistence `Model::init()` will not fire. You can explicitly associate model with persistence like this:

```
$m = new User();

// ...

$m->setPersistence($db); // links with persistence
```

Multiple models can be associated with the same persistence. Here are also some examples of static persistence:

```
$m = new Model(new Persistence\Static_(['john', 'peter', 'steve']));
$m = $m->load(1);
echo $m->get('name'); // peter
```

See `Persistence\Static_`

property `Atk4\Data\Model::$persistence`

Refers to the persistence driver in use by current model. Calling certain methods such as `save()`, `addCondition()` or `action()` will rely on this property.

property `Atk4\Data\Model::$persistenceData`

DO NOT USE: Array containing arbitrary data by a specific persistence layer.

property `Atk4\Data\Model::$table`

If `$table` property is set, then your persistence driver will use it as default table / collection when loading data. If you omit the table, you should specify it when associating model with database:

```
$m = new User($db, 'user');
```

This also overrides current table value.

Atk4\Data\Model::withPersistence(*\$persistence*)

Creates a duplicate of a current model and associate new copy with a specified persistence. This method is useful for moving model data from one persistence to another.

4.4 Populating Data

Atk4\Data\Model::insert(*\$row*)

Inserts a new record into the database and returns `$id`. It does not affect currently loaded record and in practice would be similar to:

```
$entity = $m->createEntity();
$entity->setMulti($row);
$entity->save();

return $entity;
```

The main goal for `insert()` method is to be as fast as possible, while still performing data validation. After inserting method will return cloned model.

Atk4\Data\Model::import(*\$data*)

Similar to `insert()` however works across array of rows. This method will not return any IDs or models and is optimized for importing large amounts of data.

The method will still convert the data needed and operate with joined tables as needed. If you wish to access tables directly, you'll have to look into `Persistence::insert($m, $data)`;

4.5 Working with selective fields

When you normally work with your model then all fields are available and will be loaded / saved. You may, however, specify that you wish to load only a sub-set of fields.

Atk4\Data\Model::setOnlyFields(\$fields)

Specify array of fields. Only those fields will be accessible and will be loaded / saved. Attempt to access any other field will result in exception.

Null restore to full set of fields. This will also unload active record.

property Atk4\Data\Model::\$onlyFields

Contains list of fields to be loaded / accessed.

4.6 Setting and Getting active record data

When your record is loaded from database, record data is stored inside the \$data property:

property Atk4\Data\Model::\$data

Contains the data for an active record.

Model allows you to work with the data of single a record directly. You should use the following syntax when accessing fields of an active record:

```
$m->set('name', 'John');
$m->set('surname', 'Peter');
// or
$m->setMulti(['name' => 'John', 'surname' => 'Peter']);
```

When you modify active record, it keeps the original value in the \$dirty array:

Atk4\Data\Model::set(\$field, \$value)

Set field to a specified value. The original value will be stored in \$dirty property.

Atk4\Data\Model::setMulti(\$fields)

Set multiple field values.

Atk4\Data\Model::setNull(\$field)

Set value of a specified field to NULL, temporarily ignoring normalization routine. Only use this if you intend to set a correct value shortly after.

Atk4\Data\Model::unset(\$field)

Restore field value to it's original:

```
$m->set('name', 'John');
echo $m->get('name'); // John

$m->_unset('name');
echo $m->get('name'); // original value is shown
```

This will restore original value of the field.

Atk4\Data\Model::get()

Returns one of the following:

- If value was set() to the field, this value is returned

- If field was loaded from database, return original value
- if field had default set, returns default
- returns null.

Atk4\Data\Model::isDirty()

Return true if field contains unsaved changes (dirty):

```
$m->isDirty('name'); // returns false
$m->set('name', 'Other Name');
$m->isDirty('name'); // returns true
```

```
if ($m->isDirty('name') || $m->isDirty('surname')) {
    $m->set('full_name', $m->get('name') . ' ' . $m->get('surname'));
}
```

When the code above is placed in beforeSave hook, it will only be executed when certain fields have been changed. If your recalculations are expensive, it's pretty handy to rely on "dirty" fields to avoid some complex logic.

property Atk4\Data\Model::\$dirty

Contains list of modified fields since last loading and their original values.

Atk4\Data\Model::hasField(\$field)

Returns true if a field with a corresponding name exists.

Atk4\Data\Model::getField(\$field)

Finds a field with a corresponding name. Throws exception if field not found.

Full example:

```
$m = new Model_User($db, 'user');

// fields can be added after model is created
$m->addField('salary', ['default' => 1000]);

echo $m->isDirty('salary'); // false
echo $m->get('salary'); // 1000

// next we load record from $db
$m = $m->load(1);

echo $m->get('salary'); // 2000 (from db)
echo $m->isDirty('salary'); // false, was not changed

$m->set('salary', 3000);

echo $m->get('salary'); // 3000 (changed)
echo $m->isDirty('salary'); // true

$m->_unset('salary'); // return to original value

echo $m->get('salary'); // 2000
echo $m->isDirty('salary'); // false

$m->set('salary', 3000);
```

(continues on next page)

```
$m->save();  
  
echo $m->get('salary'); // 3000 (now in db)  
echo $m->isDirty('salary'); // false
```

protected Atk4\Data\Model::normalizeFieldName()

Verify and convert first argument got get / set;

4.7 Title Field, ID Field and Model Caption

Those are three properties that you can specify in the model or pass it through defaults:

```
class MyModel ..  
    public ?string $titleField = 'full_name';
```

or as defaults:

```
$m = new MyModel($db, ['titleField' => 'full_name']);
```

4.7.1 ID Field

property Atk4\Data\Model::\$idField

If your data storage uses field different than `id` to keep the ID of your records, then you can specify that in `$idField` property.

ID value of loaded entity cannot be changed. If you want to duplicate a record, you need to create a new entity and save it.

4.7.2 Title Field

property Atk4\Data\Model::\$titleField

This field by default is set to `'name'` will act as a primary title field of your table. This is especially handy if you use model inside UI framework, which can automatically display value of your title field in the header, or inside drop-down.

If you don't have field `'name'` but you want some other field to be title, you can specify that in the property. If `titleField` is not needed, set it to `false` or point towards a non-existent field.

See `HasOne::addTitle()`

public Atk4\Data\Model::getTitle()

Return title field value of currently loaded record.

public Atk4\Data\Model::getTitles()

Returns array of title field values of all model records in format `[id => title]`.

4.7.3 Model Caption

property Atk4\Data\Model::\$caption

This is caption of your model. You can use it in your UI components.

public Atk4\Data\Model::getModelCaption()

Returns model caption. If caption is not set, then try to generate one from model class name.

4.8 Setting limit and sort order

public Atk4\Data\Model::setLimit(\$count, \$offset = null)

Sets limit on how many records to select. Will select only \$count records starting from \$offset record.

public Atk4\Data\Model::setOrder(\$field, \$desc = null)

Sets sorting order of returned data records. Here are some usage examples. All these syntaxes work the same:

```
$m->setOrder('name, salary desc');
$m->setOrder(['name', 'salary desc']);
$m->setOrder(['name', 'salary' => true]);
$m->setOrder(['name' => false, 'salary' => true]);
$m->setOrder([ ['name'], ['salary', 'desc'] ]);
$m->setOrder([ ['name'], ['salary', true] ]);
$m->setOrder([ ['name'], ['salary desc'] ]);
// and there can be many more similar combinations how to call this
```

Keep in mind - true means desc, desc means descending. Otherwise it will be ascending order by default.

You can also use \Atk4\Data\Persistence\Sql\Expression or array of expressions instead of field name here. Or even mix them together:

```
$m->setOrder($m->expr('[net] * [vat]'));
$m->setOrder([$m->expr('[net] * [vat]'), $m->expr('[closing] - [opening]')]);
$m->setOrder(['net', $m->expr('[net] * [vat]', 'ref_no')]);
```


TYPECASTING

Typecasting is evoked when you are attempting to save or load the record. Unlike strict types and normalization, typecasting is a persistence-specific operation. Here is the sequence and sample:

```
$m->addField('birthday', ['type' => 'date']);  
// Type has a number of pre-defined values. Using 'date'  
// instructs AD that we will be using it for storing dates  
// through DateTime class  
  
$m->set('birthday', 'Jan 1 1960');  
// If non-compatible value is provided, it will be converted  
// into a proper date through Normalization process. After  
// this line value of 'birthday' field will be DateTime.  
  
$m->save();  
// At this point typecasting converts the "DateTime" value  
// into UTC date-time representation for SQL or "MongoDate"  
// type if you're persisting with MongoDB. This does not affect  
// value of a model field.
```

Typecasting is necessary to save the values inside the database and restore them back just as they were before.

The purpose of a flexible typecasting system is to allow you to store your date in a compatible format or even fine-tune it to match your database settings (e.g. timezone) without affecting your domain code.

You must remember that type-casting is a two-way operation. If you are introducing your own types, you will need to make sure they can be saved and loaded correctly.

Some types such as boolean may support additional options like:

```
$m->addField('is_married', [  
    'type' => 'boolean',  
    'enum' => ['No', 'Yes'],  
]);  
  
$m->set('is_married', 'Yes'); // normalizes into true  
$m->set('is_married', true); // better way because no need to normalize  
  
$m->save(); // stores as "Yes" because of type-casting
```

5.1 Value types

Any type can have a value of `null`:

```
$m->set('is_married', null);  
if (!$m->get('is_married')) {  
    // either null or false  
}
```

If value is passed which is not compatible with field type, Agile Data will try to normalize value:

```
$m->addField('age', ['type' => 'integer']);  
$m->addField('name', ['type' => 'string']);  
  
$m->set('age', '49.8');  
$m->set('name', '    John');  
  
echo $m->get('age'); // 49 - normalization cast value to integer  
echo $m->get('name'); // 'John' - normalization trims value
```

5.1.1 Undefined type

If you do not set type for a field, Agile Data will not normalize and type-cast its value.

Because of the loose PHP types, you can encounter situations where undefined type is changed from `'4'` to `4`. This change is still considered “dirty”.

If you use numeric value with a type-less field, the response from SQL does not distinguish between integers and strings, so your value will be stored as “string” inside the model.

The same can be said about forms, which submit all their data through POST request that has no types, so undefined type fields should work relatively good with the standard setup of Agile Data + Agile Toolkit + SQL.

5.1.2 Type of IDs

Many databases will allow you to use different types for ID fields. In SQL the `'id'` column will usually be “integer”, but sometimes it can be of a different type.

5.1.3 Supported types

- `string` - for storing short strings, such as name of a person. Normalize will trim the value.
- `text` - for storing long strings, such as notes or description. Normalize will trim the value.
- `boolean` - normalize will cast value to boolean.
- `smallint`, `integer`, `bigint` - normalize will cast value to integer.
- `atk4_money` - normalize will round value with 4 digits after dot.
- `float` - normalize will cast value to float.
- `date` - normalize will convert value to DateTime object.
- `datetime` - normalize will convert value to DateTime object.

- `time` - normalize will convert value to `DateTime` object.
- `json` - no normalization by default
- `object` - no normalization by default

5.1.4 Types and UI

UI framework such as Agile Toolkit will typically rely on field type information to properly present data for views (forms and tables) without you having to explicitly specify the `ui` property.

5.1.5 Array and Object types

Some types may require serialization for some persistencies, for instance `'array'` and `'object'` types cannot be stored in SQL natively. `json` type can be used to store these in JSON.

This is handy when mapping JSON data into native PHP structures.

LOADING AND SAVING (PERSISTENCE)

`class Atk4\Data\Model`

Model object represents your real-life business objects such as “Invoice” or “Client”. The rest of your application works with “Model” objects only and have no knowledge of what database you are using and how data is stored in there. This decouples your app from the data storage (Persistence). If in the future you will want to change database server or structure of your database, you can do it without affecting your application.

Data Persistence frameworks (like ATK Data) provide the bridge between “Model” and the actual database. There is balance between performance, simplicity and consistency. While other persistence frameworks insist on strict isolation, ATK Data prefers practicality and simplicity.

ATK Data couples Model and Persistence, they have some intimate knowledge of each-other and work as a unit. Persistence object is created first and by the time Model is created, you specify persistence to the model.

During the lifecycle of the Model it can work with various records, save, load, unload data etc, but it will always remain linked with that same persistence.

6.1 Associating with Persistence

Create your persistence object first:

```
$db = \Atk4\Data\Persistence::connect($dsn);
```

There are two ways to link your model up with the persistence:

```
$m = new Model_Invoice($db);  
$m = new Model_Invoice();  
$m->setPersistence($db);
```

`Atk4\Data\Model::load()`

Load active record from the DataSet:

```
$m = $m->load(10);  
echo $m->get('name');
```

If record not found, will throw exception.

`Atk4\Data\Model::save($data = [])`

Store active record back into DataSet. If record wasn't loaded, store it as a new record:

```
$m = $m->load(10);  
$m->set('name', 'John');  
$m->save();
```

You can pass argument to save() to set() and save():

```
$m->unload();  
$m->save(['name' => 'John']);
```

Atk4\Data\Model::tryLoad()

Same as load() but will return null if record is not found:

```
$m = $m->tryLoad(10);
```

Atk4\Data\Model::unload()

Remove active record and restore model to default state:

```
$m = $m->load(10);  
$m->unload();  
  
$m->set('name', 'New User');  
$m->save(); // creates new user
```

Atk4\Data\Model::delete(\$id = null)

Remove current record from DataSet. You can optionally pass ID if you wish to delete a different record. If you pass ID of a currently loaded record, it will be unloaded.

6.1.1 Inserting Record with a specific ID

When you add a new record with save(), insert() or import, you can specify ID explicitly:

```
$m->set('id', 123);  
$m->save();  
  
// or $m->insert(['Record with ID=123', 'id' => 123]);
```

However if you change the ID for record that was loaded, then your database record will also have its ID changed. Here is example:

```
$m = $m->load(123);  
$m->setId(321);  
$m->save();
```

After this your database won't have a record with ID 123 anymore.

6.2 Type Converting

PHP operates with a handful of scalar types such as integer, string, booleans etc. There are more advanced types such as DateTime. Finally user may introduce more useful types.

Agile Data ensures that regardless of the selected database, types are converted correctly for saving and restored as they were when loading:

```
$m->addField('is_admin', ['type' => 'boolean']);
$m->set('is_admin', false);
$m->save();

// SQL database will actually store `0`

$m = $m->load();

$m->get('is_admin'); // converted back to `false`
```

Behind a two simple lines might be a long path for the value. The various components are essential and as developer you must understand the full sequence:

```
$m->set('is_admin', false);
$m->save();
```

6.2.1 Strict Types an Normalization

PHP does not have strict types for variables, however if you specify type for your model fields, the type will be enforced.

Calling “set()” or using array-access to set the value will start by casting the value to an appropriate data-type. If it is impossible to cast the value, then exception will be generated:

```
$m->set('is_admin', '1'); // ok, will be normalized/stored as `true`

$m->set('is_admin', 123); // throws exception
```

It’s not only the ‘type’ property, but ‘enum’ can also imply restrictions:

```
$m->addField('access_type', ['enum' => ['readOnly', 'full']]);

$m->set('access_type', 'full'); // ok
$m->set('access_type', 'half-full'); // throws exception
```

There are also non-trivial types in Agile Data:

```
$m->addField('salary', ['type' => 'atk4_money']);
$m->set('salary', 20); // converts to `20.00 EUR`

$m->addField('date', ['type' => 'date']);
$m->set('date', time()); // converts to DateTime class
```

Finally, you may create your own custom field types that follow a more complex logic:

```
$m->add(new Field_Currency(), 'balance');
$m->set('balance', 12_200.0);

// may transparently work with 2 columns: 'balance_amount' and 'balance_currency_id' for
↳ example
```

Loaded/saved data are always normalized unless the field value normalization is intercepted a hook.

Final field flag that is worth mentioning is called `Field::$readOnly` and if set, then value of a field may not be modified directly:

```
$m->addField('ref_no', ['readOnly' => true]);
$m = $m->load(123);

$m->get('ref_no'); // perfect for reading field that is populated by trigger

$m->set('ref_no', 'foo'); // exception
```

Note that `readOnly` can still have a default value:

```
$m->addField('created', [
    'readOnly' => true,
    'type' => 'datetime',
    'default' => new DateTime(),
]);

$m->save(); // stores creation time just fine and also will load it
```

Note: If you have been following our “Domain” vs “Persistence” then you can probably see that all of the above functionality described in this section apply only to the “Domain” model.

6.2.2 Typecasting

For full documentation on type-casting see [Typecasting](#)

6.2.3 Validation

Validation in application always depends on business logic. For example, if you want age field to be above 14 for the user registration you may have to ask yourself some questions:

- Can user store 12 inside a age field?
- If yes, Can user persist age with value of 12?
- If yes, Can user complete registration with age of 12?

If 12 cannot be stored at all, then exception would be generated during `set()`, before you even get a chance to look at other fields.

If storing of 12 in the model field is OK validation can be called from `beforeSave()` hook. This might be a better way if your validation rules depends on multiple field conditions which you need to be able to access.

Finally you may allow persistence to store 12 value, but validate before a user-defined operation. `completeRegistration` method could perform the validation. In this case you can create a confirmation page, that actually stores your in-complete registration inside the database.

You may also make a decision to store registration-in-progress inside a session, so your validation should be aware of this logic.

Agile Data relies on 3rd party validation libraries, and you should be able to find more information on how to integrate them.

6.2.4 Multi-column fields

Lets talk more about this currency field:

```
$m->add(new Field_Currency(), 'balance');
$m->set('balance', 12_200.0);
```

It may be designed to split up the value by using two fields in the database: `balance_amount` and `balance_currency_id`. Both values must be loaded otherwise it will be impossible to re-construct the value.

On other hand, we would prefer to hide those two columns for the rest of application.

Finally, even though we are storing “id” for the currency we want to make use of References.

Your `init()` method for a `Field_Currency` might look like this:

```
protected function init(): void
{
    parent::init();

    $this->neverPersist = true;

    $f = $this->shortName; // balance

    $this->getOwner()->addField(
        $f . '_amount',
        ['type' => 'atk4_money', 'system' => true]
    );

    $this->getOwner()->hasOne(
        $f . '_currency_id',
        [
            'model' => [Currency::class],
            'system' => true,
        ]
    );
}
```

There are more work to be done until `Field_Currency` could be a valid field, but I wanted to draw your attention to the use of field flags:

- `system` flag is used to hide `balance_amount` and `balance_currency_id` in UI.
- `neverPersist` flag is used because there are no `balance` column in persistence.

6.2.5 Dates and Time

There are 3 datetime formats supported:

- **date**: Converts into YYYY-MM-DD using UTC timezone for SQL. Defaults to DateTime() class in PHP, but supports string input (parsed as date in a current timezone) or unix timestamp.
- **time**: converts into HH:MM:SS using UTC timezone for storing in SQL. Defaults to DateTime() class in PHP, but supports string input (parsed as date in current timezone) or unix timestamp. Will discard date from timestamp.
- **datetime**: stores both date and time. Uses UTC in DB. Defaults to DateTime() class in PHP. Supports string input parsed by strtotime() or unix timestamp.

6.2.6 Customizations

Process which converts field values in native PHP format to/from database-specific formats is called *Typecasting*. Persistence driver implements a necessary type-casting through the following two methods:

`Atk4\Data\Model::typecastLoadRow($model, $row)`

Convert persistence-specific row of data to PHP-friendly row of data.

`Atk4\Data\Model::typecastSaveRow($model, $row)`

Convert native PHP-native row of data into persistence-specific.

Row persisting may rely on additional methods, such as:

`Atk4\Data\Model::typecastLoadField(Field $field, $value)`

Convert persistence-specific row of data to PHP-friendly row of data.

`Atk4\Data\Model::typecastSaveField(Field $field, $value)`

Convert native PHP-native row of data into persistence-specific.

6.3 Duplicating and Replacing Records

In normal operation, once you store a record inside your database, your interaction will always update this existing record. Sometimes you want to perform operations that may affect other records.

6.3.1 Create copy of existing record

`Atk4\Data\Model::duplicate($id = null)`

Normally, active record stores “id”, but when you call duplicate() it forgets current ID and as result it will be inserted as new record when you execute save() next time.

If you pass the \$id parameter, then the new record will be saved under a new ID:

```
// assume DB with only one record with ID = 123

// load and duplicate that record
$m->load(123)->duplicate()->save();

// now you have 2 records:
// one with ID = 123 and another with ID = {next db generated id}
echo $m->executeCountQuery();
```

6.3.2 Duplicate then save under a new ID

Assuming you have 2 different records in your database: 123 and 124, how can you take values of 123 and write it on top of 124?

Here is how:

```
$m->load(123)->duplicate()->setId(124)->save();
```

Now the record 124 will be replaced with the data taken from record 123. For SQL that means calling 'replace into x'.

Warning: There is no special treatment for joins() when duplicating records, so your new record will end up referencing the same joined record. If the join is reverse then your new record may not load.

This will be properly addressed in a future version of Agile Data.

6.4 Working with Multiple DataSets

When you load a model, conditions are applied that make it impossible for you to load record from outside of a data-set. In some cases you do want to store the model outside of a data-set. This section focuses on various use-cases like that.

6.4.1 Cloning versus New Instance

When you clone a model, the new copy will inherit pretty much all the conditions and any in-line modifications that you have applied on the original model. If you decide to create new instance, it will provide a vanilla copy of model without any in-line modifications.

6.4.2 Looking for duplicates

We have a model 'Order' with a field 'ref', which must be unique within the context of a client. However, orders are also stored in a 'Basket'. Consider the following code:

```
$basket->ref('Order')->insert(['ref' => 123]);
```

You need to verify that the specific client wouldn't have another order with this ref, how do you do it?

Start by creating a beforeSave handler for Order:

```
$this->onHookShort(Model::HOOK_BEFORE_SAVE, function () {
    if ($this->isDirty('ref')) {
        $m = (new static())
            ->addCondition('client_id', $this->get('client_id')) // same client
            ->addCondition($this->idField, '!=', $this->getId()) // has another order
            ->tryLoadBy('ref', $this->get('ref')) // with same ref
        if ($m !== null) {
            throw (new Exception('Order with ref already exists for this client'))
                ->addMoreInfo('client', $this->get('client_id'))
                ->addMoreInfo('ref', $this->get('ref'))
        }
    }
});
```

6.4.3 Archiving Records

In this use case you are having a model 'Order', but you have introduced the option to archive your orders. The method `archive()` is supposed to mark order as archived and return that order back. Here is the usage pattern:

```
$o->addCondition('is_archived', false); // to restrict loading of archived orders
$o = $o->load(123);
$archive = $o->archive();
$archive->set('note', $archive->get('note') . "\nArchived on $date.");
$archive->save();
```

With Agile Data API building it's quite common to create a method that does not actually persist the model.

The problem occurs if you have added some conditions on the `$o` model. It's quite common to use `$o` inside a UI element and exclude Archived records. Because of that, saving record as archived may cause exception as it is now outside of the result-set.

There are two approaches to deal with this problem. The first involves disabling after-save reloading:

```
public function archive()
{
    $this->getModel()->reloadAfterSave = false;
    $this->set('is_archived', true);

    return $this;
}
```

After-save reloading would fail due to `is_archived = false` condition so disabling reload is a hack to get your record into the database safely.

The other, more appropriate option is to re-use a vanilla Order record:

```
public function archive()
{
    $this->save(); // just to be sure, no dirty stuff is left over

    $archive = new static();
    $archive = $archive->load($this->getId());
    $archive->set('is_archived', true);

    $this->unload(); // active record is no longer accessible

    return $archive;
}
```

6.5 Working with Multiple Persistencies

Normally when you load the model and save it later, it ends up in the same database from which you have loaded it. There are cases, however, when you want to store the record inside a different database. As we are looking into use-cases, you should keep in mind that with Agile Data Persistence can be pretty much anything including 'RestAPI', 'File', 'Memcache' or 'MongoDB'.

Important: Instance of a model can be associated with a single persistence only. Once it is associated, it stays like

that. To store a model data into a different persistence, a new instance of your model will be created and then associated with a new persistence.

```
Atk4\Data\Model::withPersistence($persistence)
```

6.5.1 Creating Cache with Memcache

Assuming that loading of a specific items from the database is expensive, you can opt to store them in a MemCache. Caching is not part of core functionality of Agile Data, so you will have to create logic yourself, which is actually quite simple.

You can use several designs. I will create a method inside my application class to load records from two persistencies that are stored inside properties of my application:

```
public function loadQuick(Model $class, $id)
{
    // first, try to load it from MemCache
    $m = (clone $class)->setPersistence($this->mdb)->tryLoad($id);

    if ($m === null) {
        // fall-back to load from SQL
        $m = $this->sql->add(clone $class)->load($id);

        // store into MemCache too
        $m = $m->withPersistence($this->mdb)->save();
    }

    $m->onHook(Model::HOOK_BEFORE_SAVE, function (Model $entity) {
        $entity->withPersistence($this->sql)->save();
    });

    $m->onHook(Model::HOOK_BEFORE_DELETE, function (Model $entity) {
        $entity->withPersistence($this->sql)->delete();
    });

    return $m;
}
```

The above logic provides a simple caching framework for all of your models. To use it with any model:

```
$m = $app->loadQuick(new Order(), 123);

$m->set('completed', true);
$m->save();
```

To look in more details into the actual method, I have broken it down into chunks:

```
// first, try to load it from MemCache:
$m = (clone $class)->setPersistence($this->mdb)->tryLoad($id);
```

The `$class` will be an uninitialized instance of a model (although you can also use a string). It will first be associated with the MemCache DB persistence and we will attempt to load a corresponding ID. Next, if no record is found in the cache:

```
if ($m === null) {
    // fall-back to load from SQL
    $m = $this->sql->add(clone $class)->load($id);

    // store into MemCache too
    $m = $m->withPersistence($this->mdb)->save();
}
```

Load the record from the SQL database and store it into `$m`. Next, save `$m` into the MemCache persistence by replacing (or creating new) record. The `$m` at the end will be associated with the MemCache persistence for consistency with cached records. The last two hooks are in order to replicate any changes into the SQL database also:

```
$m->onHook(Model::HOOK_BEFORE_SAVE, function (Model $entity) {
    $entity->withPersistence($this->sql)->save();
});

$m->onHook(Model::HOOK_BEFORE_DELETE, function (Model $entity) {
    $entity->withPersistence($this->sql)->delete();
});
```

I have to note that `withPersistence()` transfers the dirty flags into a new model, so SQL record will be updated with the record that you have modified only.

If saving into SQL is successful the memcache persistence will be also updated.

6.5.2 Using Read / Write Replicas

In some cases your application have to deal with read and write replicas of the same database. In this case all the operations would be done on the read replica, except for certain changes.

In theory you can use hooks (that have option to cancel default action) to create a comprehensive system-wide solution, I'll illustrate how this can be done with a single record:

```
$m = new Order($readReplica);

$m->set('completed', true);

$m->withPersistence($writeReplica)->save();
$dirtyRef = &$m->getDirtyRef();
$dirtyRef = [];

// possibly the update is delayed
// $m->reload();
```

By changing 'completed' field value, it creates a dirty field inside `$m`, which will be saved inside a `$writeReplica`. Although the proper approach would be to reload the `$m`, if there is chance that your update to a write replica may not propagate to read replica, you can simply reset the dirty flags.

If you need further optimization, make sure `reloadAfterSave` is disabled for the write replica:

```
$m->withPersistence($writeReplica)->setDefaults(['reloadAfterSave' => false])->save();
```

or use:

```
$m->withPersistence($writeReplica)->saveAndUnload();
```

6.5.3 Archive Copies into different persistence

If you wish that every time you save your model the copy is also stored inside some other database (for archive purposes) you can implement it like this:

```
$m->onHook(Model::HOOK_BEFORE_SAVE, function (Model $entity) {
    $sarc = $this->withPersistence($entity->getApp()->archive_db);

    // add some audit fields
    $sarc->addField('original_id', ['type' => 'bigint'])->set($this->getId());
    $sarc->addField('saved_by')->set($this->getApp()->user);

    $sarc->saveAndUnload();
});
```

6.5.4 Store a specific record

If you are using authentication mechanism to log a user in and you wish to store his details into Session, so that you don't have to reload every time, you can implement it like this:

```
if (!isset($_SESSION['ad'])) {
    $_SESSION['ad'] = []; // initialize
}

$sess = new \Atk4\Data\Persistence\Array_($_SESSION['ad']);
$loggedUser = new User($sess);
$loggedUser = $loggedUser->load('active_user');
```

This would load the user data from Array located inside a local session. There is no point storing multiple users, so I'm using id='active_user' for the only user record that I'm going to store there.

How to add record inside session, e.g. log the user in? Here is the code:

```
$u = new User($db);
$u = $u->load(123);

$u->withPersistence($sess)->save();
```

6.6 Actions

Action is a multi-row operation that will affect all the records inside DataSet. Actions will not affect records outside of DataSet (records that do not match conditions)

```
Atk4\Data\Model::action($action, $args = [])
```

Prepares a special object representing "action" of a persistence layer based around your current model.

6.6.1 Action Types

Actions can be grouped by their result. Some action will be executed and will not produce any results. Others will respond with either one value or multiple rows of data.

- no results
- single value
- single row
- single column
- array of hashes

Action can be executed at any time and that will return an expected result:

```
$m = Model_Invoice();  
$val = (int) $m->action('count')->getOne(); // same as $val = $m->executeCountQuery()
```

Most actions are sufficiently smart to understand what type of result you are expecting, so you can have the following code:

```
$m = Model_Invoice();  
$val = $m->action('count')();
```

When used inside the same Persistence, sometimes actions can be used without executing:

```
$m = Model_Product($db);  
$m->addCondition('name', $productName);  
$productIdAction = $m->action('field', ['id']);  
  
$m = Model_Invoice($db);  
$m->insert(['qty' => 20, 'product_id' => $productIdAction]);
```

Insert operation will check if you are using same persistence. If the persistence object is different, it will execute action and will use result instead.

Being able to embed actions inside next query allows Agile Data to reduce number of queries issued.

6.6.2 SQL Actions

Currently only read-only actions are supported by Persistence\Sql:

- select - produces query that returns DataSet (array of hashes)

There are ability to execute aggregation functions:

```
echo $m->action('fx', ['max', 'salary'])->getOne();
```

and finally you can also use count:

```
echo $m->executeCountQuery(); // same as echo $m->action('count')->getOne()
```

6.6.3 SQL Actions on Linked Records

In conjunction with `Model::refLink()` you can produce expressions for creating sub-selects. The functionality is nicely wrapped inside `HasMany::addField()`:

```
$client->hasMany('Invoice', ['model' => [Model_Invoice::class]])
  ->addField('total_gross', ['aggregate' => 'sum', 'field' => 'gross']);
```

This operation is actually consisting of 3 following operations:

1. Related model is created and linked up using `refLink` that essentially places a condition between `$client` and `$invoice` assuming they will appear inside same query.
2. Action is created from `$invoice` using `'fx'` and requested method / field.
3. Expression is created with name `'total_gross'` that uses Action.

Here is a way how to intervene with the process:

```
$client->hasMany('Invoice', ['model' => [Model_Invoice::class]]);
$client->addExpression('last_sale', ['expr' => function (Model $m) {
    return $m->refLink('Invoice')
        ->setOrder('date desc')
        ->setLimit(1)
        ->action('field', ['total_gross'], 'getOne');
}], 'type' => 'float']);
```

The code above uses `refLink` and also creates expression, but it tweaks the action used.

6.6.4 Action Matrix

SQL actions apply the following:

- insert: init, mode
- update: init, mode, conditions, limit, order, hook
- delete: init, mode, conditions
- select: init, fields, conditions, limit, order, hook
- count: init, field, conditions, hook,
- field: init, field, conditions
- fx: init, field, conditions

FETCHING RESULTS

```
class Atk4\Data\Model
```

Model linked to a persistence is your “window” into DataSet and you get several ways which allow you to fetch the data.

7.1 Iterate through model data

```
Atk4\Data\Model::getIterator()
```

Create your persistence object first then iterate it:

```
$db = \Atk4\Data\Persistence::connect($dsn);
$m = new Model_Client($db);

foreach ($m as $id => $entity) {
    echo $id . ': ' . $entity->get('name') . "\n";
}
```

Note: changing query parameter during iteration will has no effect until you finish iterating.

7.1.1 Raw Data Fetching

If you do not care about the hooks and simply wish to get the data, you can fetch it:

```
foreach ($m->getPersistence()->prepareIterator($m) as $row) {
    var_dump($row); // array
}
```

The \$row will also contain value for “id” and it’s up to you to find it yourself if you need it.

```
Atk4\Data\Model::export()
```

Will fetch and output array of hashes which will represent entirety of data-set. Similarly to other methods, this will have the data mapped into your fields for you and server-side expressions executed that are embedded in the query.

By default - onlyFields will be presented as well as system fields.

7.1.2 Fetching data through action

You can invoke and iterate action (particularly SQL) to fetch the data:

```
foreach ($m->action('select')->getRowsIterator() as $row) {  
    var_dump($row); // array  
}
```

This has the identical behavior to `$m->getPersistence()->prepareIterator($m)`.

```
class Atk4\Data\Field
```

FIELD

Field represents a model property that can hold information about your entity. In Agile Data we call it a Field, to distinguish it from object properties. Fields inside a model normally have a corresponding instance of Field class.

See `Model::addField()` on how fields are added. By default, persistence sets the property `_defaultSeedAddField` which should correspond to a field object that has enough capabilities for performing field-specific mapping into persistence-logic.

```
class Atk4\Data\Field
```

Field represents a property of your business entity or column if you think of your data in a tabular way. Once you have defined Field for your Model, you can set and read value of that field:

```
$model->addField('name');  
$model->set('name', 'John');  
  
echo $model->get('name'); // John
```

8.1 Purpose of Field

Implementation of Field in Agile Data is a very powerful and distinctive feature. While `Model::$data` store your field values, the job of `Field` is to interpret that value, normalize it, type-cast it, validate it and decide on how to store or present it.

The implementation of Fields is tightly integrated with `Model` and Persistence.

8.1.1 Field Type

```
property Atk4\Data\Field::$type
```

Probably a most useful quality of Field is that it has a clear type:

```
$model->addField('age', ['type' => 'integer']);  
$model->set('age', '123');  
  
var_dump($model->get('age')); // int(123)
```

Agile Data defines some basic types to make sure that values:

- can be safely stored and manipulated.
- can be saved (Persistence)

- can be presented to user (UI)

A good example would be a date type:

```
$model->addField('birth', ['type' => 'date']);
$model->set('birth', new DateTime('2014-01-10'));

$model->save();
```

When used with SQL persistence the value will be automatically converted into a format preferred by the database 2014-10-01. Because PHP has only a single type for storing date, time and datetime, this can lead to various problems such as handling of timezones or DST. Agile Data takes care of those issues for you automatically.

Conversions between types is what we call *Typecasting* and there is a documentation section dedicated to it.

Finally, because Field is a class, it can be further extended. For some interesting examples, check out PasswordField. I'll explain how to create your own field classes and where they can be beneficial.

Valid types are: string, integer, boolean, datetime, date, time.

You can specify unsupported type too. It will be untouched by Agile Data so you would have to implement your own handling of a new type.

Persistence implements two methods:

- Persistence::typecastSaveRow()
- Persistence::typecastLoadRow()

Those are responsible for converting PHP native types to persistence specific formats as defined in fields. Those methods will also change name of the field if needed (see Field::actual)

8.1.2 Basic Properties

Fields have properties, which define its behavior. Some properties apply on how the values are handled or restrictions on interaction, other values can even help with data visualization. For example if *Field::\$enum* is used with Agile UI form, it will be displayed as radio button or a drop-down:

```
$model->addField('gender', ['enum' => ['F', 'M']]);

// Agile UI code:
$app = new \Atk4\Ui\App('my app');
$app->initLayout('Centered');
Form::addTo($app)->setModel($model);
```

You will also not be able to set value which is not one of the `enum` values even if you don't use UI.

This allows you to define your data fields once and have those rules respected everywhere in your app - in your manual code, in UI and in API.

property Atk4\Data\Field::\$default

When no value is specified for a field, default value is used when inserting. This value will also appear pre-filled inside a Form.

property Atk4\Data\Field::\$enum

Specifies array containing all the possible options for the value. You can set only to one of the values (loosely typed comparison is used).

property Atk4\Data\Field::\$values

Specifies array containing all the possible options for the value. Similar with \$enum, but difference is that this array is a hash array so array keys will be used as values and array values will be used as titles for these values.

property Atk4\Data\Field::\$nullable

Set this to false if field value must NOT be NULL. Attempting to set field value to “NULL” will result in exception. Example:

```
$model->set('age', 0);
$model->save();

$model->set('age', null); // exception
```

property Atk4\Data\Field::\$required

Set this to true for field that may not contain “empty” value. You can’t use NULL or any value that is considered empty/false by PHP. Some examples that are not allowed are:

- empty string ‘’
- 0 numerical value or 0.0
- boolean false

Example:

```
$model->set('age', 0); // exception
$model->set('age', null); // exception
```

property Atk4\Data\Field::\$readOnly

Modifying field that is read-only through set() methods (or array access) will result in exception. *Field\SqlExpressionField* is read-only by default.

property Atk4\Data\Field::\$actual

Specify name of the Table Row Field under which field will be persisted.

property Atk4\Data\Field::\$join

This property will point to Join object if field is associated with a joined table row.

property Atk4\Data\Field::\$system

System flag is intended for fields that are important to have inside hooks or some core logic of a model. System fields will always be appended to *Model::setOnlyFields*, however by default they will not appear on forms or grids (see *Field::isVisible*, *Field::isEditable*).

Adding condition on a field will also make it system.

property Atk4\Data\Field::\$neverPersist

Field will never be loaded or saved into persistence. You can use this flag for fields that physically are not located in the database, yet you want to see this field in beforeSave hooks.

property Atk4\Data\Field::\$neverSave

This field will be loaded normally, but will not be saved in a database. Unlike “readOnly” which has a similar effect, you can still change the value of this field. It will simply be ignored on save. You can create some logic in beforeSave hook to read this value.

property Atk4\Data\Field::\$ui

This field contains certain arguments that may be needed by the UI layer to know if user should be allowed to edit this field.

Atk4\Data\Field::set()

Set the value of the field. Same as `$model->set($fieldName, $value);`

Atk4\Data\Field::setNull()

Set field value to NULL. This will bypass “nullable” and “required” checks and should only be used if you are planning to set a different value to the field before executing `save()`.

If you do not set non-null value to a not-nullable field, `save()` will fail with exception.

Example:

```
$model['age'] = 0;
$model->save();

$model->getField('age')->setNull(); // no exception
$model->save(); // still getting exception here
```

Atk4\Data\Field::get()

Get the value of the field. Same as `$model->get($fieldName);`

8.1.3 UI Presentation

Agile Data does not deal directly with formatting your data for the user. There may be various items to consider, for instance the same date can be presented in a short or long format for the user.

The UI framework such as Agile Toolkit can make use of the `Field::$ui` property to allow user to define default formats or input parsing rules, but Agile Data does not regulate the `Field::$ui` property and different UI frameworks may use it differently.

Atk4\Data\Field::isEditable()

Returns true if UI should render this field as editable and include inside forms by default.

Atk4\Data\Field::isVisible()

Returns true if UI should render this field in Grid and other readOnly display views by default.

Atk4\Data\Field::isHidden()

Returns true if UI should not render this field in views.

CONDITIONS AND DATASET

```
class Atk4\Data\Model
```

When model is associated with the database, you can specify a default table either explicitly or through a \$table property inside a model:

```
$m = new Model_User($db, 'user');  
$m = $m->load(1);  
echo $m->get('gender'); // "M"
```

Following this line, you can load ANY record from the table. It's possible to narrow down set of "loadable" records by introducing a condition:

```
$m = new Model_User($db, 'user');  
$m->addCondition('gender', 'F');  
$m = $m->load(1); // exception, user with ID=1 is M
```

Conditions serve important role and must be used to intelligently restrict logically accessible data for a model before you attempt the loading.

9.1 Basic Usage

```
Atk4\Data\Model::addCondition($field, $operator = null, $value = null)
```

There are many ways to execute addCondition. The most basic one that will be supported by all the drivers consists of 2 arguments or if operator is '=':

```
$m->addCondition('gender', 'F');  
$m->addCondition('gender', '=', 'F'); // exactly same
```

Once you add a condition, you can't get rid of it, so if you want to preserve the state of your model, you need to use clone:

```
$m = new Model_User($db, 'user');  
$girls = (clone $m)->addCondition('gender', 'F');  
  
$m = $m->load(1); // success  
$girls = $girls->load(1); // exception
```

9.1.1 Operations

Most database drivers will support the following additional operations:

```
>, <, >=, <=, !=, in, not in, like, not like, regexp, not regexp
```

The operation must be specified as second argument:

```
$m = new Model_User($db, 'user');  
$girls = (clone $m)->addCondition('gender', 'F');  
$notGirls = (clone $m)->addCondition('gender', '!=', 'F');
```

When you use 'in' or 'not in' you must pass value as array or DSQL expression:

```
$m = new Model_User($db, 'user');  
$girlsOrBoys = (clone $m)->addCondition('gender', 'in', ['F', 'M']);
```

9.1.2 Multiple Conditions

You can set multiple conditions on the same field even if they are contradicting:

```
$m = new Model_User($db, 'user');  
$noone = (clone $m)  
    ->addCondition('gender', 'F')  
    ->addCondition('gender', 'M');
```

Normally, however, you would use a different fields:

```
$m = new Model_User($db, 'user');  
$girlSue = (clone $m)  
    ->addCondition('gender', 'F')  
    ->addCondition('name', 'Sue');
```

You can have as many conditions as you like.

9.1.3 Adding OR Conditions

In Agile Data all conditions are additive. This is done for security - no matter what condition you are adding, it will not allow you to circumvent previously added condition.

You can, however, add condition that contains multiple clauses joined with OR operator:

```
$m->addCondition(Model\Scope::createOr(  
    ['name', 'John'],  
    ['surname', 'Smith'],  
));
```

This will add condition that will match against records with either name=John OR surname=Smith. If you are building multiple conditions against the same field, you can use this format:

```
$m->addCondition('name', 'in', ['John', 'Joe']);
```

For all other cases you can implement them with `Model::expr`:

```
$m->addCondition($m->expr('(day([birth_date]) = day([registration_date]) or day([birth_
↳date]) = [10])', 10));
```

This rather unusual condition will show user records who have registered on same date when they were born OR if they were born on 10th. (This is really silly condition, please don't judge, if you have a better example, I'd love to hear).

9.1.4 Defining your classes

Although I have used in-line addition of the arguments, normally you would want to set those conditions inside the `init()` method of your model:

```
class Model_Girl extends Model_User
{
    protected function init(): void
    {
        parent::init();

        $this->addCondition('gender', 'F');
    }
}
```

Note that the field 'gender' should be defined inside `Model_User::init()`.

9.2 Vendor-dependent logic

There are many other ways to set conditions, but you must always check if they are supported by the driver that you are using.

9.2.1 Field Matching

Supported by: SQL (planned for Array, Mongo)

Usage:

```
$m->addCondition('name', $m->getField('surname'));
```

Will perform a match between two fields.

9.2.2 Expression Matching

Supported by: SQL (planned for Array)

Usage:

```
$m->addCondition($m->expr('[name] > [surname]');
```

Allow you to define an arbitrary expression to be used with fields. Values inside [blah] should correspond to field names.

9.2.3 SQL Expression Matching

Atk4\Data\Model::**expr**(\$template, \$arguments = [])

Basically is a wrapper to create DSQL Expression, however this will find any usage of identifiers inside the template that do not have a corresponding value inside \$arguments and replace it with the field:

```
$m->expr('[age] > 20'); // same as
$m->expr('[age] > 20', ['age' => $m->getField('age')]); // same as
```

Supported by: SQL

Usage:

```
$m->addCondition($m->expr('[age] between [min_age] and [max_age]'));
```

Allow you to define an arbitrary expression using SQL language.

9.2.4 Custom Parameters in Expressions

Supported by: SQL

Usage:

```
$m->addCondition(
    $m->expr('[age] between [min_age] and [max_age]'),
    ['min_age' => 10, 'max_age' => 30]
);
```

Allow you to pass parameters into expressions. Those can be nested and consist of objects as well as actions:

```
$m->addCondition(
    $m->expr('[age] between [min_age] and [max_age]'),
    [
        'min_age' => $m->action('min', ['age']),
        'max_age' => $m->expr('(20 + [])', [20]),
    ]
);
```

This will result in the following condition:

```
WHERE
`age` between
(select min(`age`) from `user`)
and
(20 + :a)
```

where the other 20 is passed through parameter.

9.2.5 Expression as first argument

Supported by: SQL, (Planned: Array, Mongo)

The \$field of addCondition() can be passed as either an expression or any object implementing Atk4\Data\Persistence\Sql\Expressionable interface. Same logic applies to the \$value:

```
$m->addCondition($m->getField('name'), '!=', $this->getField('surname'));
```

9.3 Advanced Usage

9.3.1 Model Scope

Using the Model::addCondition method is the basic way to limit the model scope of records. Under the hood Agile Data utilizes a special set of classes (Condition and Scope) to apply the conditions as filters on records retrieved. These classes can be used directly and independently from Model class.

Atk4\Data\Model::scope()

This method provides access to the model scope enabling conditions to be added:

```
$contact->scope()->addCondition($condition); // adding condition to a model
```

class Atk4\Data\Model\Scope

Scope object has a single defined junction (AND or OR) and can contain multiple nested Condition and/or Scope objects referred to as nested conditions. This makes creating Model scopes with deep nested conditions possible, e.g ((Name like 'ABC%' and Country = 'US') or (Name like 'CDE%' and (Country = 'DE' or Surname = 'XYZ')))

Scope can be created using new Scope() statement from an array or joining Condition objects or condition arguments arrays:

```
// $condition1 will be used as nested condition
$condition1 = new Condition('name', 'like', 'ABC%');

// $condition2 will be converted to Condition object and used as nested condition
$condition2 = ['country', 'US'];

// $scope1 is created using AND as junction and $condition1 and $condition2 as nested_
↳conditions
$scope1 = Scope::createAnd($condition1, $condition2);

$condition3 = new Condition('country', 'DE');
$condition4 = ['surname', 'XYZ'];

// $scope2 is created using OR as junction and $condition3 and $condition4 as nested_
↳conditions
$scope2 = Scope::createOr($condition3, $condition4);

$condition5 = new Condition('name', 'like', 'CDE%');

// $scope3 is created using AND as junction and $condition5 and $scope2 as nested_
↳conditions
```

(continues on next page)

(continued from previous page)

```
$scope3 = Scope::createAnd($condition5, $scope2);  
  
// $scope is created using OR as junction and $scope1 and $scope3 as nested conditions  
$scope = Scope::createOr($scope1, $scope3);
```

Scope is an independent object not related to any model. Applying scope to model is using the `Model::scope()->add($condition)` method:

```
$contact->scope()->add($condition); // adding condition to a model  
$contact->scope()->add($conditionXYZ); // adding more conditions
```

`Atk4\Data\Model\Scope::__construct($conditions = [], $junction = Scope::AND)`

Creates a Scope object from an array:

```
// below will create 2 conditions and nest them in a compound conditions with AND  
↪ junction  
$scope1 = new Scope([  
    ['name', 'like', 'ABC%'],  
    ['country', 'US'],  
]);
```

`Atk4\Data\Model\Scope::negate()`

Negate method has behind the full map of conditions so any condition object can be negated, e.g negating '=>' results in '<', etc. For compound conditions this method is using De Morgan's laws, e.g:

```
// using $scope1 defined above  
// results in "(Name not like 'ABC%') or (Country does not equal 'US')"  
$scope1->negate();
```

`Atk4\Data\Model\Scope::createAnd(...$conditions)`

Merge \$conditions using AND as junction. Returns the resulting Scope object.

`Atk4\Data\Model\Scope::createOr(...$conditions)`

Merge \$conditions using OR as junction. Returns the resulting Scope object.

`Atk4\Data\Model\Scope::simplify()`

Peels off single nested conditions. Useful for converting (((field = value))) to field = value.

`Atk4\Data\Model\Scope::clear()`

Clears the condition from nested conditions.

`Atk4\Data\Model\Scope::isOr()`

Checks if scope components are joined by OR

`Atk4\Data\Model\Scope::isAnd()`

Checks if scope components are joined by AND

class `Atk4\Data\Model\Scope\Condition`

Condition represents a simple condition in a form [field, operation, value], similar to the functionality of the `Model::addCondition` method

```
Atk4\Data\Model\Scope\Condition::__construct($key, $operator = null, $value = null)
```

Creates condition object based on provided arguments. It acts similar to Model::addCondition

\$key can be Model field name, Field object or Expression object \$operator can be one of the supported operators >, <, >=, <=, !=, in, not in, like, not like, regexp, not regexp \$value can be Field object, Expression object or any scalar value

If \$value is omitted as argument then \$operator is considered as \$value and '=' is used as operator

```
Atk4\Data\Model\Scope\Condition::negate()
```

Negates the condition, e.g:

```
// results in "name != 'John'"
$condition = (new Condition('name', 'John'))->negate();
```

```
Atk4\Data\Model\Scope\Condition::toWords(Model $model = null)
```

Converts the condition object to human readable words. Condition must be assigned to a model or model argument provided:

```
// results in 'Contact where Name is John'
(new Condition('name', 'John'))->toWords($contactModel);
```

9.3.2 Conditions on Referenced Models

Agile Data allows for adding conditions on related models for retrieval of type 'model has references where'.

Setting conditions on references can be done utilizing the Model::refLink method but there is a shorthand format directly integrated with addCondition method using "/" to chain the reference names:

```
$contact->addCondition('company/country', 'US');
```

This will limit the \$contact model to those whose company is in US. 'company' is the name of the reference in \$contact model and 'country' is a field in the referenced model.

If a condition must be set directly on the existence or number of referenced records the special symbol "#" can be utilized to indicate the condition is on the number of records:

```
$contact->addCondition('company/tickets/#', '>', 3);
```

This will limit the \$contact model to those whose company have more than 3 tickets. 'company' and 'tickets' are the name of the chained references ('company' is a reference in the \$contact model and 'tickets' is a reference in Company model)

SQL EXTENSIONS

Databases that support SQL language can use *Persistence\Sql*. This driver will format queries to the database using SQL language.

In addition to normal operations you can extend and customize various queries.

10.1 Default Model Classes

When using *Persistence\Sql* model building will use different classes for fields, expressions, joins etc:

- addField - different class is no longer used/needed
- hasOne - *Reference\HasOneSql* (allow importing fields)
- addExpression - *Field\SqlExpressionField* (define expression through DSQL)
- join - *Persistence\Sql\Join* (join tables query-time)

10.1.1 SQL Field

class Atk4\Data\Field

property Atk4\Data\Field::\$actual

Persistence\Sql supports field name mapping. Your field could have different column name in your schema:

```
$this->addField('name', ['actual' => 'first_name']);
```

This will apply to load / save operations as well as query mapping.

Atk4\Data\Field::getDsqlExpression()

SQL Fields can be used inside other SQL expressions:

```
$q = $connection->expr('[age] + [birth_year]', [
    'age' => $m->getField('age'),
    'birth_year' => $m->getField('birth_year'),
]);
```

10.1.2 SQL Reference

class `Atk4\Data\Reference\HasOneSql`

Extends `Reference\HasOne`

`Atk4\Data\Reference\HasOneSql::addField()`

Allows importing field from a referenced model:

```
$model->hasOne('country_id', ['model' => [Country::class]])
->addField('country_name', 'name');
```

Second argument could be array containing additional settings for the field:

```
$model->hasOne('account_id', ['model' => [Account::class]])
->addField('account_balance', ['balance', 'type' => 'atk4_money']);
```

Returns new field object.

`Atk4\Data\Reference\HasOneSql::addFields()`

Allows importing multiple fields:

```
$model->hasOne('country_id', ['model' => [Country::class]])
->addFields(['country_name', 'country_code']);
```

You can specify defaults to be applied on all fields:

```
$model->hasOne('account_id', ['model' => [Account::class]])
->addFields([
    'opening_balance',
    'balance',
], ['type' => 'atk4_money']);
```

You can also specify aliases:

```
$model->hasOne('account_id', ['model' => [Account::class]])
->addFields([
    'opening_balance',
    'account_balance' => 'balance',
], ['type' => 'atk4_money']);
```

If you need to pass more details to individual field, you can also use sub-array:

```
$model->hasOne('account_id', ['model' => [Account::class]])
->addFields([
    [
        'opening_balance', 'caption' => 'The Opening Balance',
        'account_balance' => 'balance',
    ],
], ['type' => 'atk4_money']);
```

Returns `$this`.

`Atk4\Data\Reference\HasOneSql::ref()`

While similar to `Reference\HasOne::ref` this implementation implements deep traversal:

```
$countryModel = $customerModel->addCondition('is_vip', true)
->ref('country_id'); // $model was not loaded!
```

`Atk4\Data\Reference\HasOneSql::refLink()`

Creates a model for related entity with applied condition referencing field of a current model through SQL expression rather than value. This is usable if you are creating sub-queries.

`Atk4\Data\Reference\HasOneSql::addTitle()`

Similar to `addField`, but will import “title” field and will come up with good name for it:

```
$model->hasOne('country_id', ['model' => [Country::class]])
    ->addTitle();

// creates 'country' field as sub-query for country.name
```

You may pass defaults:

```
$model->hasOne('country_id', ['model' => [Country::class]])
    ->addTitle(['caption' => 'Country Name']);
```

Returns new field object.

10.1.3 Expressions

`class Atk4\Data\Field\SqlExpressionField`

Extends `Field`

Expression will map into the SQL code, but will perform as read-only field otherwise.

property `Atk4\Data\Field\SqlExpressionField::$expr`

Stores expression that you define through DSQL expression:

```
$model->addExpression('age', ['expr' => 'year(now()) - [birth_year]']);
// tag [birth_year] will be automatically replaced by respective model field
```

`Atk4\Data\Field\SqlExpressionField::getDsqlExpression()`

SQL Expressions can be used inside other SQL expressions:

```
$model->addExpression('can_buy_alcohol', ['expr' => 'if([age] > 25, 1, 0)', 'type' =>
    'boolean']);
```

Adding expressions to model will make it automatically reload itself after save as default behavior, see `Model::$reloadAfterSave`.

10.2 Transactions

`class Atk4\Data\Persistence\Sql`

`Atk4\Data\Persistence\Sql::atomic()`

This method allows you to execute code within a ‘START TRANSACTION / COMMIT’ block:

```
class Invoice
{
    public function applyPayment(Payment $p)
```

(continues on next page)

```
{
    $this->getModel()->getPersistence()->atomic(function () use ($p) {
        $this->set('paid', true);
        $this->save();

        $p->set('applied', true);
        $p->save();
    });
}
```

Callback format of this method allows a more intuitive syntax and nested execution of various blocks. If any exception is raised within the block, then transaction will be automatically rolled back. The return of `atomic()` is same as return of user-defined callback.

10.3 Custom Expressions

Atk4\Data\Persistence\Sql::expr()

This method is also injected into the model, that is associated with `Persistence\Sql` so the most convenient way to use this method is by calling `$model->expr('foo')`.

This method is quite similar to `\Atk4\Data\Persistence\Sql\Query::expr()` method.

There is, however, one difference. Expression class requires all named arguments to be specified. Use of `Model::expr()` allows you to specify field names and those field expressions will be automatically substituted. Here is long / short format:

```
$q = $connection->expr('[age] + [birth_year]', [
    'age' => $m->getField('age'),
    'birth_year' => $m->getField('birth_year'),
]);

// identical to

$q = $m->expr('[age] + [birth_year]');
```

This method is automatically used by `Field\SqlExpressionField`.

10.4 Actions

The most basic action you can use with SQL persistence is 'select':

```
$action = $model->action('select');
```

10.4.1 Action: select

This action returns a basic select query. You may pass one argument - array containing list of fields:

```
$action = $model->action('select', ['name', 'surname']);
```

Passing false will not include any fields into select (so that you can include them yourself):

```
$action = $model->action('select', [false]);
$action->field('count(*)', 'c');
```

10.4.2 Action: count

Returns query for count(*):

```
$action = $model->action('count');
$snt = $action->getOne();
// for materialized count use:
$snt = $model->executeCountQuery();
```

You can also specify alias:

```
$action = $model->action('count', ['alias' => 'cc']);
$data = $action->getRow();
$snt = $data->get('cc');
```

10.4.3 Action: field

Get query for a specific field:

```
$action = $model->action('field', ['age']);
$age = $action->limit(1)->getOne();
```

You can also specify alias:

```
$action = $model->action('field', ['age', 'alias' => 'the_age']);
$age = $action->limit(1)->getRow()['the_age'];
```

10.4.4 Action: fx

Executes single-argument SQL function on field:

```
$action = $model->action('fx', ['avg', 'age']);
$ageAvg = $action->getOne();
```

This method also supports alias. Use of alias is handy if you are using those actions as part of other query (e.g. UNION)

10.5 Stored Procedures

SQL servers allow to create and use stored procedures and there are several ways to invoke them:

1. CALL procedure. No data / output.
2. Specify OUT parameters.
3. Stored FUNCTION, e.g. `select myfunc(123)`
4. Stored procedures that return data.

Agile Data has various ways to deal with above scenarios:

1. Custom expression through DSQL
2. Model Method
3. Model Field
4. Model Source

Here I'll try to look into each of those approaches but closely pay attention to the following:

- Abstraction and concern separation.
- Security and protecting against injection.
- Performance and scalability.
- When to refactor away stored procedures.

10.5.1 Compatibility Warning

Agile Data is designed to be cross-database agnostic. That means you should be able to swap your SQL to NoSQL or RestAPI at any moment. My relying on stored procedures you will loose portability of your application.

We do have our legacy applications to maintain, so Stored Procedures and SQL extensions are here to stay. By making your Model rely on those extensions you will loose ability to use the same model with non-sql persistencies.

Sometimes you can fence the code like this:

```
if ($this->getPersistence() instanceof \Atk4\Data\Persistence\Sql) {  
    .. sql code ..  
}
```

Or define your pure model, then extend it to add SQL capabilities. Note that using single model with cross-persistencies should still be possible, so you should be able to retrieve model data from stored procedure then cache it.

10.5.2 as a Model method

In short this should allow you to build and execute any SQL statement:

```
$this->expr('call get_nominal_sheet([], [], \'2014-10-01\', \'2015-09-30\', 0)', [  
    $this->getApp()->system->getId(),  
    $this->getApp()->system['contractor_id'],  
)->executeQuery();
```

Depending on the statement you can also use your statement to retrieve data:

```
$data = $this->expr('call get_client_report_data([client_id])', [
    'client_id' => $clientId,
])->getRows();
```

This can be handy if you wish to create a method for your Model to abstract away the data:

```
class Client extends \Atk4\Data\Model
{
    protected function init(): void
    {
        ...
    }

    public function getReportData($arg)
    {
        $this->assertIsLoaded();

        return $this->expr('call get_client_report_data([client_id], [arg])', [
            'arg' => $arg,
            'client_id' => $clientId,
        ]->getRows();
    }
}
```

Here is another example using PHP generator:

```
class Client extends \Atk4\Data\Model
{
    protected function init(): void
    {
        ...
    }

    public function fetchReportData($arg)
    {
        $this->assertIsLoaded();

        foreach ($this->expr('call get_client_report_data([client_id], [arg])', [
            'arg' => $arg,
            'client_id' => $clientId,
        ]) as $row) {
            yield $row;
        }
    }
}
```

10.5.3 as a Model Field

Important: Not all SQL vendors may support this approach.

`Model::addExpression` is a SQL extension that allow you to define any expression for your field query. You can use SQL stored function for data fetching like this:

```
class Category extends \Atk4\Data\Model
{
    public $table = 'category';

    protected function init(): void
    {
        parent::init();

        $this->hasOne('parent_id', ['model' => [self::class]]);
        $this->addField('name');

        $this->addExpression('path', ['expr' => 'get_path([id])']);
    }
}
```

This should translate into SQL query:

```
select parent_id, name, get_path(id) from category;
```

where once again, stored function is hidden.

10.5.4 as an Action

Important: Not all SQL vendors may support this approach.

Method `Persistence\Sql::action` and `Model::action` generates queries for most of model operations. By re-defining this method, you can significantly affect the query building of an SQL model:

```
class CompanyProfit extends \Atk4\Data\Model
{
    public $companyId; // inject company ID, which will act as a condition/argument
    public bool $readOnly = true; // instructs rest of the app, that this model is read-
    ↪only

    protected function init(): void
    {
        parent::init();

        $this->addField('date_period');
        $this->addField('profit');
    }

    public function action(string $mode, array $args = [])
```

(continues on next page)

(continued from previous page)

```

{
    if ($mode === 'select') {
        // must return DSQL object here
        return $this->expr('call get_company_profit([company_id])', [
            'company_id' => $this->companyId,
        ]);
    }

    if ($mode === 'count') {
        // optionally - expression for counting data rows, for pagination support
        return $this->expr('select count(*) from (call get_company_profit([company_
↵id]))', [
            'company_id' => $this->companyId,
        ]);
    }

    throw (new \Atk4\Data\Exception('You may only perform "select" or "count" action_
↵on this model'))
        ->addMoreInfo('action', $mode);
}
}

```

10.5.5 as a Temporary Table

A most convenient (although inefficient) way for stored procedures is to place output data inside a temporary table. You can perform an actual call to stored procedure inside Model::init() then set \$table property to a temporary table:

```

class NominalReport extends \Atk4\Data\Model
{
    public $table = 'temp_nominal_sheet';
    public bool $readOnly = true; // instructs rest of the app, that this model is read-
↵only

    protected function init(): void
    {
        parent::init();

        $res = $this->expr('call get_nominal_sheet([], [], \'2014-10-01\', \'2015-09-30\
↵', 0)'), [
            $this->getApp()->system->getId(),
            $this->getApp()->system['contractor_id'],
        ]->executeQuery();

        $this->addField('date', ['type' => 'date']);
        $this->addField('items', ['type' => 'integer']);

        ...
    }
}

```

10.5.6 as a Model Source

Important: Not all SQL vendors may support this approach.

Technically you can also specify expression as a \$table property of your model:

```
class ClientReport extends \Atk4\Data\Model
{
    public $table; // will be set in init()
    public bool $readOnly = true; // instructs rest of the app, that this model is read-
    ↪only

    protected function init(): void
    {
        parent::init();

        $this->table = $this->expr('call get_report_data()');

        $this->addField('date', ['type' => 'date']);
        $this->addField('items', ['type' => 'integer']);

        ...
    }
}
```

Technically this will give you select date, items from (call get_report_data()).

STATIC PERSISTENCE

```
class Atk4\Data\Persistence\Static_
```

Static Persistence extends `Persistence\Array_` to implement a user-friendly way of specifying data through an array.

11.1 Usage

This is most useful when working with “sample” code, where you want to see your results quick:

```
$table->setModel(new Model(new Persistence\Static_([
    ['VAT_rate' => '12.0%', 'VAT' => '36.00', 'Net' => '300.00'],
    ['VAT_rate' => '10.0%', 'VAT' => '52.00', 'Net' => '520.00'],
])););
```

Lets unwrap the example:

```
Atk4\Data\Persistence\Static_::__construct()
```

Constructor accepts array as an argument, but the array could be in various forms:

- can be array of strings ['one', 'two']
- can be array of hashes. First hash will be examined to pick up fields
- can be array of arrays. Will name columns as 'field1', 'field2', 'field3'.

If you are using any fields without keys (numeric keys) it's important that all your records have same number of elements.

Static Persistence will also make attempt to deduce a “title” field and will set it automatically for the model. If you have a field with key “name” then it will be used. Alternative it will check key “title”.

If neither are present you can still manually specify title field for your model.

Finally, static persistence (unlike `Persistence\Array_`) will automatically populate fields for the model and will even attempt to deduce field types.

Currently it recognizes integer, date, boolean, float, array and object types. Other fields will appear as-is.

11.1.1 Saving Records

Models that you specify against static persistence will not be marked as “Read Only” (`Model::$readOnly`), and you will be allowed to save data back. The data will only be stored inside persistence object and will be discarded at the end of your PHP script.

REFERENCES

```
class Atk4\Data\Model
```

```
Atk4\Data\Model::ref($link, $details = [])
```

Models can relate one to another. The logic of traversing references, however, is slightly different to the traditional ORM implementation, because in Agile Data traversing also imposes *Conditions and DataSet*

There are two basic types of references: `hasOne()` and `hasMany()`, but it's also possible to add other reference types. The basic ones are really easy to use:

```
$m = new Model_User($db, 'user');
$m->hasMany('Orders', ['model' => [Model_Order::class]]);
$m = $m->load(13);

$sordersForUser13 = $m->ref('Orders');
```

As mentioned - `$sordersForUser13` will have its `DataSet` automatically adjusted so that you could only access orders for the user with ID=13. The following is also possible:

```
$m = new Model_User($db, 'user');
$m->hasMany('Orders', ['model' => [Model_Order::class]]);
$m->addCondition('is_vip', true);

$sordersForVips = $m->ref('Orders');
$sordersForVips = $sordersForVips->loadAny();
```

Condition on the base model will be carried over to the orders and you will only be able to access orders that belong to VIP users. The query for loading order will look like this:

```
select * from order where user_id in (
    select id from user where is_vip = 1
) limit 1
```

Argument `$defaults` will be passed to the new model that will be used to create referenced model. This will not work if you have specified reference as existing model that has a persistence set. (See `Reference::getModel()`)

12.1 Persistence

Agile Data supports traversal between persistencies. The code above does not explicitly assign database to `Model_Order`. But what if destination model does not reside inside the same database?

You can specify it like this:

```
$m = new Model_User($dbArrayCache, 'user');
$m->hasMany('Orders', ['model' => [Model_Order::class, $dbSql]]);
$m->addCondition('is_vip', true);

$ordersForVips = $m->ref('Orders');
```

Now that a different databases are used, the queries can no longer be joined so Agile Data will carry over list of IDs instead:

```
$ids = select id from user where is_vip = 1
select * from order where user_id in ($ids)
```

Since we are using `$dbArrayCache`, then field values will actually be retrieved from memory.

Note: This is not implemented as of 1.1.0, see <https://github.com/atk4/data/issues/158>

12.2 Safety and Performance

When using `ref()` on `hasMany` reference, it will always return a fresh clone of the model. You can perform actions on the clone and next time you execute `ref()` you will get a fresh copy.

If you are worried about performance you can keep 2 models in memory:

```
$order = new Order($db);
$client = $order->ref('client_id');

foreach ($order as $o) {
    $c = $client->load($o->get('client_id'));
}
```

Warning: This code is seriously flawed and is called “N+1 Problem”. Agile Data discourages you from using this and instead offers you many other tools: field importing, model joins, field actions and `refLink()`.

12.3 hasMany Reference

```
Atk4\Data\Model::hasMany($link[, 'model' => $model ])
```

There are several ways how to link models with hasMany:

```
$m->hasMany('Orders', ['model' => [Model_Order::class]]); // using seed

$m->hasMany('Order', ['model' => static function (Persistence $persistence, array
↳ $defaults) { // using callback
    return new Model_Order();
}]);
```

12.3.1 Dealing with many-to-many references

It is possible to perform reference through an 3rd party table:

```
$i = new Model_Invoice();

$i->hasMany('Payments', ['model' => static function () {
    return (new Model_Payment())
        ->join('invoice_payment.payment_id')
        ->addFields(['amount_allocated', 'invoice_id']);
}]);
```

Now you can fetch all the payments associated with the invoice through:

```
$paymentsForInvoice1 = $i->load(1)->ref('Payments');
```

12.3.2 Dealing with NON-ID fields

Sometimes you have to use non-ID references. For example, we might have two models describing list of currencies and for each currency we might have historic rates available. Both models will relate through `currency.code = exchange.currency_code`:

```
$c = new Model_Currency();
$c->hasMany('Exchanges', ['model' => [Model_ExchangeRate::class], 'theirField' =>
↳ 'currency_code', 'ourField' => 'code']);
$c->addCondition('is_convertible', true);
$e = $c->ref('Exchanges');
```

This will produce the following query:

```
select * from exchange
where currency_code in
    (select code from currency where is_convertible = 1)
```

12.3.3 Concatenating Fields

You may want to display want to list your related entities by concatenating. For example:

```
$user->hasMany('Tags', ['model' => [Tag::class]])
->addField('tags', ['concat' => ',', 'field' => 'name']);
```

This will create a new field for your user, `tags` which will contain all comma-separated tag names.

12.3.4 Add Aggregate Fields

Reference `hasMany` makes it a little simpler for you to define an aggregate fields:

```
$u = new Model_User($dbArrayCache, 'user');

$u->hasMany('Orders', ['model' => [Model_Order::class]])
->addField('amount', ['aggregate' => 'sum']);
```

It's important to define aggregation functions here. This will add another field inside `$m` that will correspond to the sum of all the orders. Here is another example:

```
$u->hasMany('PaidOrders', ['model' => static function (Persistence $persistence) {
    return (new Model_Order($persistence))
        ->addCondition('is_paid', true);
}])
->addField('paid_amount', ['aggregate' => 'sum', 'field' => 'amount']);
```

You can also define multiple fields, although you must remember that this will keep making your query bigger and bigger:

```
$invoice->hasMany('Invoice_Line', ['model' => [Model_Invoice_Line::class]])
->addFields([
    ['total_vat', 'aggregate' => 'sum'],
    ['total_net', 'aggregate' => 'sum'],
    ['total_gross', 'aggregate' => 'sum'],
]);
```

Imported fields will preserve format of the field they reference. In the example, if `'Invoice_line'` field `total_vat` has type `money` then it will also be used for a sum.

You can also specify a type yourself:

```
->addField('paid_amount', ['aggregate' => 'sum', 'field' => 'amount', 'type' => 'atk4_
↪money']);
```

Aggregate fields are always declared read-only, and if you try to change them (`$m->set('paid_amount', 123);`), you will receive exception.

12.3.5 Available Aggregation Functions

The mathematical aggregate `sum` will automatically default to 0 if no respective rows were provided. The default SQL behavior is to return NULL, but this does go well with the cascading formulas:

```
coalesce(sum([field]), 0);
```

For other functions, such as `min`, `max`, `avg` and non mathematical aggregates such as `group_concat` no zero-coalesce will be used. Expect that result could be zero or null.

When you specify 'aggregate' => 'count' field defaults to *.

12.3.6 Aggregate Expressions

Sometimes you want to use a more complex formula, and you may do so by specifying expression into 'aggregate':

```
->addField('len', ['expr' => 'sum(length([name]))']),
```

You can reference fields by using square brackets here. Also you may pass args containing your optional arguments:

```
->addField('len', [
    'expr' => 'sum(if([date] = [exp_date], 1, 0))',
    'args' => ['exp_date' => '2003-03-04'],
]),
```

Alternatively you may also specify either 'aggregate':

```
$book->hasMany('Pages', ['model' => [Page::class]])
->addField('page_list', [
    'aggregate' => $book->getReference('Pages')->createTheirModel()->expr('group_
concat([number], [])', ['-']),
]);
```

or 'field':

```
->addField('paid_amount', ['aggregate' => 'count', 'field' => new \Atk4\Data\Persistence\
Sql\Expression('*')]);
```

Note: as of 1.3.4 count's field defaults to * - no need to specify explicitly.

12.4 hasMany / refLink

Atk4\Data\Model::refLink(\$link)

Normally ref() will return a usable model back to you, however if you use refLink then the conditioning will be done differently. refLink is useful when defining sub-queries:

```
$m = new Model_User($dbArrayCache, 'user');
$m->hasMany('Orders', ['model' => [Model_Order::class]]);
$m->addCondition('is_vip', true);
```

(continues on next page)

(continued from previous page)

```
$sum = $m->refLink('Orders')->action('fx0', ['sum', 'amount']);
$m->addExpression('sum_amount')->set($sum);
```

The refLink would define a condition on a query like this:

```
select * from `order` where user_id = `user`.id
```

And it will not be viable on its own, however if you use it inside a sub-query, then it now makes sense for generating expression:

```
select
  (select sum(amount) from `order` where user_id = `user`.id) sum_amount
from user
where is_vip = 1
```

12.5 hasOne reference

```
Atk4\Data\Model::hasOne($link[, 'model' => $model])
    $model can be an array containing options: [$model, ...]
```

This reference allows you to attach a related model to a foreign key:

```
$o = new Model_Order($db, 'order');
$u = new Model_User($db, 'user');

$o->hasOne('user_id', ['model' => $u]);
```

This reference is similar to hasMany, but it does behave slightly different. Also this reference will define a system new field user_id if you haven't done so already.

12.5.1 Traversing loaded model

If your \$o model is loaded, then traversing into user will also load the user, because we specifically know the ID of that user. No conditions will be set:

```
echo $o->load(3)->ref('user_id')['name']; // will show name of the user, of order #3
```

12.5.2 Traversing DataSet

If your model is not loaded then using ref() will traverse by conditioning DataSet of the user model:

```
$o->unload(); // just to be sure!
$o->addCondition('status', 'failed');
$u = $o->ref('user_id');

$u = $u->loadAny(); // will load some user who has at least one failed order
```

The important point here is that no additional queries are generated in the process and the loadAny() will look like this:

```
select * from user where id in
(select user_id from order where status = 'failed')
```

By passing options to `hasOne()` you can also differentiate field name:

```
$o->addField('user_id', ['type' => 'bigint']);
$o->hasOne('User', ['model' => $u, 'ourField' => 'user_id']);

$o->load(1)->ref('User')['name'];
```

You can also use `theirField` if you need non-id matching (see example above for `hasMany()`).

12.5.3 Importing Fields

You can import some fields from related model. For example if you have list of invoices, and each invoice contains “currency_id”, but in order to get the currency name you need another table, you can use this syntax to easily import the field:

```
$i = new Model_Invoice($db)
$c = new Model_Currency($db);

$i->hasOne('currency_id', ['model' => $c])
  ->addField('currency_name', 'name');
```

This code also resolves problem with a duplicate ‘name’ field. Since you might have a ‘name’ field inside ‘Invoice’ already, you can name the field ‘currency_name’ which will reference ‘name’ field inside Currency. You can also import multiple fields but keep in mind that this may make your query much longer. The argument is associative array and if key is specified, then the field will be renamed, just as we did above:

```
$u = new Model_User($db)
$a = new Model_Address($db);

$u->hasOne('address_id', ['model' => $a])
  ->addFields([
    'address_1',
    'address_2',
    'address_3',
    'address_notes' => ['notes', 'type' => 'text'],
  ]);
```

Above, all `address_` fields are copied with the same name, however field ‘notes’ from Address model will be called ‘address_notes’ inside user model.

Important: When importing fields, they will preserve type, e.g. if you are importing ‘date’ then the type of your imported field will also be date. Imported fields are also marked as “read-only” and attempt to change them will result in exception.

12.5.4 Importing hasOne Title

When you are using `hasOne()` in most cases the referenced object will be addressed through “ID” but will have a human-readable field as well. In the example above `Model_Currency` has a title field called `name`. Agile Data provides you an easier way how to define currency title:

```
$i = new Invoice($db)

$i->hasOne('currency_id', ['model' => [Currency::class]])
  ->addTitle();
```

This would create ‘currency’ field containing name of the currency:

```
$i = $i->load(20);

echo 'Currency for invoice 20 is ' . $i->get('currency'); // EUR
```

Unlike `addField()` which creates fields read-only, title field can in fact be modified:

```
$i->set('currency', 'GBP');
$i->save();

// will update $i->get('currency_id') to the corresponding ID for currency with name GBP
```

This behavior is awesome when you are importing large amounts of data, because the lookup for the `currency_id` is entirely done in a database.

By default name of the field will be calculated by removing “_id” from the end of `hasOne` field, but to override this, you can specify name of the title field explicitly:

```
$i->hasOne('currency_id', ['model' => [Currency::class]])
  ->addTitle(['field' => 'currency_name']);
```

12.6 User-defined Reference

`Atk4\Data\Model::addReference($link, $callback)`

Sometimes you would want to have a different type of relation between models, so with `addReference` you can define whatever reference you want:

```
$m->addReference('Archive', ['model' => static function (Persistence $persistence) use (
  ↪ $m) {
    return new $m(null, ['table' => $m->table . '_archive']);
  }]);
```

The above example will work for a table structure where a main table `user` is shadowed by a archive table `user_archive`. Structure of both tables are same, and if you wish to look into an archive of a `User` you would do:

```
$user->ref('Archive');
```

Note that you can create one-to-many or many-to-one relations, by using your custom logic. No condition will be applied by default so it’s all up to you:

```

$model->addReference('Archive', ['model' => static function (Persistence $persistence) use (
  ↪ $m) {
    $archive = new $m(null, ['table' => $m->table . '_archive']);

    $m->addField('original_id', ['type' => 'bigint']);

    if ($m->isLoaded()) {
      $archive->addCondition('original_id', $m->getId());
      // only show record of currently loaded record
    }
  }
});

```

12.7 Reference Discovery

You can call `Model::getReferences()` to fetch all the references of a model:

```

$references = $model->getReferences();
$reference = $references['owner_id'];

```

or if you know the reference you'd like to fetch, you can use `Model::getReference()`:

```

$reference = $model->getReference('owner_id');

```

While `Model::ref()` returns a related model, `Model::getReference()` gives you the reference object itself so that you could perform some changes on it, such as import more fields with `Model::addField()`.

`Model::hasReference()` can be used to check if particular reference exists in model:

```

if ($model->hasReference('owner_id')) {
    $reference = $model->getReference('owner_id');
}

```

12.8 Deep traversal

When operating with data-sets you can define references that use deep traversal:

```

echo $o->load(1)->ref('user_id')->ref('address_id')['address_1'];

```

The above example will actually perform 3 load operations, because as I have explained above, `Model::ref()` loads related model when called on a loaded model. To perform a single query instead, you can use:

```

echo $o->addCondition('id', 1)->ref('user_id')->ref('address_id')->loadAny()['address_1
  ↪'];

```

Here `addCondition('id', 1)` will only set a condition without actually loading the record and traversal will encapsulate sub-queries resulting in a query like this:

```

select * from address where id in
  (select address_id from user where id in
    (select user_id from order where id = 1 ))

```

12.9 Reference Aliases

When related entity relies on the same table it is possible to run into problem when SQL is confused about which table to use.

```
select name, (select name from item where item.parent_id = item.id) parent_name from item
```

To avoid this problem Agile Data will automatically alias tables in sub-queries. Here is how it works:

```
$item->hasMany('parent_item_id', ['model' => [Model_Item::class]])
->addField('parent', 'name');
```

When generating expression for 'parent', the sub-query will use alias pi consisting of first letters in 'parent_item_id'. (except _id). You can actually specify a custom table alias if you want:

```
$item->hasMany('parent_item_id', ['model' => [Model_Item::class], 'tableAlias' => 'mypi
->'])
->addField('parent', 'name');
```

Additionally you can pass tableAlias as second argument into *Model::ref()* or *Model::refLink()*. This can help you in creating a recursive models that relate to itself. Here is example:

```
class Model_Item3 extends \Atk4\Data\Model
{
    public $table = 'item';

    protected function init(): void
    {
        parent::init();

        $m = new Model_Item3();

        $this->addField('name');
        $this->addField('age');
        $i2 = $this->join('item2.item_id');
        $i2->hasOne('parent_item_id', ['model' => $m, 'tableAlias' => 'parent'])
            ->addTitle();

        $this->hasMany('Child', ['model' => $m, 'theirField' => 'parent_item_id',
->'tableAlias' => 'child'])
            ->addField('child_age', ['aggregate' => 'sum', 'field' => 'age']);
    }
}
```

Loading model like that can produce a pretty sophisticated query:

```
select
    `pp`.`id`, `pp`.`name`, `pp`.`age`, `pp_i`.`parent_item_id`,
    (select `parent`.`name`
     from `item` `parent`
     left join `item2` as `parent_i` on `parent_i`.`item_id` = `parent`.`id`
     where `parent`.`id` = `pp_i`.`parent_item_id`
    ) `parent_item`,
    (select sum(`child`.`age`) from `item` `child`
```

(continues on next page)

(continued from previous page)

```

left join `item2` as `child_i` on `child_i`.`item_id` = `child`.`id`
where `child_i`.`parent_item_id` = `pp`.`id`
) `child_age`, `pp`.`id` `i`
from `item` `pp` left join `item2` as `pp_i` on `pp_i`.`item_id` = `pp`.`id`

```

12.10 Various ways to specify options

When calling `hasOne()->addFields()` there are various ways to pass options:

- `addFields(['name', 'dob'])` - no options are passed, use defaults. Note that reference will not fetch the type of foreign field due to performance consideration.
- `addFields(['first_name' => 'name'])` - this indicates aliasing. Field name will be added as `first_name`.
- `addFields(['dob', 'type' => 'date'])` - wrap inside array to pass options to field
- `addFields(['the_date' => ['dob', 'type' => 'date']])` - combination of aliasing and options
- `addFields(['dob', 'dod'], ['type' => 'date'])` - passing defaults for multiple fields

12.11 References with New Records

Agile Data takes extra care to help you link your new records with new related entities. Consider the following two models:

```

class Model_User extends \Atk4\Data\Model
{
    public $table = 'user';

    protected function init(): void
    {
        parent::init();

        $this->addField('name');

        $this->hasOne('contact_id', ['model' => [Model_Contact::class]]);
    }
}

class Model_Contact extends \Atk4\Data\Model
{
    public $table = 'contact';

    protected function init(): void
    {
        parent::init();

        $this->addField('address');
    }
}

```

This is a classic one to one reference, but let's look what happens when you are working with a new model:

```
$m = new Model_User($db);  
  
$m->set('name', 'John');  
$m->save();
```

In this scenario, a new record will be added into 'user' with 'contact_id' equal to null. The next example will traverse into the contact to set it up:

```
$m = new Model_User($db);  
  
$m->set('name', 'John');  
$m->ref('address_id')->save(['address' => 'street']);  
$m->save();
```

When entity which you have referenced through ref() is saved, it will automatically populate \$m->get('contact_id') field and the final \$m->save() will also store the reference.

ID setting is implemented through a basic hook. Related model will have afterSave hook, which will update address_id field of the \$m.

12.12 Reference Classes

References are implemented through several classes:

class Atk4\Data\Reference\HasOne

Defines generic reference, that is typically created by *Model::addReference*

property Atk4\Data\Reference\HasOne::\$tableAlias

Alias for related table. Because multiple references can point to the same table, ability to have unique alias is pretty good.

You don't have to change this property, it is generated automatically.

property Atk4\Data\Reference\HasOne::\$link

What should we pass into owner->ref() to get through to this reference. Each reference has a unique identifier, although it's stored in Model's elements as '#ref-xx'.

property Atk4\Data\Reference\HasOne::\$model

May store reference to related model, depending on implementation.

property Atk4\Data\Reference\HasOne::\$ourField

This is an optional property which can be used by your implementation to store field-level relationship based on a common field matching.

property Atk4\Data\Reference\HasOne::\$their_field

This is an optional property which can be used by your implementation to store field-level relationship based on a common field matching.

Atk4\Data\Reference\HasOne::getModel()

Returns referenced model without conditions.

Atk4\Data\Reference\HasOne::ref()

Returns referenced model WITH conditions. (if possible)

EXPRESSIONS

```
class Atk4\Data\Model
```

You already know that you can define fields inside your Model with `addField`. While a regular field maps to physical field inside your database, sometimes you want to do something different - execute expression or function inside SQL and use result as an output.

Expressions solve this problem by adding a read-only field to your model that corresponds to an expression:

```
Atk4\Data\Model::addExpression($name, $seed)
```

Example will calculate “total_gross” by adding up values for “net” and “vat”:

```
$m = new Model_Invoice($db);
$m->addField('total_net');
$m->addField('total_vat');
$m->addExpression('total_gross', ['expr' => '[total_net] + [total_vat]']);

$m = $m->load(1);

echo $m->get('total_gross');
```

The query using during `load()` will look like this:

```
select
  `id`, `total_net`, `total_vat`,
  (`total_net` + `total_vat`) `total_gross`
from `invoice`
```

13.1 Defining Expression

The simplest format to define expression is by simply passing a string. The argument is executed through `Model::expr()` which automatically substitutes values for the other fields including other expressions.

There are other ways how you can specify expression:

```
$m->addExpression('total_gross', [
  'expr' => $m->expr('[total_net] + [total_vat] + [fee]', ['fee' => $fee]),
]);
```

This format allow you to supply additional parameters inside expression. You should always use parameters instead of appending values inside your expression string (for safety)

You can also use expressions to pass a select action for a specific field:

13.2 No-table Model Expression

Agile Data allows you to define a model without table. While this may have no purpose initially, it does come in handy in some cases, when you need to unite a few statistical queries. Let's start by looking at a very basic example:

```
$m = new Model($db, ['table' => false]);
$m->addExpression('now', ['expr' => 'now()']);
$m = $m->loadAny();
echo $m->get('now');
```

In this example the query will look like this:

```
select (1) `id`, (now()) `now` limit 1
```

so that `$m->getId()` will always be 1 which will make it a model that you can actually use consistently throughout the system. The real benefit from this can be gained when you need to pull various statistical values from your database at once:

```
$m = new Model($db, ['table' => false]);
$m->addExpression('total_orders', ['expr' => (new Model_Order($db))->action('count')]);
$m->addExpression('total_payments', ['expr' => (new Model_Payment($db))->action('count
↵')]);
$m->addExpression('total_received', ['expr' => (new Model_Payment($db))->action('fx0', [
↵'sum', 'amount'])]);

$data = $m->loadOne()->get();
```

Of course you can also use a DSQL for this:

```
$q = $db->dsql();
$q->field(new Model_Order($db)->action('count'), 'total_orders');
$q->field(new Model_Payment($db)->action('count'), 'total_payments');
$q->field(new Model_Payment($db)->action('fx0', ['sum', 'amount']), 'total_received');
$data = $q->getRow();
```

You can decide for yourself based on circumstances.

13.3 Expression Callback

You can use a callback method when defining expression:

```
$m->addExpression('total_gross', ['expr' => function (Model $m, Expression $q) {
    return '[total_net] + [total_vat]';
}, 'type' => 'float']);
```

13.4 Model Reloading after Save

When you add SQL Expressions into your model, that means that some of the fields might be out of sync and you might need your SQL to recalculate those expressions.

To simplify your life, Agile Data implements smart model reloading. Consider the following model:

```
class Model_Math extends \Atk4\Data\Model
{
    public $table = 'math';

    protected function init(): void
    {
        parent::init();

        $this->addField('a');
        $this->addField('b');

        $this->addExpression('sum', ['expr' => '[a] + [b]']);
    }
}

$m = new Model_Math($db);
$m->set('a', 4);
$m->set('b', 6);

$m->save();

echo $m->get('sum');
```

When `$m->save()` is executed, Agile Data will perform reloading of the model. This is to ensure that expression 'sum' would be re-calculated for the values of 4 and 6 so the final line will output a desired result - 10;

Reload after save will only be executed if you have defined any expressions inside your model, however you can affect this behavior:

```
$m = new Model_Math($db, ['reloadAfterSave' => false]);
$m->set('a', 4);
$m->set('b', 6);

$m->save();

echo $m->get('sum'); // outputs null

$m->reload();
echo $m->get('sum'); // outputs 10
```

Now it requires an explicit reload for your model to fetch the result. There is another scenario when your database defines default fields:

```
alter table math change b b int default 10;
```

Then try the following code:

```
class Model_Math extends \Atk4\Data\Model
{
    public $table = 'math';

    protected function init(): void
    {
        parent::init();

        $this->addField('a');
        $this->addField('b');
    }
}

$m = new Model_Math($db);
$m->set('a', 4);

$m->save();

echo $m->get('a')+$m->get('b');
```

This will output 4, because model didn't reload itself due to lack of any expressions. This time you can explicitly enable reload after save:

```
$m = new Model_Math($db, ['reloadAfterSave' => true]);
$m->set('a', 4);

$m->save();

echo $m->get('a')+$m->get('b'); // outputs 14
```

Note: If your model is using `reloadAfterSave`, but you wish to insert data without additional query - use `Model::insert()` or `Model::import()`.

MODEL FROM MULTIPLE JOINED TABLES

```
class Atk4\Data\Model\Join
```

Sometimes model logically contains information that is stored in various places in the database. Your database may want to split up logical information into tables for various reasons, such as to avoid repetition or to better optimize indexes.

14.1 Join Basics

Agile Data allows you to map multiple table fields into a single business model by using joins:

```
$user->addField('username');  
$jContact = $user->join('contact');  
$jContact->addField('address');  
$jContact->addField('county');  
$jContact->hasOne('Country', ['model' => [Country::class]]);
```

This code will load data from two tables simultaneously and if you do change any of those fields they will be update in their respective tables. With SQL the load query would look like this:

```
select  
    u.username, c.address, c.county, c.country_id  
    (select name from country where country.id = c.country_id) country  
from user u  
join contact c on c.id = u.contact_id  
where u.id = $id
```

If driver is unable to query both tables simultaneously, then it will load one record first, then load other record and will collect fields together:

```
$user = $user->load($id);  
$contact = $contact->load($user->get('contact_id'));
```

When saving the record, Joins will automatically record data correctly:

```
insert into contact (address, county, country_id) values ($, $, $);  
@join_c = last_insert_id();  
insert into user (username, contact_id) values ($, @join_c)
```

14.1.1 Strong and Weak joins

When you are joining tables, then by default a strong join is used. That means that both records are not-nullable and when adding records, they will both be added and linked.

Weak join is used if you do not really want to modify the other table. For example it can be used to pull country information based on user.country_id but you wouldn't want that adding a new user would create a new country:

```
$user->addField('username');
$user->addField('country_id', ['type' => 'bigint']);
$jCountry = $user->join('country', ['weak' => true, 'prefix' => 'country_']);
$jCountry->addField('code');
$jCountry->addField('name');
$jCountry->addField('default_currency', ['prefix' => false]);
```

After this you will have the following fields in your model:

- username
- country_id
- country_code [readOnly]
- country_name [readOnly]
- default_currency [readOnly]

14.1.2 Join relationship definitions

When defining joins, you need to outline two fields that must match. In our earlier examples, we the master table was “user” that contained reference to “contact”. The condition would look like this `user.contact_id=contact.id`. In some cases, however, a relation should be reversed:

```
$jContact = $user->join('contact.user_id');
```

This will result in the following join condition: `user.id=contact.user_id`. The first argument to join defines both the table that we need to join and can optionally define the field in the foreign table. If field is set, we will assume that it's a reverse join.

Reverse joins are saved in the opposite order - primary table will be saved first and when id of a primary table is known, foreign table record is stored and ID is supplied. You can pass option 'masterField' to the join() which will specify which field to be used for matching. By default the field is calculated like this: `foreignTable . '_id'`. Here is usage example:

```
$user->addField('username');
$jCreditCard = $user->join('credit_card', [
    'prefix' => 'cc_',
    'masterField' => 'default_credit_card_id',
]);
$jCreditCard->addField('number'); // creates cc_number
$jCreditCard->addField('name'); // creates cc_name
```

Master field can also be specified as an object of a Field class.

There are more options that you can pass inside join(), but those are vendor-specific and you'll have to look into documentation for `sql\Join` and `mongo\Join` respectfully.

14.1.3 Method Proxying

Once your join is defined, you can call several methods on the join objects, that will create fields, other joins or expressions but those would be associated with a foreign table.

`Atk4\Data\Model\Join::addField()`

same as `Model::addField` but associates field with foreign table.

`Atk4\Data\Model\Join::join()`

same as `Model::join` but links new table with this foreign table.

`Atk4\Data\Model\Join::hasOne()`

same as `Model::hasOne` but reference ID field will be associated with foreign table.

`Atk4\Data\Model\Join::hasMany()`

same as `Model::hasMany` but condition for related model will be based on foreign table field and `Reference::$theirField` will be set to `$foreignTable . '_id'`.

`Atk4\Data\Model\Join::containsOne()`

same as `Model::hasOne` but the data will be stored in a field inside foreign table.

Not yet implemented !

`Atk4\Data\Model\Join::containsMany()`

same as `Model::hasMany` but the data will be stored in a field inside foreign table.

Not yet implemented !

14.1.4 Create and Delete behavior

Updating joined records are simple, but when it comes to creation and deletion, there are some conditions. First we look at dependency. If master table contains id of a foreign table, then foreign table record must be created first, so that we can store its ID in a master table. If the join is reversed, the master record is created first and then foreign record is inserted along with the value of master id.

When it comes to deleting record, there are three possible conditions:

1. `[deleteBehavior = cascade, reverse = false]`

If we are using strong join and master table contains ID of foreign table, then foreign master table record is deleted first. Foreign table record is deleted after. This is done to avoid error with foreign constraints.

2. `[deleteBehavior = cascade, reverse = true]`

If we are using strong join and foreign table contains ID of master table, then foreign table record is deleted first followed by the master table record.

3. `[deleteBehavior = ignore, reverse = false]`

If we are using weak join and the master table contains ID of foreign table, then master table is deleted first. Foreign table record is not deleted.

4. `[deleteBehavior = setnull, reverse = true]`

If we are using weak join and foreign table contains ID of master table, then foreign table is updated to set ID of master table to NULL first. Then the master table record is deleted.

Based on the way how you define join an appropriate strategy is selected and Join will automatically decide on `$deleteBehavior` and `$reverse` values. There are situations, however when it's impossible to determine in which order the operations have to be performed. A good example is when you define both master/foreign fields.

In this case system will default to “reverse=false” and will delete master record first, however you can specify a different value for “reverse”.

Sometimes it’s also sensible to set `deleteBehavior = ignore` and perform your own delete operation yourself.

14.1.5 Implementation Detail

Joins are implemented like this:

- all the fields that has ‘joinName’ property set will not be saved into default table by default driver
- join will add either `beforeInsert` or `afterInsert` hook inside your model. When save is executed, it will execute additional query to update foreign table.
- while `$model->getId()` stores the ID of the main table active record, `$join->id` stores ID of the foreign record and will be used when updating.
- option ‘deleteBehavior’ is ‘cascade’ for strong joins and ‘ignore’ for weak joins, but you can set some other value. If you use “setnull” value and you are using reverse join, then foreign table record will not be updated, but value of the foreign field will be set to null.

```
class Atk4\Data\Persistence\Sql\Join
```

14.2 SQL-specific joins

When your model is associated with SQL-capable driver, then instead of using `Join` class, the `Persistence\Sql\Join` is used instead. This class is designed to improve loading technique, because SQL vendors can query multiple tables simultaneously.

Vendors that cannot do JOINS will have to implement compatibility by pulling data from collections in a correct order.

14.2.1 Implementation Details

- although some SQL vendors allow `update .. join ..` syntax, this will not be used. That is done to ensure better compatibility.
- when field has the ‘joinName’ option set, trying to convert this field into expression will prefix the field properly with the foreign table alias.
- join will be added in all queries
- strong join can potentially reduce your data-set as it exclude table rows that cannot be matched with foreign table row.

14.2.2 Specifying complex ON logic

When you’re dealing with SQL drivers, you can specify `Persistence\Sql\Expression` for your “on” clause:

```
$stats = $user->join('stats', [  
    'on' => $user->expr('year({}) = _st.year'),  
    'foreignAlias' => '_st',  
]);
```

You can also specify `'on' => false` then the ON clause will not be used at all and you'll have to add additional `where()` condition yourself.

`foreignAlias` can be specified and will be used as table alias and prefix for all fields. It will default to `'_' . $this->foreignTable`. Agile Data will also resolve situations when multiple tables have same first character so the prefixes will be named `'_c'`, `'_c_2'`, `'_c_3'` etc.

Additional arguments accepted by SQL joins are:

- `'kind'` - will be "inner" for strong join and "left" for weak join, but you can specify other kind of join, for example, "right".

MODEL AGGREGATES

```
class Atk4\Data\Model\AggregateModel
```

In order to create model aggregates the AggregateModel model needs to be used:

15.1 Grouping

AggregateModel model can be used for grouping:

```
$aggregate = new AggregateModel($orders)->setGroupBy(['country_id']);
```

\$aggregate above is a new object that is most appropriate for the model's persistence and which can be manipulated in various ways to fine-tune aggregation. Below is one sample use:

```
$aggregate = new AggregateModel($orders);
$aggregate->addField('country');
$aggregate->setGroupBy(['country_id'], [
    'count' => ['expr' => 'count(*)', 'type' => 'bigint'],
    'total_amount' => ['expr' => 'sum([amount])', 'type' => 'atk4_money'],
],
);

// $aggregate will have following rows:
// ['country' => 'UK', 'count' => 20, 'total_amount' => 123.2];
// ...
```

Below is how opening balance can be built:

```
$ledger = new GeneralLedger($db);
$ledger->addCondition('date', '<', $from);

// we actually need grouping by nominal
$ledger->setGroupBy(['nominal_id'], [
    'opening_balance' => ['expr' => 'sum([amount])', 'type' => 'atk4_money'],
]);
```


HOOKS

Hook is a mechanism for adding callbacks. The core features of Hook sub-system (explained in detail here <https://atk4-core.readthedocs.io/en/develop/hook.html>) include:

- ability to define “spots” in PHP code, such as “beforeLoad”.
- ability to add callbacks to be executed when PHP goes over the spot.
- prioritization of callbacks
- ability to pass arguments to callbacks
- ability to collect response from callbacks
- ability to break hooks (will stop any other hook execution)

Model implements hook trait and defines various hooks which will allow you to execute code before or after various operations, such as save, load etc.

16.1 Model Operation Hooks

All of model operations (adding, updating, loading and deleting) have two hooks - one that executes before operation and another that executes after.

Those hooks are database-agnostic, so regardless where you save your model data, your `beforeSave` hook will be triggered.

If database has transaction support, then hooks will be executed while inside the same transaction:

- begin transaction
- beforeSave hook
- actual save
- reload (see `Model::$reloadAfterSave`)
- afterSave hook
- commit transaction

In case of error:

- do rollback
- call `onRollback` hook

If your `afterSave` hook creates exception, then the entire operation will be rolled back.

16.1.1 Example with beforeSave

The next code snippet demonstrates a basic usage of a `beforeSave` hook. This one will update field values just before record is saved:

```
$m->onHook(Model::HOOK_BEFORE_SAVE, function (Model $entity) {
    $entity->set('name', strtoupper($entity->get('name')));
    $entity->set('surname', strtoupper($entity->get('surname')));
});

$m->insert(['name' => 'John', 'surname' => 'Smith']);

// will save into DB: ['name' => 'JOHN', 'surname' => 'SMITH']
```

16.1.2 Arguments

When you define a callback, then you'll receive reference to model from all the hooks. It's important that you use this argument instead of `$this` to perform operation, otherwise you can run into problems with cloned models.

Callbacks does non expect anything to be returned, but you can modify fields of the model.

16.1.3 Interrupting

You can also break all “before” hooks which will result in cancellation of the original action:

```
$m->breakHook(false);
```

If you break `beforeSave`, then the save operation will not take place, although model will assume the operation was successful.

You can also break `beforeLoad` hook which can be used to skip rows:

```
$model->onHook(Model::HOOK_AFTER_LOAD, function (Model $entity) {
    if ($entity->get('date') < $entity->date_from) {
        $entity->breakHook(false); // will not yield such data row
    }
    // otherwise yields data row
});
```

This will also prevent data from being loaded. If you return false from `afterLoad` hook, then record which we just loaded will be instantly unloaded. This can be helpful in some cases, although you should still use `Model::addCondition` where possible as it is much more efficient.

16.1.4 Insert/Update Hooks

Insert/Update are triggered from inside save() method but are based on current state of `Model::isLoading`:

- `beforeInsert($m, &$data)` (creating new records only)
- `afterInsert($m, $id)`
- `beforeUpdate($m, &$data)` (updating existing records only. Not executed if model is not dirty)
- `afterUpdate($m)`

The `$data` argument will contain array of actual data (field => value) to be saved, which you can use to withdraw certain fields from actually being saved into the database (by unsetting it's value).

Note that altering data via `$m->set()` does not work in `beforeInsert` and `beforeUpdate` hooks, only by altering `$data`.

`afterInsert` will receive either `$id` of new record or null if model couldn't provide ID field. Also, `afterInsert` is actually called before reloading is done (when `Model::$reloadAfterSave` is set).

For some examples, see *Soft Delete*

16.1.5 beforeSave, afterSave Hook

A good place to hook is `beforeSave` as it will be fired when adding new records or modifying existing ones:

- `beforeSave($m)` (saving existing or new records. Not executed if model is not dirty)
- `afterSave($m, $isUpdate)` (same as above, `$isUpdate` is boolean true if it was update and false otherwise)

You might consider "save" to be a higher level hook, as `beforeSave` is called pretty early on during saving the record and `afterSave` is called at the very end of save.

You may actually drop validation exception inside save, insert or update hooks:

```
$m->onHook(Model::HOOK_BEFORE_SAVE, function (Model $entity) {
    if ($entity->get('name') === 'Yagi') {
        throw new \Atk4\Data\ValidationException(['name' => "We don't serve like you"]);
    }
});
```

16.1.6 Loading, Deleting

Those are relatively simple hooks:

- `beforeLoad($m, $id)` (`$m` will be unloaded). Break for custom load or skip.
- `afterLoad($m)`. (`$m` will contain data). Break to unload and skip.

For the deletion it's pretty similar:

- `beforeDelete($m, $id)`. Unload and Break to preserve record.
- `afterDelete($m, $id)`.

A good place to clean-up delete related records would be inside `afterDelete`, although if your database consistency requires those related records to be cleaned up first, use `beforeDelete` instead.

For some examples, see *Soft Delete*

16.1.7 Hook execution sequence

- beforeSave
- beforeInsert [only if insert]
 - beforeInsertQuery [sql only] (query)
 - afterInsertQuery (query, affectedRows)
- beforeUpdate [only if update]
 - beforeUpdateQuery [sql only] (query)
 - afterUpdateQuery (query, affectedRows)
- afterUpdate [only if existing record, model is reloaded]
- afterInsert [only if new record, model not reloaded yet]
- beforeUnload
- afterUnload
- afterSave (bool \$isUpdate) [after insert or update, model is reloaded]

16.1.8 How to prevent actions

In some cases you want to prevent default actions from executing. Suppose you want to check ‘memcache’ before actually loading the record from the database. Here is how you can implement this functionality:

```
$m->onHook(Model::HOOK_BEFORE_LOAD, function (Model $entity, $id) {
    $data = $entity->getApp()->cacheFetch($entity->getModel()->table, $id);
    if ($data) {
        $dataRef = &$entity->getDataRef();
        $dataRef = $data;
        $entity->setId($id);

        $entity->breakHook($entity);
    }
});
```

\$app property is injected through your \$db object and is passed around to all the models. This hook, if successful, will prevent further execution of other beforeLoad hooks and by specifying argument as ‘false’ it will also prevent call to \$persistence for actual loading of the data.

Similarly you can prevent deletion if you wish to implement soft-delete or stop insert/modify from occurring.

16.1.9 onRollback Hook

This hook is executed right after transaction fails and rollback is done. This can be used in various situations.

Save information into auditLog about failure:

```
$m->onHook(Model::HOOK_ROLLBACK, function (Model $m) {
    $m->auditLog->registerFailure();
});
```

Upgrade schema:

```

use Atk4\Data\Persistence\Sql\Exception as SQLException;

$m->onHook(Model::HOOK_ROLLBACK, function (Model $m, \Throwable $exception) {
    if ($exception instanceof SQLException) {
        $m->schema->upgrade();
        $m->breakHook(false); // exception will not be thrown
    }
});

```

In first example we will register failure in audit log, but afterwards still throw exception. In second example we will upgrade model schema and will not throw exception at all because we break hook and return false boolean value.

16.2 Persistence Hooks

Persistence has a few spots which it actually executes through `$model->hook()`, so depending on where you save the data, there are some more hooks available.

16.2.1 Persistence\Sql

Those hooks can be used to affect queries before they are executed. None of these are breakable:

- `beforeUpdateQuery($m, Query $query)`
- `afterUpdateQuery($m, Query $query, int $affectedRows)`. Executed before retrieving data.
- `beforeInsertQuery($m, Query $query)`
- `afterInsertQuery($m, Query $query, int $affectedRows)`. Executed before retrieving data.

The delete has only “before” hook:

- `beforeDeleteQuery($m, Query $query)`

Finally for queries there is hook `initSelectQuery($model, $query, $type)`. It can be used to enhance queries generated by “action” for:

- “count”
- “update”
- “delete”
- “select”
- “field”
- “fx” or “fx0”

16.3 Other Hooks:

ADVANCED TOPICS

Agile Data allow you to implement various tricks.

17.1 SubTypes

Disjoint subtypes is a concept where you give your database just a little bit of OOP by allowing to extend additional types without duplicating columns. For example, if you are implementing “Account” and “Transaction” models. You may want to have multiple transaction types. Some of those types would even require additional fields. The pattern suggest you should add a new table “transaction_transfer” and store extra fields there. In your code:

```
class Transaction_Transfer extends Transaction
{
    protected function init(): void
    {
        parent::init();

        $j = $this->join('transaction_transfer.transaction_id');
        $j->addField('destination_account');
    }
}
```

As you implement single Account and multiple Transaction types, you want to relate both:

```
$account->hasMany('Transactions', ['model' => [Transaction::class]]);
```

There are however two difficulties here:

1. sometimes you want to operate with specific sub-type.
2. when iterating, you want to have appropriate class, not Transaction()

17.1.1 Best practice for specifying relation type

Although there is no magic behind it, I recommend that you use the following code pattern when dealing with multiple types:

```
$account->hasMany('Transactions', ['model' => [Transaction::class]]);
$account->hasMany('Transactions:Deposit', ['model' => [Transaction\Deposit::class]]);
$account->hasMany('Transactions:Transfer', ['model' => [Transaction\Transfer::class]]);
```

You can then use type-specific reference:

```
$account->ref('Transaction:Deposit')->insert(['amount' => 10]);
```

and the code would be clean. If you introduce new type, you would have to add extra line to your “Account” model, but it will not be impacting anything, so that should be pretty safe.

17.1.2 Type substitution on loading

Another technique is for ATK Data to replace your object when data is being loaded. You can treat “Transaction” class as a “shim”:

```
$obj = $account->ref('Transactions')->load(123);
```

Normally \$obj would be instance of Transaction class, however we want this class to be selected based on transaction type. Therefore a more broad record for ‘Transaction’ should be loaded first and then, if necessary, replaced with the correct class transparently, so that the code above would work without a change.

Another scenario which could benefit by type substitution would be:

```
foreach ($account->ref('Transactions') as $tr) {
    echo get_class($tr) . "\n";
}
```

ATK Data allow class substitution during load and iteration by breaking “afterLoad” hook. Place the following inside Transaction::init():

```
$this->onHookShort(Model::HOOK_AFTER_LOAD, function () {
    if (get_class($this) != $this->getClassName()) {
        $cl = $this->getClassName();
        $m = new $cl($this->getModel()->getPersistence());
        $m = $m->load($this->getId());

        $this->breakHook($m);
    }
});
```

You would need to implement method “getClassName” which would return DESIRED class of the record.

17.2 Audit Fields

If you wish to have a certain field inside your models that will be automatically changed when the record is being updated, this can be easily implemented in Agile Data.

I will be looking to create the following fields:

- created_dts
- updated_dts
- created_by_user_id
- updated_by_user_id

To implement the above, I’ll create a new class:

```

class ControllerAudit
{
    use \Atk4\Core\InitializerTrait {
        init as private _init;
    }
    use \Atk4\Core\TrackableTrait;
    use \Atk4\Core\AppScopeTrait;
}

```

TrackableTrait means that I'll be able to add this object inside model with `$model->add(new ControllerAudit())` and that will automatically populate \$owner, and \$app values (due to AppScopeTrait) as well as execute `init()` method, which I want to define like this:

```

protected function init(): void
{
    $this->_init();

    if (isset($this->getOwner()->no_audit)) {
        return;
    }

    $this->getOwner()->addField('created_dts', ['type' => 'datetime', 'default' => new \
    \DateTime()]);

    $this->getOwner()->hasOne('created_by_user_id', ['model' => [User::class]]);
    if (isset($this->getApp()->user) && $this->getApp()->user->isLoaded()) {
        $this->getOwner()->getField('created_by_user_id')->default = $this->getApp()->
    \user->getId();
    }

    $this->getOwner()->hasOne('updated_by_user_id', ['model' => [User::class]]);

    $this->getOwner()->addField('updated_dts', ['type' => 'datetime']);

    $this->getOwner()->onHook(Model::HOOK_BEFORE_UPDATE, function (Model $entity, array
    \>$data) {
        if (isset($this->getApp()->user) && $this->getApp()->user->isLoaded()) {
            $data['updated_by'] = $this->getApp()->user->getId();
        }
        $data['updated_dts'] = new \DateTime();
    });
}

```

In order to add your defined behavior to the model. The first check actually allows you to define models that will bypass audit altogether:

```

$u1 = new Model_User($db); // Model_User::init() includes audit
$u2 = new Model_User($db, ['no_audit' => true]); // will exclude audit features

```

Next we are going to define 'created_dts' field which will default to the current date and time.

The default value for our 'created_by_user_id' field would depend on a currently logged-in user, which would typically be accessible through your application. AppScope allows you to pass \$app around through all the objects, which means

that your Audit Controller will be able to get the current user.

Of course if the application is not defined, no default is set. This would be handy for unit tests where you could manually specify the value for this field.

The last 2 fields (update_*) will be updated through a hook - beforeUpdate() and will provide the values to be saved during save(). beforeUpdate() will not be called when new record is inserted, so those fields will be left as “null” after initial insert.

If you wish, you can modify the code and insert historical records into other table.

17.3 Soft Delete

Most of the data frameworks provide some way to enable ‘soft-delete’ for tables as a core feature. Design of Agile Data makes it possible to implement soft-delete through external controller. There may be a 3rd party controller for comprehensive soft-delete, but in this section I’ll explain how you can easily build your own soft-delete controller for Agile Data (for educational purposes).

Start by creating a class:

```
class ControllerSoftDelete
{
    use \Atk4\Core\InitializerTrait {
        init as private _init;
    }
    use \Atk4\Core\TrackableTrait;

    protected function init(): void
    {
        $this->_init();

        if (property_exists($this->getOwner(), 'no_soft_delete')) {
            return;
        }

        $this->getOwner()->addField('is_deleted', ['type' => 'boolean']);

        if (property_exists($this->getOwner(), 'deleted_only') && $this->getOwner()->
        deleted_only) {
            $this->getOwner()->addCondition('is_deleted', true);
            $this->getOwner()->addMethod('restore', \Closure::fromCallable([$this,
        restore]));
        } else {
            $this->getOwner()->addCondition('is_deleted', false);
            $this->getOwner()->addMethod('softDelete', \Closure::fromCallable([$this,
        softDelete]));
        }
    }

    public function softDelete(Model $entity)
    {
        $entity->assertIsLoaded();

        $id = $entity->getId();
    }
}
```

(continues on next page)

(continued from previous page)

```

    if ($entity->hook('beforeSoftDelete') === false) {
        return $entity;
    }

    $entity->saveAndUnload(['is_deleted' => true]);

    $entity->hook('afterSoftDelete', [$id]);

    return $entity;
}

public function restore(Model $entity)
{
    $entity->assertIsLoaded();

    $id = $entity->getId();
    if ($entity->hook('beforeRestore') === false) {
        return $entity;
    }

    $entity->saveAndUnload(['is_deleted' => false]);

    $entity->hook('afterRestore', [$id]);

    return $entity;
}
}

```

This implementation of soft-delete can be turned off by setting model's property 'deleted_only' to true (if you want to recover a record).

When active, a new field will be defined 'is_deleted' and a new dynamic method will be added into a model, allowing you to do this:

```

$m = new Model_Invoice($db);
$m = $m->load(10);
$m->softDelete();

```

The method body is actually defined in our controller. Notice that we have defined 2 hooks - beforeSoftDelete and afterSoftDelete that work similarly to beforeDelete and afterDelete.

beforeSoftDelete will allow you to "break" it in certain cases to bypass the rest of method, again, this is to maintain consistency with the rest of before* hooks in Agile Data.

Hooks are called through the model, so your call-back will automatically receive first argument \$m, and afterSoftDelete will pass second argument - \$id of deleted record.

I am then setting reloadAfterSave value to false, because after I set 'is_deleted' to false, \$m will no longer be able to load the record - it will fall outside of the DataSet. (We might implement a better method for saving records outside of DataSet in the future).

After softDelete active record is unloaded, mimicking behavior of delete().

It's also possible for you to easily look at deleted records and even restore them:

```
$m = new Model_Invoice($db, ['deleted_only' => true]);
$m = $m->load(10);
$m->restore();
```

Note that you can call `$m->delete()` still on any record to permanently delete it.

17.3.1 Soft Delete that overrides default delete()

In case you want `$m->delete()` to perform soft-delete for you - this can also be achieved through a pretty simple controller. In fact I'm reusing the one from before and just slightly modifying it:

```
class ControllerSoftDelete2 extends ControllerSoftDelete
{
    protected function init(): void
    {
        parent::init();

        $this->getOwner()->onHook(Model::HOOK_BEFORE_DELETE, \Closure::fromCallable([
            =>$this, 'softDelete']), null, 100);
    }

    public function softDelete(Model $entity)
    {
        parent::softDelete();

        $entity->hook(Model::HOOK_AFTER_DELETE);

        $entity->breakHook(false); // this will cancel original delete()
    }
}
```

Implementation of this controller is similar to the one above, however instead of creating `softDelete()` it overrides the `delete()` method through a hook. It will still call 'afterDelete' to mimic the behavior of regular `delete()` after the record is marked as deleted and unloaded.

You can still access the deleted records:

```
$m = new Model_Invoice($db, ['deleted_only' => true]);
$m = $m->load(10);
$m->restore();
```

Calling `delete()` on the model with 'deleted_only' property will delete it permanently.

17.4 Creating Unique Field

Database can has UNIQUE constraint, but this does work if you use DataSet. For instance, you may be only able to create one 'Category' with name 'Book', but what if there is a soft-deleted record with same name or record that belongs to another user?

With Agile Data you can create controller that will ensure that certain fields inside your model are unique:

```
class ControllerUniqueFields
{
    use \Atk4\Core\InitializerTrait {
        init as private _init;
    }
    use \Atk4\Core\TrackableTrait;

    protected ?array $fields = null;

    protected function init(): void
    {
        $this->_init();

        // by default make 'name' unique
        if (!$this->fields) {
            $this->fields = [$this->getOwner()->titleField];
        }

        $this->getOwner()->onHook(Model::HOOK_BEFORE_SAVE, \Closure::fromCallable([$this,
→ 'beforeSave']));
    }

    protected function beforeSave(Model $entity)
    {
        foreach ($this->fields as $field) {
            if ($entity->getDirtyRef()[$field]) {
                $modelCloned = clone $entity->getModel();
                $modelCloned->addCondition($entity->idField != $this->id);
                $entityCloned = $modelCloned->tryLoadBy($field, $entity->get($field));

                if ($entityCloned !== null) {
                    throw (new \Atk4\Data\Exception('Duplicate record exists'))
                        ->addMoreInfo('field', $field)
                        ->addMoreInfo('value', $entity->get($field));
                }
            }
        }
    }
}
```

As expected - when you add a new model the new values are checked against existing records. You can also slightly modify the logic to make addCondition additive if you are verifying for the combination of matched fields.

17.5 Using WITH cursors

Many SQL database engines support defining WITH cursors to use in select, update and even delete statements.

```
class Atk4\Data\Model
```

```
Atk4\Data\Model::addCteModel(string $name, Model $model, bool $recursive = false)
```

Agile toolkit data models also support these cursors. Usage is like this:

```
$invoices = new Invoice();

$contacts = new Contact();
$contacts->addCteModel('inv', $invoices);
$contacts->join('inv.cid');
```

```
with
  `inv` as (select `contact_id`, `ref_no`, `total_net` from `invoice`)
select
  *
from `contact`
  join `inv` on `inv`.`contact_id`=`contact`.`id`
```

Note: Supported since MySQL 8.x, MariaDB supported it earlier.

17.6 Creating Many to Many relationship

Depending on the use-case many-to-many relationships can be implemented differently in Agile Data. I will be focusing on the practical approach. My system has “Invoice” and “Payment” document and I’d like to introduce “invoice_payment” that can link both entities together with fields (‘invoice_id’, ‘payment_id’, and ‘amount_closed’). Here is what I need to do:

17.6.1 1. Create Intermediate Entity - InvoicePayment

Create new Model:

```
class Model_InvoicePayment extends \Atk4\Data\Model
{
    public $table = 'invoice_payment';

    protected function init(): void
    {
        parent::init();

        $this->hasOne('invoice_id', ['model' => [Model_Invoice::class]]);
        $this->hasOne('payment_id', ['model' => [Model_Payment::class]]);
        $this->addField('amount_closed');
    }
}
```

17.6.2 2. Update Invoice and Payment model

Next we need to define reference. Inside Model_Invoice add:

```
$this->hasMany('InvoicePayment', ['model' => [Model_InvoicePayment::class]]);

$this->hasMany('Payment', ['model' => static function (Persistence $persistence) {
    $p = new Model_Payment($persistence);
    $j = $p->join('invoice_payment.payment_id');
    $j->addField('amount_closed');
    $j->hasOne('invoice_id', ['model' => [Model_Invoice::class]]);

    return $p;
}, 'theirField' => 'invoice_id']);

$this->onHookShort(Model::HOOK_BEFORE_DELETE, function () {
    foreach ($this->ref('InvoicePayment') as $payment) {
        $payment->delete();
    }
});
```

You'll have to do a similar change inside Payment model. The code for '\$j->' have to be duplicated until we implement method Join->importModel().

17.6.3 3. How to use

Here are some use-cases. First lets add payment to existing invoice. Obviously we cannot close amount that is bigger than invoice's total:

```
$i->ref('Payment')->insert([
    'amount' => $paid,
    'amount_closed' => min($paid, $i->get('total')),
    'payment_code' => 'XYZ',
]);
```

Having some calculated fields for the invoice is handy. I'm adding total_payments that shows how much amount is closed and amount_due:

```
// define field to see closed amount on invoice
$this->hasMany('InvoicePayment', ['model' => [Model_InvoicePayment::class]])
    ->addField('total_payments', ['aggregate' => 'sum', 'field' => 'amount_closed']);
$this->addExpression('amount_due', ['expr' => '[total] - coalesce([total_payments], 0)
↪'];
```

Note that I'm using coalesce because without InvoicePayments the aggregate sum will return NULL. Finally let's build allocation method, that allocates new payment towards a most suitable invoice:

```
// add to Model_Payment
public function autoAllocate()
{
    $client = $this->ref['client_id'];
    $invoices = $client->ref('Invoice');
```

(continues on next page)

(continued from previous page)

```

// we are only interested in unpaid invoices
$invoices->addCondition('amount_due', '>', 0);

// prioritize older invoices
$invoices->setOrder('date');

while ($this->get('amount_due') > 0) {
    // see if any invoices match by 'reference'
    $invoice = $invoices->tryLoadBy('reference', $this->get('reference'));

    if ($invoice === null) {
        // otherwise load any unpaid invoice
        $invoice = $invoices->tryLoadAny();

        if ($invoice === null) {
            // couldn't load any invoice
            return;
        }
    }

    // how much we can allocate to this invoice
    $alloc = min($this->get('amount_due'), $invoice->get('amount_due'))
    $this->ref('InvoicePayment')->insert(['amount_closed' => $alloc, 'invoice_id' =>
->$invoice->getId()]);

    // reload ourselves to refresh amount_due
    $this->reload();
}
}

```

The method here will prioritize oldest invoices unless it finds the one that has a matching reference. Additionally it will allocate your payment towards multiple invoices. Finally if invoice is partially paid it will only allocate what is due.

17.7 Creating Related Entity Lookup

Sometimes when you add a record inside your model you want to specify some related records not through ID but through other means. For instance, when adding invoice, I want to make it possible to specify 'Category' through the name, not only category_id. First, let me illustrate how can I do that with category_id:

```

class Model_Invoice extends \Atk4\Data\Model
{
    protected function init(): void
    {
        parent::init();

        ...

        $this->hasOne('category_id', ['model' => [Model_Category::class]]);

        ...
    }
}

```

(continues on next page)

(continued from previous page)

```

    }
}

$m = new Model_Invoice($db);
$m->insert(['total' => 20, 'client_id' => 402, 'category_id' => 6]);

```

So in situations when client_id and category_id is not known (such as import or API call) this approach will require us to perform 2 extra queries:

```

$m = new Model_Invoice($db);
$m->insert([
    'total' => 20,
    'client_id' => $m->ref('client_id')->loadBy('code', $clientCode)->getId(),
    'category_id' => $m->ref('category_id')->loadBy('name', $category)->getId(),
]);

```

The ideal way would be to create some “non-persistable” fields that can be used to make things easier:

```

$m = new Model_Invoice($db);
$m->insert([
    'total' => 20,
    'client_code' => $clientCode,
    'category' => $category,
]);

```

Here is how to add them. First you need to create fields:

```

$this->addField('client_code', ['neverPersist' => true]);
$this->addField('client_name', ['neverPersist' => true]);
$this->addField('category', ['neverPersist' => true]);

```

I have declared those fields with neverPersist so they will never be used by persistence layer to load or save anything. Next I need a beforeSave handler:

```

$this->onHookShort(Model::HOOK_BEFORE_SAVE, function () {
    if ($this->isDirty('client_code') && !$this->isDirty('client_id')) {
        $client = $this->ref('client_id');
        $this->set('client_id', $client->getId());
    }

    if ($this->isDirty('client_name') && !$this->isDirty('client_id')) {
        $client = $this->ref('client_id');
        $this->set('client_id', $client->getId());
    }

    if ($this->isDirty('category') && !$this->isDirty('category_id')) {
        $category = $this->ref('category_id');
        $this->set('category_id', $category->getId());
    }
});

```

Note that isset() here will be true for modified fields only and behaves differently from PHP’s default behavior. See documentation for Model::isset

This technique allows you to hide the complexity of the lookups and also embed the necessary queries inside your “insert” query.

17.8 Inserting Hierarchical Data

In this example I’ll be building API that allows me to insert multi-model information. Here is usage example:

```
$invoice->insert([
  'client' => 'Joe Smith',
  'payment' => [
    'amount' => 15,
    'ref' => 'half upfront',
  ],
  'lines' => [
    ['descr' => 'Book', 'qty' => 3, 'price' => 5]
    ['descr' => 'Pencil', 'qty' => 1, 'price' => 10]
    ['descr' => 'Eraser', 'qty' => 2, 'price' => 2.5],
  ],
]);
```

Not only ‘insert’ but ‘set’ and ‘save’ should be able to use those fields for ‘payment’ and ‘lines’, so we need to first define those as ‘neverPersist’. If you curious about client lookup by-name, I have explained it in the previous section. Add this into your Invoice Model:

```
$this->addField('payment', ['neverPersist' => true]);
$this->addField('lines', ['neverPersist' => true]);
```

Next both payment and lines need to be added after invoice is actually created, so:

```
$this->onHookShort(Model::HOOK_AFTER_SAVE, function (bool $isUpdate) {
  if ($this->isDirty('payment')) {
    $this->ref('Payment')->insert($this->get('payment'));
  }

  if ($this->isDirty('lines')) {
    $this->ref('Line')->import($this->get('lines'));
  }
});
```

You should never call save() inside afterSave hook, but if you wish to do some further manipulation, you can reload a clone:

```
$entityCloned = clone $entity;
$entityCloned->reload();
if ($entityCloned->get('amount_due') == 0) {
  $entityCloned->save(['status' => 'paid']);
}
```

17.9 Related Record Conditioning

Sometimes you wish to extend one Model into another but related field type can also change. For example let's say we have Model_Invoice that extends Model_Document and we also have Model_Client that extends Model_Contact.

In theory Document's 'contact_id' can be any Contact, however when you create 'Model_Invoice' you wish that 'contact_id' allow only Clients. First, lets define Model_Document:

```
$this->hasOne('client_id', ['model' => [Model_Contact::class]]);
```

One option here is to move 'Model_Contact' into model property, which will be different for the extended class:

```
$this->hasOne('client_id', ['model' => [$this->client_class]]);
```

Alternatively you can replace model in the init() method of Model_Invoice:

```
$this->getReference('client_id')->model = [Model_Client::class];
```

You can also use array here if you wish to pass additional information into related model:

```
$this->getReference('client_id')->model = ['Model_Client', 'no_audit' => true];
```

Combined with our "Audit" handler above, this should allow you to relate with deleted clients.

17.10 Narrowing Down Existing References

Agile Data allow you to define multiple references between same entities, but sometimes that can be quite useful. Consider adding this inside your Model_Contact:

```
$this->hasMany('Invoice', ['model' => [Model_Invoice::class]]);

$this->hasMany('OverdueInvoice', ['model' => static function (Persistence $persistence) {
    return (new Model_Invoice($persistence))
        ->addCondition('due', '<', new \DateTime());
}]);
```

This way if you extend your class into 'Model_Client' and modify the 'Invoice' reference to use different model:

```
$this->getReference('Invoice')->model = [Model_Invoice_Sale::class];
```

The 'OverdueInvoice' reference will be also properly adjusted.

LOADING AND SAVING CSV FILES

`class Atk4\Data\Persistence\Csv`

Agile Data can operate with CSV files for data loading, or saving. The capabilities of `Persistence\Csv` are limited to the following actions:

- open any CSV file, use column mapping
- identify which column is corresponding for respective field
- support custom mapper, e.g. if date is stored in a weird way
- support for CSV files that have extra lines on top/bottom/etc
- listing/iterating
- adding a new record

18.1 Setting Up

When creating new persistence you must provide a valid URL for the file that might be stored either on a local system or use a remote file scheme (ftp://...). The file will not be actually opened unless you perform load/save operation:

```
$p = new Persistence\Csv('myfile.csv');  
  
$u = new Model_User($p);  
$u = $u->tryLoadAny(); // actually opens file and finds first record
```

18.2 Exporting and Importing data from CSV

You can take a model that is loaded from other persistence and save it into CSV like this. The next example demonstrates a basic functionality of SQL database export to CSV file:

```
$db = new Persistence\Sql($connection);  
$csv = new Persistence\Csv('dump.csv');  
  
$m = new Model_User($db);  
  
foreach (new Model_User($db) as $m) {  
    $m->withPersistence($csv)->save();  
}
```

Theoretically you can do few things to tweak this process. You can specify which fields you would like to see in the CSV:

```
foreach (new Model_User($db) as $m) {
    $m->withPersistence($csv)
        ->setOnlyFields(['id', 'name', 'password'])
        ->save();
}
```

Additionally if you want to use a different column titles, you can:

```
foreach (new Model_User($db) as $m) {
    $mCsv = $m->withPersistence($csv);
    $mCsv->setOnlyFields(['id', 'name', 'password'])
    $mCsv->getField('name')->actual = 'First Name';
    $mCsv->save();
}
```

Like with any other persistence you can use typecasting if you want data to be stored in any particular format.

The examples above also create object on each iteration, that may appear as a performance inefficiency. This can be solved by re-using Csv model through iterations:

```
$m = new Model_User($db);
$mCsv = $m->withPersistence($csv);
$mCsv->setOnlyFields(['id', 'name', 'password'])
$mCsv->getField('name')->actual = 'First Name';

foreach ($m as $mCsv) {
    $mCsv->save();
}
```

This code can be further simplified if you use import() method:

```
$m = new Model_User($db);
$mCsv = $m->withPersistence($csv);
$mCsv->setOnlyFields(['id', 'name', 'password'])
$mCsv->getField('name')->actual = 'First Name';
$mCsv->import($m);
```

Naturally you can also move data in the other direction:

```
$m = new Model_User($db);
$mCsv = $m->withPersistence($csv);
$mCsv->setOnlyFields(['id', 'name', 'password'])
$mCsv->getField('name')->actual = 'First Name';

$m->import($mCsv);
```

Only the last line changes and the data will now flow in the other direction.

AGILE DSQL DOCUMENTATION

Contents:

19.1 Overview

DSQL is a dynamic SQL query builder. You can write multi-vendor queries in PHP profiting from better security, clean syntax and most importantly – sub-query support. With DSQL you stay in control of when queries are executed and what data is transmitted. DSQL is easily composable – build one query and use it as a part of other query.

19.1.1 Goals of DSQL

- simple and concise syntax
- consistently scalable (e.g. 5 levels of sub-queries, 10 with joins and 15 parameters? no problem)
- “One Query” paradigm
- support for PDO vendors as well as NoSQL databases (with query language similar to SQL)
- small code footprint (over 50% less than competing frameworks)
- free, licensed under MIT
- no dependencies
- follows design paradigms:
 - “PHP the Agile way”
 - “Functional ORM”
 - “Open to extend”
 - “Vendor Transparency”

19.1.2 DSQL by example

The simplest way to explain DSQL is by example:

```
$query = $connection->dsql();
$query->table('employees')
    ->where('birth_date', '1961-05-02')
    ->field('count(*)');
echo 'Employees born on May 2, 1961: ' . $query->getOne();
```

The above code will execute the following query:


```
select count(*) from `salary` where `birth_date` = :a
:a = "1961-05-02"
```

DSQL can also execute queries with multiple sub-queries, joins, expressions grouping, ordering, unions as well as queries on result-set.

- See *Quickstart* if you would like to start learning DSQL.

19.1.3 DSQL is Part of Agile Toolkit

DSQL is a stand-alone and lightweight library with no dependencies and can be used in any PHP project, big or small.



DSQL is also a part of Agile Toolkit framework and works best with Agile Models. Your project may benefit from a higher-level data abstraction layer, so be sure to look at the rest of the suite.

19.2 Quickstart

When working with DSQL you need to understand the following basic concepts:

19.2.1 Basic Concepts

- Expression (see `expr`)

Expression object, represents a part of a SQL query. It can be used to express advanced logic in some part of a query, which *Query* itself might not support or can express a full statement. Never try to look for “raw” queries, instead build expressions and think about escaping.

- Query (see `query`)

Object of a *Query* class can be used for building and executing valid SQL statements such as SELECT, INSERT, UPDATE, etc. After creating *Query* object you can call various methods to add “table”, “where”, “from” parts of your query.

- Connection

Represents a connection to the database. If you already have a PDO object you can feed it into *Expression* or *Query*, but for your comfort there is a *Connection* class with very little overhead.

19.2.2 Getting Started

We will start by looking at the *Query* building, because you do not need a database to create a query:

```
$query = $connection->dsql();
```

Once you have a query object, you can add parameters by calling some of it's methods:

```
$query
->table('employees')
->where('birth_date', '1961-05-02')
->field('count(*)');
```

Finally you can get the data:

```
$count = $query->getOne();
```

While DSQL is simple to use for basic queries, it also gives a huge power and consistency when you are building complex queries. Unlike other query builders that sometimes rely on “hacks” (such as method `whereOr()`) and claim to be useful for “most” database operations, with DSQL, you can use DSQL to build ALL of your database queries.

This is hugely beneficial for frameworks and large applications, where various classes need to interact and inject more clauses/fields/joins into your SQL query.

DSQL does not resolve conflicts between similarly named tables, but it gives you all the options to use aliases.

The next example might be a bit too complex for you, but still read through and try to understand what each section does to your base query:

```
// establish a query looking for a maximum salary
$salary = $connection->dsql();

// create few expression objects
$maxSalary = $salary->expr('max(salary)');
$months = $salary->expr('TimeStampDiff(month, from_date, to_date)');

// configure our basic query
$salary
->table('salary')
->field(['emp_no', 'max_salary' => $maxSalary, 'months' => $months])
->group('emp_no')
->order('-max_salary');

// define sub-query for employee "id" with certain birth-date
$employees = $salary->dsql()
->table('employees')
->where('birth_date', '1961-05-02')
->field('emp_no');

// use sub-select to condition salaries
```

(continues on next page)

(continued from previous page)

```

$salary->where('emp_no', $employees);

// join with another table for more data
$salary
  ->join('employees.emp_id')
  ->field('employees.first_name');

// finally, fetch result
foreach ($salary as $row) {
  echo 'Data: ' . json_encode($row) . "\n";
}

```

The above query resulting code will look like this:

```

SELECT
  `emp_no`,
  max(salary) `max_salary`,
  TimeStampDiff(month, from_date, to_date) `months`
FROM
  `salary`
JOIN
  `employees` on `employees`.`emp_id` = `salary`.`emp_id`
WHERE
  `salary`.`emp_no` in(select `id` from `employees` where `birth_date` = :a)
GROUP BY `emp_no`
ORDER BY max_salary desc

:a = "1961-05-02"

```

Using DSQL in higher level ORM libraries and frameworks allows them to focus on defining the database logic, while DSQL can perform the heavy-lifting of query building and execution.

19.2.3 Creating Objects and PDO

DSQL classes does not need database connection for most of it's work. Once you create new instance of *Expression* or *Query* you can perform operation and finally call *Expression::render()* to get the final query string with params:

```

use Atk4\Data\Persistence\Sql\Query;

$q = (new Query())->table('user')->where('id', 1)->field('name');
[$query, $params] = $q->render();

```

When used in application you would typically generate queries with the purpose of executing them, which makes it very useful to create a *Connection* object. The usage changes slightly:

```

$c = Atk4\Data\Persistence\Sql\Connection::connect($dsn, $user, $password);
$q = $c->dsql()->table('user')->where('id', 1)->field('name');

$name = $q->getOne();

```

You no longer need “use” statement and *Connection* class will automatically do some of the hard work to adopt query building for your database vendor. There are more ways to create connection, see Advanced Connections section.

The format of the `$dsn` is the same as with for [DBAL connection](#). If you need to execute query that is not supported by DSQL, you should always use expressions:

```
$tables = $c->expr('show tables like []', [$likeStr])->getRows();
```

DSQL classes are mindful about your SQL vendor and its quirks, so when you're building sub-queries with `Query::dsl`, you can avoid some nasty problems:

```
$sqliteConnection->dsl()->table('user')->mode('truncate')->executeStatement();
```

The above code will work even though SQLite does not support truncate. That's because DSQL takes care of this.

19.2.4 Query Building

Each `Query` object represents a query to the database in-the-making. Calling methods such as `Query::table` or `Query::where` affect part of the query you're making. At any time you can either execute your query or use it inside another query.

`Query` supports majority of SQL syntax out of the box. Some unusual statements can be easily added by customizing template for specific query and we will look into examples in [Extending Query Class](#)

19.2.5 Query Mode

When you create a new `Query` object, it is going to be a `SELECT` query by default. If you wish to execute update operation instead, you can simply call `Query::mode` to change it. For more information see [Query Modes](#). You can actually perform multiple operations:

```
$q = $c->dsl()->table('employee')->where('emp_no', 1234);
$backupData = $q->getRows();
$q->mode('delete')->executeStatement();
```

A good practice is to re-use the same query object before you branch out and perform the action:

```
$q = $c->dsl()->table('employee')->where('emp_no', 1234);

if ($confirmed) {
    $q->mode('delete')->executeStatement();
} else {
    echo 'Are you sure you want to delete ' . $q->field('count(*)') . ' employees?';
}
```

19.2.6 Fetching Result

When you are selecting data from your database, DSQL will prepare and execute statement for you. Depending on the connection, there may be some magic involved, but once the query is executed, you can start streaming your data:

```
foreach ($query->table('employee')->where('dep_no', 123) as $employee) {
    echo $employee['first_name'] . "\n";
}
```

When iterating you'll have `Doctrine\DBAL\Result`. Remember that DSQL can support vendors, `$employee` will always contain associative array representing one row of data. (See also [Manual Query Execution](#)).

19.3 Connection

DSQL supports various database vendors natively but also supports 3rd party extensions. For current status on database support see: *Vendor support and Extensions*.

class Atk4\Data\Persistence\Sql\Connection

Connection class is handy to have if you plan on building and executing queries in your application. It's more appropriate to store connection in a global variable or global class:

```
$app->db = Atk4\Data\Persistence\Sql\Connection::connect($dsn, $user, $pass, $defaults);
```

```
static Atk4\Data\Persistence\Sql\Connection::connect($dsn, $user = null, $password = null, $defaults = [])
```

Determine which Connection class should be used for specified \$dsn, establish connection to DB by creating new object of this connection class and return.

Parameters

- **\$dsn** (string) – DSN, see <https://www.doctrine-project.org/projects/doctrine-dbal/en/latest/reference/configuration.html>
- **\$user** (string) – username
- **\$password** (string) – password
- **\$defaults** (array) – Other default properties for connection class.

Returns new Connection

This should allow you to access this class from anywhere and generate either new Query or Expression class:

```
$query = $app->db->dsql();  
  
// or  
  
$expr = $app->db->expr('show tables');
```

Atk4\Data\Persistence\Sql\Connection::**expr**(\$template, \$arguments)

Creates new Expression class and sets *Expression::\$connection*.

Parameters

- **\$arguments** (array) – Other default properties for connection class.

Returns new Expression

Atk4\Data\Persistence\Sql\Connection::**dsql**(\$defaults)

Creates new Query class and sets *Query::\$connection*.

Parameters

- **\$defaults** (array) – Other default properties for connection class.

Returns new Query

Here is how you can use all of this together:

```

$dsn = 'mysql:host=localhost;port=3307;dbname=testdb';

$connection = Atk4\Data\Persistence\Sql\Connection::connect($dsn, 'root', 'root');

echo 'Time now is: ' . $connection->expr('select now()');

```

`connect` will determine appropriate class that can be used for this DSN string. This can be a PDO class or it may try to use a 3rd party connection class.

Connection class is also responsible for executing queries. This is only used if you connect to vendor that does not use PDO.

`Atk4\Data\Persistence\Sql\Connection::execute(Expression $expr) → \Doctrine\DBAL\Result`

Creates new Expression class and sets `Expression::$connection`.

Parameters

- **\$expr** (*Expression*) – Expression (or query) to execute

Returns *Doctrine\DBAL\Result*

`Atk4\Data\Persistence\Sql\Connection::registerConnectionClass($connectionClass, $connectionType)`

Adds connection class to the registry for resolving in `Connection::resolveConnectionClass` method.

Parameters

- **\$connectionClass** (string) – The connection class to be used for the driver type
- **\$connectionType** (string) – Alias of the connection

Developers can register custom classes to handle driver types using the `Connection::registerConnectionClass` method:

```

Connection::registerConnectionClass(Custom\MySQL\Connection::class, 'pdo_mysql');

```

`Atk4\Data\Persistence\Sql\Connection::connectDbalConnection(array $dsn)`

The method should establish connection with DB and return the underlying connection object used by the Connection class. By default PDO is used but the method can be overridden to return custom object to be used for connection to DB.

class `Atk4\Data\Persistence\Sql\Expression`

19.4 Expressions

Expression class implements a flexible way for you to define any custom expression then execute it as-is or as a part of another query or expression. Expression is supported anywhere in DSQL to allow you to express SQL syntax properly.

Quick Example:

```

$query->where('time', $query->expr(
    'between "[]" and "[]"',
    [$fromTime, $toTime]
));

// produces: .. where `time` between :a and :b

```

Another use of expression is to supply field instead of value and vice versa:

```
$query->where($query->expr(
    '[] between time_from and time_to',
    [$time]
));

// produces: where :a between time_from and time_to
```

Yet another curious use for the DSQL library is if you have certain object in your ORM implementing `Expressionable` interface. Then you can also use it within expressions:

```
$query->where($query->expr(
    '[] between [] and []',
    [$time, $model->getElement('time_form'), $model->getElement('time_to')]
));

// produces: where :a between `time_from` and `time_to`
```

Another uses for expressions could be:

- Sub-Queries
- SQL functions, e.g. IF, CASE
- nested AND / OR clauses
- vendor-specific queries - “describe table”
- non-traditional constructions, UNIONS or SELECT INTO

19.4.1 Properties, Arguments, Parameters

Be careful when using those similar terms as they refer to different things:

- Properties refer to object properties, e.g. `$expr->template`, see *Other Properties*
- Arguments refer to template arguments, e.g. `select * from [table]`, see *Expression Template*
- Parameters refer to the way of passing user values within a query `where id=:a` and are further explained below.

Parameters

Because some values are un-safe to use in the query and can contain dangerous values they are kept outside of the SQL query string and are using PDO's `bindValue` instead. DSQL can consist of multiple objects and each object may have some parameters. During rendering those parameters are joined together to produce one complete query.

19.4.2 Creating Expression

```
$expr = $connection->expr('NOW()');
```

You can also use `expr()` method to create expression, in which case you do not have to define “use” block:

```
$query->where('time', '>', $query->expr('NOW()));  
  
// produces: .. where `time` > NOW()
```

You can specify some of the expression properties through first argument of the constructor:

```
$expr = $connection->expr(['template' => 'NOW()']);
```

Scroll down for full list of properties.

19.4.3 Expression Template

When you create a template the first argument is the template. It will be stored in `$template` property. Template string can contain arguments in a square brackets:

- `coalesce([], [])` is same as `coalesce([0], [1])`
- `coalesce([one], [two])`

Arguments can be specified immediately through an array as a second argument into constructor or you can specify arguments later:

```
$expr = $connection->expr(  
    'coalesce([name], [surname])',  
    ['name' => $name, 'surname' => $surname]  
);  
  
// is the same as  
  
$expr = $connection->expr('coalesce([name], [surname])');  
$expr['name'] = $name;  
$expr['surname'] = $surname;
```

19.4.4 Nested expressions

Expressions can be nested several times:

```
$age = $connection->expr('coalesce([age], [default_age])');  
$age['age'] = $connection->expr("year(now()) - year(birth_date)");  
$age['default_age'] = 18;  
  
$query->table('user')->field($age, 'calculated_age');  
  
// select coalesce(year(now()) - year(birth_date), :a) `calculated_age` from `user`
```

When you include one query into another query, it will automatically take care of all user-defined parameters (such as value 18 above) which will make sure that SQL injections could not be introduced at any stage.

19.4.5 Rendering

An expression can be rendered into a valid SQL code by calling `render()` method. The method will return an array with string and params.

`Atk4\Data\Persistence\Sql\Expression::render()`

Converts *Expression* object to an array with string and params. Parameters are replaced with `:a`, `:b`, etc.

19.4.6 Executing Expressions

If your expression is a valid SQL query, (such as `show databases`) you might want to execute it. Expression class offers you various ways to execute your expression. Before you do, however, you need to have `$connection` property set. (See [Connecting to Database](#) on more details). In short the following code will connect your expression with the database:

```
$expr = $connection->expr();
```

If you are looking to use connection *Query* class, you may want to consider using a proper vendor-specific subclass:

```
$query = new \Atk4\Data\Persistence\Sql\Mysql\Query('connection' => $connection);
```

Finally, you can pass connection class into `executeQuery` directly.

`Atk4\Data\Persistence\Sql\Expression::executeQuery($connection = null)`

Executes expression using current database connection or the one you specify as the argument:

```
$stmt = $expr->executeQuery($connection);
```

returns `Doctrine\DBAL\Result`.

`Atk4\Data\Persistence\Sql\Expression::expr($template, $arguments)`

Creates a new *Expression* object that will inherit current `$connection` property. Also if you are creating a vendor-specific expression/query support, this method must return instance of your own version of Expression class.

The main principle here is that the new object must be capable of working with database connection.

`Atk4\Data\Persistence\Sql\Expression::getRows()`

Executes expression and return whole result-set in form of array of hashes:

```
$data = $connection->expr('show databases')->getRows();  
echo json_encode($data);
```

The output would be

```
[  
  { "Database": "mydb1" },  
  { "Database": "mysql" },  
  { "Database": "test" },  
]
```

`Atk4\Data\Persistence\Sql\Expression::getRow()`

Executes expression and returns first row of data from result-set as a hash:

```
$data = $connection->expr('SELECT @@global.time_zone, @@session.time_zone')->
↳getRow()

echo json_encode($data);
```

The output would be

```
{ "@@global.time_zone": "SYSTEM", "@@session.time_zone": "SYSTEM" }
```

Atk4\Data\Persistence\Sql\Expression::getOne()

Executes expression and return first value of first row of data from result-set:

```
$time = $connection->expr('NOW()')->getOne();
```

19.4.7 Magic an Debug Methods

Atk4\Data\Persistence\Sql\Expression::__debugInfo()

This method is used to prepare a sensible information about your query when you are executing `var_dump($expr)`. The output will be HTML-safe.

Atk4\Data\Persistence\Sql\Expression::getDebugQuery()

Outputs query as a string by placing parameters into their respective places. The parameters will be escaped, but you should still avoid using generated query as it can potentially make you vulnerable to SQL injection.

This method will use HTML formatting if argument is passed.

In order for HTML parsing to work and to make your debug queries better formatted, install `doctrine/sql-formatter`:

```
composer require doctrine/sql-formatter
```

19.4.8 Escaping Methods

The following methods are useful if you're building your own code for rendering parts of the query. You must not call them in normal circumstances.

Atk4\Data\Persistence\Sql\Expression::consume(\$expr, string \$escapeMode)

Makes `$expr` part of `$this` expression. Argument may be either a string (which will be escaped) or another *Expression* or *Query*. If specified *Query* is in "select" mode, then it's automatically placed inside parentheses:

```
$query->consume('first_name', Expression::ESCAPE_PARAM); // :a
$query->consume('first_name', Expression::ESCAPE_IDENTIFIER); // `first_name`
$query->consume($otherQuery, Expression::ESCAPE_PARAM); // will merge parameters,
↳and return string
```

Atk4\Data\Persistence\Sql\Expression::escapeIdentifier(string \$value)

Always surrounds `$value` with back-ticks.

This escaping method is automatically used for `{...}` expression template tags.

`Atk4\Data\Persistence\Sql\Expression::escapeIdentifierSoft(string $value)`

Surrounds `$value` with back-ticks.

This escaping method is automatically used for `{{...}}` expression template tags .

It will smartly escape `table.field` type of strings resulting in `table.field`.

Will do nothing if it finds `"*`, `"` or `"("` character in `$value``:

```
$query->escapeIdentifierSoft('first_name'); // `first_name`
$query->escapeIdentifierSoft('first.name'); // `first`.`name`
$query->escapeIdentifierSoft('(2 + 2)'); // (2 + 2)
$query->escapeIdentifierSoft('*'); // *
```

`Atk4\Data\Persistence\Sql\Expression::escapeParam($value)`

Converts value into parameter and returns reference. Used only during query rendering. Consider using `consume()` instead, which will also handle nested expressions properly.

This escaping method is automatically used for `[...]` expression template tags .

19.4.9 Other Properties

property `Atk4\Data\Persistence\Sql\Expression::$template`

Template which is used when rendering. You can set this with either `$connection->expr('show tables')` or `$connection->expr(['show tables'])` or `$connection->expr(['template' => 'show tables'])`.

property `Atk4\Data\Persistence\Sql\Expression::$connection`

DB connection object.

property `Atk4\Data\Persistence\Sql\Expression::$paramBase`

Normally parameters are named `:a`, `:b`, `:c`. You can specify a different param base such as `:param_00` and it will be automatically increased into `:param_01` etc.

property `Atk4\Data\Persistence\Sql\Expression::$debug`

If true, then next call of `execute` will echo results of `getDebugQuery`.

class `Atk4\Data\Persistence\Sql\Query`

19.5 Queries

Query class represents your SQL query in-the-making. Once you create object of the Query class, call some of the methods listed below to modify your query. To actually execute your query and start retrieving data, see `fetching-result` section.

You should use [Connection](#) if possible to create your query objects. All examples below are using `$c->dsql()` method which generates Query linked to your established database connection.

Once you have a query object you can execute modifier methods such as `field()` or `table()` which will change the way how your query will act.

Once the query is defined, you can either use it inside another query or expression or you can execute it in exchange for result set.

Quick Example:

```
$query = $c->dsql();

$query->field('name');
$query->where('id', 123);

$name = $query->getOne();
```

19.5.1 Method invocation principles

Methods of Query are designed to be flexible and concise. Most methods have a variable number of arguments and some arguments can be skipped:

```
$query->where('id', 123);
$query->where('id', '=', 123); // the same
```

Most methods will accept *Expression* or strings. Strings are escaped or quoted (depending on type of argument). By using *Expression* you can bypass the escaping.

There are 2 types of escaping:

- *Expression::escapeIdentifier()*. Used for field and table names. Surrounds name with `.
- *Expression::escapeParam()*. Will convert value into parameter and replace with :a

In the next example \$a is escaped but \$b is parameterized:

```
$query->where('a', 'b');

// where `a` = "b"
```

If you want to switch places and execute where "b" = 'a', then you can resort to Expressions:

```
$query->where($c->expr('{} = []', ['b', 'a']));
```

Parameters which you specify into Expression will be preserved and linked into the \$query properly.

19.5.2 Query Modes

When you create new Query it always start in “select” mode. You can switch query to a different mode using mode. If you don’t switch the mode, your Query remains in select mode and you can fetch results from it anytime.

The pattern of defining arguments for your Query and then executing allow you to re-use your query efficiently:

```
$data = ['name' => 'John', 'surname' => 'Smith']

$query = $c->dsql();
$query
  ->where('id', 123)
  ->field('id')
  ->table('user')
  ->set($data);

$row = $query->getRow();
```

(continues on next page)

(continued from previous page)

```
if ($row) {
    $query
    ->set('revision', $query->expr('revision + 1'))
    ->mode('update')
    ->executeStatement();
} else {
    $query
    ->set('revision', 1)
    ->mode('insert')
    ->executeStatement();
}
```

The example above will perform a select query first:

- select id from user where id = 123

If a single row can be retrieved, then the update will be performed:

- update user set name = "John", surname = "Smith", revision = revision + 1 where id = 123

Otherwise an insert operation will be performed:

- insert into user (name, surname, revision) values ("John", "Smith", 1)

19.5.3 Chaining

Majority of methods return `$this` when called, which makes it pretty convenient for you to chain calls by using `->fx()` multiple times as illustrated in last example.

You can also combine creation of the object with method chaining:

```
$age = $c->dsql()->table('user')->where('id', 123)->field('age')->getOne();
```

19.5.4 Using query as expression

You can use query as expression where applicable. The query will get a special treatment where it will be surrounded in parentheses. Here are few examples:

```
$q = $c->dsql()
    ->table('employee');

$q2 = $c->dsql()
    ->field('name')
    ->table($q);

$q->getRows();
```

This query will perform select name from (select * from employee):

```
$q1 = $c->dsql()
    ->table('sales')
    ->field('date')
```

(continues on next page)

(continued from previous page)

```

->field('amount', null, 'debit');

$q2 = $c->dsql()
->table('purchases')
->field('date')
->field('amount', null, 'credit');

$u = $c->dsql('[] union []', [$q1, $q2]);

$q = $c->dsql()
->field('date, debit, credit')
->table($u, 'derivedTable');

$q->getRows();

```

This query will perform union between 2 table selects resulting in the following query:

```

select `date`, `debit`, `credit` from (
  (select `date`, `amount` `debit` from `sales`) union
  (select `date`, `amount` `credit` from `purchases`)
) `derivedTable`

```

19.5.5 Modifying Select Query

Setting Table

Atk4\Data\Persistence\Sql\Query::table(\$table, \$alias)

Specify a table to be used in a query.

Parameters

- **\$table** (mixed) – table such as “employees”
- **\$alias** (mixed) – alias of table

Returns

This method can be invoked using different combinations of arguments. Follow the principle of specifying the table first, and then optionally provide an alias. You can specify multiple tables at the same time by using comma or array (although you won't be able to use the alias there). Using keys in your array will also specify the aliases.

Basic Examples:

```

$c->dsql()->table('user');
// SELECT * from `user`

$c->dsql()->table('user', 'u');
// aliases table with "u"
// SELECT * from `user` `u`

$c->dsql()->table('user')->table('salary');
// specify multiple tables. Don't forget to link them by using "where"
// SELECT * from `user`, `salary`

```

(continues on next page)

(continued from previous page)

```

$c->dsql()->table(['user', 'salary']);
// identical to previous example
// SELECT * from `user`, `salary`

$c->dsql()->table(['u' => 'user', 's' => 'salary']);
// specify aliases for multiple tables
// SELECT * from `user` `u`, `salary` `s`

```

Inside your query table names and aliases will always be surrounded by backticks. If you want to use a more complex expression, use *Expression* as table:

```

$c->dsql()->table(
  $c->expr('(SELECT id FROM user UNION select id from document)'),
  'tbl'
);
// SELECT * FROM (SELECT id FROM user UNION SELECT id FROM document) `tbl`

```

Finally, you can also specify a different query instead of table, by simply passing another *Query* object:

```

$subQuery = $c->dsql();
$subQuery->table('employee');
$subQuery->where('name', 'John');

$q = $c->dsql();
$q->field('surname');
$q->table($subQuery, 'sub');

// SELECT `surname` FROM (SELECT * FROM `employee` WHERE `name` = :a) `sub`

```

Method can be executed several times on the same Query object.

Setting Fields

Atk4\Data\Persistence\Sql\Query::field(\$fields, \$alias = null)

Adds additional field that you would like to query. If never called, will default to *Query::\$defaultField*, which normally is ***.

This method has several call options. *\$field* can be array of fields and also can be an *Expression* or *Query*

Parameters

- **\$fields** (string|array|object) – Specify list of fields to fetch
- **\$alias** (string) – Optionally specify alias of field in resulting query

Returns \$this

Basic Examples:

```

$query = new Query();
$query->table('user');

$query->field('first_name');
// SELECT `first_name` from `user`

```

(continues on next page)

(continued from previous page)

```

$query->field('first_name, last_name');
// SELECT `first_name`, `last_name` from `user`

$query->field('employee.first_name')
// SELECT `employee`.`first_name` from `user`

$query->field('first_name', 'name')
// SELECT `first_name` `name` from `user`

$query->field(['name' => 'first_name'])
// SELECT `first_name` `name` from `user`

$query->field(['name' => 'employee.first_name']);
// SELECT `employee`.`first_name` `name` from `user`

```

If the first parameter of field() method contains non-alphanumeric character such as space or parenthesis, then field() will assume that you're passing an expression:

```

$query->field('now()');

$query->field('now()', 'time_now');

```

You may also pass array as first argument. In such case array keys will be used as aliases (if they are specified):

```

$query->field(['time_now' => 'now()', 'time_created']);
// SELECT now() `time_now`, `time_created` ...

$query->field($query->dsql()->table('user')->field('max(age)'), 'max_age');
// SELECT (SELECT max(age) from user) `max_age` ...

```

Method can be executed several times on the same Query object.

Setting where and having clauses

Atk4\Data\Persistence\Sql\Query::where(\$field, \$operation, \$value)

Adds WHERE condition to your query.

Parameters

- **\$field** (mixed) – field such as “name”
- **\$operation** (mixed) – comparison operation such as “>” (optional)
- **\$value** (mixed) – value or expression

Returns \$this

Atk4\Data\Persistence\Sql\Query::having(\$field, \$operation, \$value)

Adds HAVING condition to your query.

Parameters

- **\$field** (mixed) – field such as “name”
- **\$operation** (mixed) – comparison operation such as “>” (optional)

- **\$value** (mixed) – value or expression

Returns \$this

Both methods use identical call interface. They support one, two or three argument calls.

Pass string (field name), *Expression* or even *Query* as first argument.

Operator can be specified through a second parameter - \$operation. If unspecified, will default to '='.

Last argument is value. You can specify number, string, array, expression or even null (specifying null is not the same as omitting this argument). This argument will always be parameterized unless you pass expression. If you specify array, all elements will be parametrized individually.

Starting with the basic examples:

```
$q->where('id', 1);
$q->where('id', '=', 1); // same as above

$q->where('id', '>', 1);

$q->where('id', '=', null); // will render to "IS NULL" SQL
$q->where('id', null) // same as above

$q->where('now()', 1); // will not use backticks
$q->where($c->expr('now()'), 1); // same as above

$q->where('id', [1, 2]); // renders as id in (1, 2)
```

You may call where() multiple times, and conditions are always additive (uses AND). The easiest way to supply OR condition is to specify multiple conditions through array:

```
$q->where([[ 'name', 'like', '%john%' ], [ 'surname', 'like', '%john%' ]]);
// .. WHERE `name` like %john% OR `surname` like %john%
```

You can also mix and match with expressions and strings:

```
$q->where([[ 'name', 'like', '%john%' ], 'surname is null']);
// .. WHERE `name` like %john% AND `surname` is null

$q->where([[ 'name', 'like', '%john%' ], $q->expr('surname is null')]);
// .. WHERE `name` like %john% AND surname is null
```

There is a more flexible way to use OR arguments:

Atk4\Data\Persistence\Sql\Query::orExpr()

Returns new Query object with method “where()”. When rendered all clauses are joined with “OR”.

Atk4\Data\Persistence\Sql\Query::andExpr()

Returns new Query object with method “where()”. When rendered all clauses are joined with “OR”.

Here is a sophisticated example:

```
$q = $c->dsql();

$q->table('employee')->field('name');
$q->where('deleted', 0);
$q->where(
```

(continues on next page)

(continued from previous page)

```

$q
  ->orExpr()
  ->where('a', 1)
  ->where('b', 1)
  ->where(
    $q->andExpr()
      ->where('a', 2)
      ->where('b', 2)
  )
);

```

The above code will result in the following query:

```

select
  `name`
from
  `employee`
where
  deleted = 0 and
  (`a` = :a or `b` = :b or (`a` = :c and `b` = :d))

```

Technically orExpr() generates a yet another object that is composed and renders its calls to where() method:

```

$q->having(
  $q
  ->orExpr()
  ->where('a', 1)
  ->where('b', 1)
);

```

```

having
  (`a` = :a or `b` = :b)

```

Grouping results by field

Atk4\Data\Persistence\Sql\Query::group(\$field)

Group by functionality. Simply pass either field name as string or *Expression* object.

Parameters

- **\$field** (mixed) – field such as “name”

Returns

\$this

The “group by” clause in SQL query accepts one or several fields. It can also accept expressions. You can call group() with one or several comma-separated fields as a parameter or you can specify them in array. Additionally you can mix that with *Expression* or *Expressionable* objects.

Few examples:

```

$q->group('gender');

$q->group('gender, age');

```

(continues on next page)

(continued from previous page)

```
$q->group(['gender', 'age']);
$q->group('gender')->group('age');
$q->group($q->expr('year(date)'));
```

Method can be executed several times on the same Query object.

Concatenate group of values

Atk4\Data\Persistence\Sql\Query::groupConcat(\$field, \$separator = ',')

Quite often when you use `group by` in your queries you also would like to concatenate group of values.

Parameters

- **\$field** (mixed) – Field name or object
- **\$separator** (string) – Optional separator to use. It's comma by default

Different SQL engines have different syntax for doing this. In MySQL it's `group_concat()`, in Oracle it's `listagg`, but in PostgreSQL it's `string_agg`. That's why we have this method which will take care of this.

```
$q->groupConcat('phone', ';');
// group_concat('phone', ';')
```

If you need to add more parameters for this method, then you can extend this class and overwrite this simple method to support expressions like this, for example:

```
group_concat('phone' order by 'date' desc separator ';')
```

Joining with other tables

Atk4\Data\Persistence\Sql\Query::join(\$foreignTable, \$masterField, \$joinKind)

Join results with additional table using “JOIN” statement in your query.

Parameters

- **\$foreignTable** (string|array) – table to join (may include field and alias)
- **\$masterField** (mixed) – main field (and table) to join on or Expression
- **\$joinKind** (string) – ‘left’ (default), ‘inner’, ‘right’ etc - which join type to use

Returns

\$this

When joining with a different table, the results will be stacked by the SQL server so that fields from both tables are available. The first argument can specify the table to join, but may contain more information:

```
$q->join('address'); // address.id = address_id
// JOIN `address` ON `address`.`id`=`address_id`

$q->join('address a'); // specifies alias for the table
// JOIN `address` `a` ON `address`.`id`=`address_id`
```

(continues on next page)

(continued from previous page)

```
$q->join('address.user_id'); // address.user_id = id
// JOIN `address` ON `address`.`user_id`='id`
```

The second argument to join specifies which existing table/field is used in on condition:

```
$q->table('user u');
$q->join('user boss', 'u.boss_user_id');
// JOIN `user` `boss` ON `boss`.`id`='u`.`boss_user_id`
```

By default the “on” field is defined as `$table . "_id"`, as you have seen in the previous examples where join was done on “address_id”, and “credit_card_id”. If you have specified field explicitly in the foreign field, then the “on” field is set to “id”, like in the example above.

You can specify both fields like this:

```
$q->table('employees');
$q->join('salaries.emp_no', 'emp_no');
```

If you only specify field like this, then it will be automatically prefixed with the name or alias of your main table. If you have specified multiple tables, this won’t work and you’ll have to define name of the table explicitly:

```
$q->table('user u');
$q->join('user boss', 'u.boss_user_id');
$q->join('user super_boss', 'boss.boss_user_id');
```

The third argument specifies type of join and defaults to “left” join. You can specify “inner”, “straight” or any other join type that your database support.

Method can be executed several times on the same Query object.

Joining on expression

For a more complex join conditions, you can pass second argument as expression:

```
$q->table('user', 'u');
$q->join('address a', $q->expr('a.name like u.pattern'));
```

Use WITH cursors

Atk4\Data\Persistence\Sql\Query::with(*Query \$cursor*, *string \$alias*, *?array \$fields = null*, *bool \$recursive = false*)

If you want to add WITH cursor statement in your SQL, then use this method. First parameter defines sub-query to use. Second parameter defines alias of this cursor. By using third, optional argument you can set aliases for columns in cursor. And finally forth, optional argument set if cursors will be recursive or not.

You can add more than one cursor in your query.

Did you know: you can use these cursors when joining your query to other tables. Just join cursor instead.

Keep in mind that if any of cursors added in your query will be recursive, then all cursors will be set recursive. That’s how SQL want it to be.

Example:

```

$quotes = $q->table('quotes')
->field('emp_id')
->field($q->expr('sum([])', ['total_net']))
->group('emp_id');
$invoices = $q()->table('invoices')
->field('emp_id')
->field($q->expr('sum([])', ['total_net']))
->group('emp_id');
$employees = $q
->with($quotes, 'q', ['emp', 'quoted'])
->with($invoices, 'i', ['emp', 'invoiced'])
->table('employees')
->join('q.emp')
->join('i.emp')
->field(['name', 'salary', 'q.quoted', 'i.invoiced']);

```

This generates SQL below:

```

with
  `q` (`emp`, `quoted`) as (select `emp_id`, sum(`total_net`) from `quotes` group by
↳ `emp_id`),
  `i` (`emp`, `invoiced`) as (select `emp_id`, sum(`total_net`) from `invoices` group
↳ by `emp_id`)
select `name`, `salary`, `q`.`quoted`, `i`.`invoiced`
from `employees`
  left join `q` on `q`.`emp` = `employees`.`id`
  left join `i` on `i`.`emp` = `employees`.`id`

```

Limiting result-set

Atk4\Data\Persistence\Sql\Query::limit(\$cnt, \$shift)

Limit how many rows will be returned.

Parameters

- **\$cnt** (int) – number of rows to return
- **\$shift** (int) – offset, how many rows to skip

Returns \$this

Use this to limit your *Query* result-set:

```

$q->limit(5, 10);
// .. LIMIT 10, 5

$q->limit(5);
// .. LIMIT 0, 5

```

Ordering result-set

Atk4\Data\Persistence\Sql\Query::order(\$order, \$desc)

Orders query result-set in ascending or descending order by single or multiple fields.

Parameters

- **\$order** (string) – one or more field names, expression etc.
- **\$desc** (int) – pass true to sort descending

Returns \$this

Use this to order your *Query* result-set:

```
$q->order('name'); // .. order by name
$q->order('name desc'); // .. order by name desc
$q->order(['name desc', 'id asc']) // .. order by name desc, id asc
$q->order('name', true); // .. order by name desc
```

Method can be executed several times on the same Query object.

19.5.6 Insert and Replace query

Set value to a field

Atk4\Data\Persistence\Sql\Query::set(\$field, \$value)

Assigns value to the field during insert.

Parameters

- **\$field** (string) – name of the field
- **\$value** (mixed) – value or expression

Returns \$this

Example:

```
$q->table('user')->set('name', 'john')->mode('insert')->executeStatement();
// insert into user (name) values (john)

$q->table('log')->set('date', $q->expr('now()'))->mode('insert')->executeStatement();
// insert into log (date) values (now())
```

Method can be executed several times on the same Query object.

Set Insert Options

Atk4\Data\Persistence\Sql\Query::option(\$option, \$mode = 'select')

It is possible to add arbitrary options for the query. For example this will fetch unique user birthdays:

```
$q->table('user');
$q->option('distinct');
$q->field('birthday');
$birthdays = $q->getRows();
```

Other possibility is to set options for delete or insert:

```
$q->option('delayed', 'insert');  
  
// or  
  
$q->option('ignore', 'insert');
```

See your SQL capabilities for additional options (low_priority, delayed, high_priority, ignore)

19.5.7 Update Query

Set Conditions

Same syntax as for Select Query.

Set value to a field

Same syntax as for Insert Query.

Other settings

Limit and Order are normally not included to avoid side-effects, but you can modify `Query::$templateUpdate` to include those tags.

19.5.8 Delete Query

Set Conditions

Same syntax as for Select Query.

Other settings

Limit and Order are normally not included to avoid side-effects, but you can modify `Query::$templateUpdate` to include those tags.

19.5.9 Dropping attributes

If you have called `where()` several times, there is a way to remove all the where clauses from the query and start from beginning:

```
Atk4\Data\Persistence\Sql\Query::reset($tag)
```

Parameters

- **\$tag** (string) – part of the query to delete/reset.

Example:

```
$q
->table('user')
->where('name', 'John');
->reset('where')
->where('name', 'Peter');

// where name = 'Peter'
```

19.5.10 Other Methods

Atk4\Data\Persistence\Sql\Query::**dsql**(\$defaults)

Use this instead of new Query() if you want to automatically bind query to the same connection as the parent.

Atk4\Data\Persistence\Sql\Query::**expr**(\$template, \$arguments)

Method very similar to *Connection::expr* but will return a corresponding Expression class for this query.

Atk4\Data\Persistence\Sql\Query::**exprNow**(\$precision)

Method will return current_timestamp(precision) sub-query.

Atk4\Data\Persistence\Sql\Query::**option**(\$option, \$mode)

Use this to set additional options for particular query mode. For example:

```
$q
->table('test')
->field('name')
->set('name', 'John')
->option('calc_found_rows') // for default select mode
->option('ignore', 'insert') // for insert mode;

$q->executeQuery(); // select calc_found_rows `name` from `test`
$q->mode('insert')->executeStatement(); // insert ignore into `test` (`name`) values
->(`name` = 'John')
```

Atk4\Data\Persistence\Sql\Query::_**setArgs**(\$what, \$alias, \$value)

Internal method which sets value in Expression::\$args array. It doesn't allow duplicate aliases and throws Exception in such case. Argument \$what can be 'table' or 'field'.

Atk4\Data\Persistence\Sql\Query::**caseExpr**(\$operand)

Returns new Query object with CASE template. You can pass operand as parameter to create SQL like

```
CASE <operand> WHEN <expression> THEN <expression> END type of SQL statement.
```

Atk4\Data\Persistence\Sql\Query::**caseWhen**(\$when, \$then)

Set WHEN condition and THEN expression for CASE statement.

Atk4\Data\Persistence\Sql\Query::**otherwise**(\$else)

Set ELSE expression for CASE statement.

Few examples:

```
$s = $this->q()->caseExpr()
->caseWhen(['status', 'New'], 't2.expose_new')
->caseWhen(['status', 'like', '%Used%'], 't2.expose_used')
->caseElse(null);
```

```
case when "status" = 'New' then "t2"."expose_new" when "status" like '%Used%' then
↳ "t2"."expose_used" else null end
```

```
$s = $this->q()->caseExpr('status')
->caseWhen('New', 't2.expose_new')
->caseWhen('Used', 't2.expose_used')
->caseElse(null);
```

```
case "status" when 'New' then "t2"."expose_new" when 'Used' then "t2"."expose_used"↳
↳ else null end
```

19.5.11 Properties

property Atk4\Data\Persistence\Sql\Query::\$mode

Query will use one of the predefined “templates”. The mode will contain name of template used. Basically it’s array key of \$templates property. See Query Modes.

property Atk4\Data\Persistence\Sql\Query::\$defaultField

If no fields are defined, this field is used.

property Atk4\Data\Persistence\Sql\Query::\$templateSelect

Template for SELECT query. See Query Modes.

property Atk4\Data\Persistence\Sql\Query::\$templateInsert

Template for INSERT query. See Query Modes.

property Atk4\Data\Persistence\Sql\Query::\$templateReplace

Template for REPLACE query. See Query Modes.

property Atk4\Data\Persistence\Sql\Query::\$templateUpdate

Template for UPDATE query. See Query Modes.

property Atk4\Data\Persistence\Sql\Query::\$templateDelete

Template for DELETE query. See Query Modes.

property Atk4\Data\Persistence\Sql\Query::\$templateTruncate

Template for TRUNCATE query. See Query Modes.

19.6 Results

When query is executed by [Connection](#) or [PDO](#), it will return an object that can stream results back to you. The PDO class execution produces a [Doctrine\DBAL\Result](#) object which you can iterate over.

If you are using a custom connection, you then will also need a custom object for streaming results.

The only requirement for such an object is that it has to be a [Generator](#). In most cases developers will expect your generator to return sequence of id => hash representing a key/value result set.

19.7 Transactions

When you work with the DSQL, you can work with transactions. There are 2 enhancements to the standard functionality of transactions in DSQL:

1. You can start nested transactions.
2. You can use `Connection::atomic()` which has a nicer syntax.

It is recommended to always use `atomic()` in your code.

class Atk4\Data\Persistence\Sql\Connection

Atk4\Data\Persistence\Sql\Connection::atomic(\$callback)

Execute callback within the SQL transaction. If callback encounters an exception, whole transaction will be automatically rolled back:

```
$c->atomic(function () use ($c) {
    $c->dsql('user')->set('balance = balance + 10')->where('id', 10)->mode('update
    ↪')->executeStatement();
    $c->dsql('user')->set('balance = balance - 10')->where('id', 14)->mode('update
    ↪')->executeStatement();
});
```

`atomic()` can be nested. The successful completion of a top-most method will commit everything. Rollback of a top-most method will roll back everything.

Atk4\Data\Persistence\Sql\Connection::beginTransaction()

Start new transaction. If already started, will do nothing but will increase transaction depth.

Atk4\Data\Persistence\Sql\Connection::commit()

Will commit transaction, however if `Connection::beginTransaction` was executed more than once, will only decrease transaction depth.

Atk4\Data\Persistence\Sql\Connection::inTransaction()

Returns true if transaction is currently active. There is no need for you to ever use this method.

Atk4\Data\Persistence\Sql\Connection::rollback()

Roll-back the transaction, however if `Connection::beginTransaction` was executed more than once, will only decrease transaction depth.

Warning: If you roll-back internal transaction and commit external transaction, then result might be unpredictable. Please discuss this <https://github.com/atk4/dsql/issues/89>

19.8 Advanced Topics

DSQL has huge capabilities in terms of extending. This chapter explains just some of the ways how you can extend this already incredibly powerful library.

19.8.1 Advanced Connections

Connection is incredibly lightweight and powerful in DSQL. The class tries to get out of your way as much as possible.

Using DSQL without Connection

You can use *Query* and *Expression* without connection at all. Simply create expression:

```
$expr = new Mysql\Expression('show tables like []', ['foo%']);
```

or query:

```
$query = (new Mysql\Query())->table('user')->where('id', 1);
```

When it's time to execute you can specify your Connection manually:

```
$rows = $expr->getRows($connection);
foreach ($rows as $row) {
    echo json_encode($row) . "\n";
}
```

With queries you might need to select mode first:

```
$stmt = $query->mode('delete')->executeStatement($connection);
```

The `Expression::execute` is a convenient way to prepare query, bind all parameters and get `Doctrine\DBAL\Result`, but if you wish to do it manually, see [Manual Query Execution](#).

Using in Existing Framework

If you use DSQL inside another framework, it's possible that there is already a PDO object which you can use. In Laravel you can optimize some of your queries by switching to DSQL:

```
$c = new Connection(['connection' => $pdo]);

$userIds = $c->dsql()->table('expired_users')->field('user_id');
$c->dsql()->table('user')->where('id', 'in', $userIds)->set('active', 0)->mode('update')->executeStatement();

// native Laravel Database Query Builder
// $userIds = DB::table('expired_users')->lists('user_id');
// DB::table('user')->whereIn('id', $userIds)->update(['active', 0]);
```

The native query builder in the example above populates `$userIds` with array from `expired_users` table, then creates second query, which is an update. With DSQL we have accomplished same thing with a single query and without fetching results too.

```
UPDATE
  user
SET
  active = 0
WHERE
  id in (SELECT user_id from expired_users)
```

If you are creating `Connection` through constructor, you may have to explicitly specify property `Connection::$queryClass`:

```
$c = new Connection(['connection' => $pdo, 'queryClass' => Atk4\Data\Persistence\Sql\
↳ SQLite\Query::class]);
```

This is also useful, if you have created your own Query class in a different namespace and wish to use it.

19.8.2 Extending Query Class

You can add support for new database vendors by creating your own `Query` class. Let's say you want to add support for new SQL vendor:

```
class Query_MyVendor extends Atk4\Data\Persistence\Sql\Query
{
    protected string $identifierEscapeChar = '';
    protected string $expressionClass = Expression_MyVendor::class;

    // truncate is done differently by this vendor
    protected string $templateTruncate = 'delete [from] [table]';

    // also join is not supported
    public function join(
        $foreignTable,
        $masterField = null,
        $joinKind = null,
        $foreignAlias = null
    ) {
        throw new Atk4\Data\Persistence\Sql\Exception('Join is not supported by the_
↳ database');
    }
}
```

Now that our custom query class is complete, we would like to use it by default on the connection:

```
$c = \Atk4\Data\Persistence\Sql\Connection::connect($dsn, $user, $pass, ['queryClass' =>
↳ 'Query_MyVendor']);
```

Adding new vendor support through extension

If you think that more people can benefit from your custom query class, you can create a separate add-on with its own namespace. Let's say you have created `myname/dsql-myvendor`.

1. Create your own Query class inside your library. If necessary create your own Connection class too.
2. Make use of composer and add dependency to DSQL.
3. Add a nice README file explaining all the quirks or extensions. Provide install instructions.
4. Fork DSQL library.
5. Modify `Connection::connect` to recognize your database identifier and refer to your namespace.
6. Modify docs/extensions.md to list name of your database and link to your repository / composer requirement.

7. Copy `phpunit-mysql.xml` into `phpunit-myvendor.xml` and make sure that `dsql/tests/db/*` works with your database.
8. Submit pull request for only the `Connection` class and `docs/extensions.md`.

If you would like that your vendor support be bundled with DSQL, you should contact copyright@agiletoolkit.org after your external class has been around and received some traction.

Adding New Query Modes

By Default DSQL comes with the following *Query Modes*:

- `select`
- `delete`
- `insert`
- `replace`
- `update`
- `truncate`

You can add new mode if you wish. Let's look at how to add a MySQL specific query "LOAD DATA INFILE":

1. Define new property inside your *Query* class `$templateLoadData`.
2. Add public method allowing to specify necessary parameters.
3. Re-use existing methods/template tags if you can.
4. Create `_render` method if your tag rendering is complex.

So to implement our task, you might need a class like this:

```
use \Atk4\Data\Persistence\Sql\Exception;

class QueryMysqlCustom extends \Atk4\Data\Persistence\Sql\Mysql\Query
{
    protected string $templateLoadData = 'load data local infile [file] into table_
→[table]';

    public function file($file)
    {
        if (!is_readable($file)) {
            throw Exception(['File is not readable', 'file' => $file]);
        }
        $this['file'] = $file;
    }

    public function loadData(): array
    {
        return $this->mode('loadData')->getRows();
    }
}
```

Then to use your new statement, you can do:

```
$c->dsql()->file('abc.csv')->loadData();
```

19.8.3 Manual Query Execution

If you are not satisfied with `Expression::execute` you can execute query yourself.

1. `Expression::render` query, then send the 1st element into `PDO::prepare()`;
2. use new `$statement` to bind `bindValue` with the contents of 2nd element;
3. set result fetch mode and parameters;
4. `execute()` your statement

19.8.4 Exception Class

DSQL slightly extends and improves `Exception` class

```
class Atk4\Data\Persistence\Sql\Exception
```

The main goal of the new exception is to be able to accept additional information in addition to the message. We realize that often `$e->getMessage()` will be localized, but if you stick some variables in there, this will no longer be possible. You also risk injection or expose some sensitive data to the user.

```
Atk4\Data\Persistence\Sql\Exception::__construct($message, $code)
```

Create new exception

Parameters

- **\$message** (string|array) – Describes the problem
- **\$code** (int) – Error code

Usage:

```
throw new Atk4\Data\Persistence\Sql\Exception('Hello');

throw (new Atk4\Data\Persistence\Sql\Exception('File is not readable'))
->addMoreInfo('file', $file);
```

When displayed to the user the exception will hide parameter for `$file`, but you still can get it if you really need it:

```
Atk4\Data\Persistence\Sql\Exception::getParams()
```

Return additional parameters, that might be helpful to find error.

Returns array

Any DSQL-related code must always throw `Atk4\Data\Persistence\Sql\Exception`. Query-related errors will generate PDO exceptions. If you use a custom connection and doing some vendor-specific operations, you may also throw other vendor-specific exceptions.

19.9 Vendor support and Extensions

Vendor	Support	PDO	Dependency
MySQL	Full	mysql:	native, PDO
SQLite	Full	sqlite:	native, PDO
Oracle	Untested	oci:	native, PDO
PostgreSQL	Untested	pgsql:	native, PDO
MSSQL	Untested	mssql:	native, PDO

Note: Most PDO vendors should work out of the box

19.9.1 3rd party vendor support

Class	Support	PDO	Dependency
Connection_MyVendor	Full	myvendor:	https://github.com/test/myvendor

See *Adding new vendor support through extension* for more details on how to add support for your driver.

INDICES AND TABLES

- [genindex](#)
- [search](#)

PHP NAMESPACE INDEX

a

[Atk4\Data](#), 1

[Atk4\Data\Persistence\Sql](#), 160

Symbols

__construct() (*Atk4Data\Persistence\Sq\Exception* method), **165**
 __debugInfo() (*Atk4Data\Persistence\Sq\Expression* method), **145**
 _setArgs() (*Atk4Data\Persistence\Sq\Query* method), **159**

A

action() (*Atk4Data\Model* method), **59**
 actual (*Atk4Data\Field* property), **67, 77**
 addCalculatedField() (*Atk4Data\Model* method), **36**
 addCondition() (*Atk4Data\Model* method), **69**
 addCteModel() (*Atk4Data\Model* method), **126**
 addExpression() (*Atk4Data\Model* method), **35, 101**
 addField() (*Atk4Data\Model* method), **34**
 addFields() (*Atk4Data\Model* method), **35**
 addReference() (*Atk4Data\Model* method), **96**
 andExpr() (*Atk4Data\Persistence\Sq\Query* method), **152**
 Atk4\Data (namespace), **1, 8, 19, 30, 43, 47, 61, 64, 68, 75, 86, 88, 100, 104, 109, 111, 118, 131**
 Atk4\Data\Persistence\SqL (namespace), **135, 136, 139, 141, 146, 160, 161, 165**
 atomic() (*Atk4Data\Persistence\Sq\Connection* method), **161**

B

beginTransaction() (*Atk4Data\Persistence\Sq\Connection* method), **161**

C

caption (*Atk4Data\Model* property), **43**
 caseExpr() (*Atk4Data\Persistence\Sq\Query* method), **159**
 caseWhen() (*Atk4Data\Persistence\Sq\Query* method), **159**
 commit() (*Atk4Data\Persistence\Sq\Connection* method), **161**
 connect() (*Atk4Data\Persistence\Sq\Connection* method), **140**

connectDbalConnection() (*Atk4Data\Persistence\Sq\Connection* method), **141**
 connection (*Atk4Data\Persistence\Sq\Expression* property), **146**
 Connection (class in *Atk4Data\Persistence\SqL*), **140, 161**
 consume() (*Atk4Data\Persistence\Sq\Expression* method), **145**

D

data (*Atk4Data\Model* property), **40**
 debug (*Atk4Data\Persistence\Sq\Expression* property), **146**
 default (*Atk4Data\Field* property), **66**
 defaultField (*Atk4Data\Persistence\Sq\Query* property), **160**
 delete() (*Atk4Data\Model* method), **50**
 dirty (*Atk4Data\Model* property), **41**
 dsqL() (*Atk4Data\Persistence\Sq\Connection* method), **140**
 dsqL() (*Atk4Data\Persistence\Sq\Query* method), **159**
 duplicate() (*Atk4Data\Model* method), **54**

E

enum (*Atk4Data\Field* property), **66**
 escapeIdentifier() (*Atk4Data\Persistence\Sq\Expression* method), **145**
 escapeIdentifierSoft() (*Atk4Data\Persistence\Sq\Expression* method), **145**
 escapeParam() (*Atk4Data\Persistence\Sq\Expression* method), **146**
 Exception (class in *Atk4Data\Persistence\SqL*), **165**
 execute() (*Atk4Data\Persistence\Sq\Connection* method), **141**
 executeQuery() (*Atk4Data\Persistence\Sq\Expression* method), **144**
 export() (*Atk4Data\Model* method), **63**
 expr() (*Atk4Data\Model* method), **72**
 expr() (*Atk4Data\Persistence\Sq\Connection* method), **140**

- `expr()` (*Atk4Data\Persistence\SqLExpression method*), **144**
- `expr()` (*Atk4Data\Persistence\SqLQuery method*), **159**
- `Expression` (*class in Atk4Data\Persistence\SqL*), **141**
- `exprNow()` (*Atk4Data\Persistence\SqLQuery method*), **159**
- ## F
- `Field` (*class in Atk4Data*), **64, 65, 77**
- `field()` (*Atk4Data\Persistence\SqLQuery method*), **150**
- `Field\SqLExpressionField` (*class in Atk4Data*), **79**
- ## G
- `get()` (*Atk4Data\Field method*), **68**
- `get()` (*Atk4Data\Model method*), **40**
- `getDebugQuery()` (*Atk4Data\Persistence\SqLExpression method*), **145**
- `getDsqlExpression()` (*Atk4Data\Field method*), **77**
- `getField()` (*Atk4Data\Model method*), **41**
- `getIterator()` (*Atk4Data\Model method*), **63**
- `getModelCaption()` (*Atk4Data\Model method*), **43**
- `getOne()` (*Atk4Data\Persistence\SqLExpression method*), **145**
- `getParams()` (*Atk4Data\Persistence\SqLException method*), **165**
- `getRow()` (*Atk4Data\Persistence\SqLExpression method*), **144**
- `getRows()` (*Atk4Data\Persistence\SqLExpression method*), **144**
- `getTitle()` (*Atk4Data\Model method*), **42**
- `getTitles()` (*Atk4Data\Model method*), **42**
- `group()` (*Atk4Data\Persistence\SqLQuery method*), **153**
- `groupConcat()` (*Atk4Data\Persistence\SqLQuery method*), **154**
- ## H
- `hasField()` (*Atk4Data\Model method*), **41**
- `hasMany()` (*Atk4Data\Model method*), **91**
- `hasOne()` (*Atk4Data\Model method*), **94**
- `HasOnegetModel()` (*Atk4Data\ReferenceHasOne method*), **100**
- `HasOnelink` (*Atk4Data\ReferenceHasOne property*), **100**
- `HasOnemodel` (*Atk4Data\ReferenceHasOne property*), **100**
- `HasOneourField` (*Atk4Data\ReferenceHasOne property*), **100**
- `HasOneref()` (*Atk4Data\ReferenceHasOne method*), **100**
- `HasOneSqladdField()` (*Atk4Data\ReferenceHasOneSql method*), **78**
- `HasOneSqladdFields()` (*Atk4Data\ReferenceHasOneSql method*), **78**
- `HasOneSqladdTitle()` (*Atk4Data\ReferenceHasOneSql method*), **79**
- `HasOneSqlref()` (*Atk4Data\ReferenceHasOneSql method*), **78**
- `HasOneSqlrefLink()` (*Atk4Data\ReferenceHasOneSql method*), **78**
- `HasOnetableAlias` (*Atk4Data\ReferenceHasOne property*), **100**
- `HasOnetheir_filed` (*Atk4Data\ReferenceHasOne property*), **100**
- `having()` (*Atk4Data\Persistence\SqLQuery method*), **151**
- ## I
- `idField` (*Atk4Data\Model property*), **42**
- `import()` (*Atk4Data\Model method*), **39**
- `init()` (*Atk4Data\Model method*), **34**
- `insert()` (*Atk4Data\Model method*), **39**
- `inTransaction()` (*Atk4Data\Persistence\SqLConnection method*), **161**
- `isDirty()` (*Atk4Data\Model method*), **41**
- `isEditable()` (*Atk4Data\Field method*), **68**
- `isHidden()` (*Atk4Data\Field method*), **68**
- `isVisible()` (*Atk4Data\Field method*), **68**
- ## J
- `join` (*Atk4Data\Field property*), **67**
- `join()` (*Atk4Data\Persistence\SqLQuery method*), **154**
- `JoinaddField()` (*Atk4Data\ModelJoin method*), **107**
- `JoincontainsMany()` (*Atk4Data\ModelJoin method*), **107**
- `JoincontainsOne()` (*Atk4Data\ModelJoin method*), **107**
- `JoinhasMany()` (*Atk4Data\ModelJoin method*), **107**
- `JoinhasOne()` (*Atk4Data\ModelJoin method*), **107**
- `Joinjoin()` (*Atk4Data\ModelJoin method*), **107**
- ## L
- `limit()` (*Atk4Data\Persistence\SqLQuery method*), **156**
- `load()` (*Atk4Data\Model method*), **49**
- ## M
- `mode` (*Atk4Data\Persistence\SqLQuery property*), **160**
- `Model` (*class in Atk4Data*), **31, 49, 63, 69, 89, 101, 126**
- `Model\AggregateModel` (*class in Atk4Data*), **111**
- `Model\Join` (*class in Atk4Data*), **105**
- `Model\Scope` (*class in Atk4Data*), **73**
- `Model\Scope\Condition` (*class in Atk4Data*), **74**

N

neverPersist (*Atk4Data\Field* property), **67**
 neverSave (*Atk4Data\Field* property), **67**
 normalizeFieldName() (*Atk4Data\Model* method), **42**
 nullable (*Atk4Data\Field* property), **67**

O

onlyFields (*Atk4Data\Model* property), **40**
 option() (*Atk4Data\Persistence\SqLQuery* method),
157, 159
 order() (*Atk4Data\Persistence\SqLQuery* method), **157**
 orExpr() (*Atk4Data\Persistence\SqLQuery* method),
152
 otherwise() (*Atk4Data\Persistence\SqLQuery*
 method), **159**

P

paramBase (*Atk4Data\Persistence\SqLExpression* prop-
 erty), **146**
 persistence (*Atk4Data\Model* property), **39**
 Persistence\Csv (class in *Atk4Data*), **133**
 Persistence\SqL (class in *Atk4Data*), **79**
 Persistence\SqL\Join (class in *Atk4Data*), **108**
 Persistence\Static_ (class in *Atk4Data*), **87**
 persistenceData (*Atk4Data\Model* property), **39**

Q

Query (class in *Atk4Data\Persistence\SqL*), **146**

R

readOnly (*Atk4Data\Field* property), **67**
 ref() (*Atk4Data\Model* method), **89**
 Reference\HasOne (class in *Atk4Data*), **100**
 Reference\HasOneSql (class in *Atk4Data*), **78**
 refLink() (*Atk4Data\Model* method), **93**
 registerConnectionClass()
 (*Atk4Data\Persistence\SqLConnection*
 method), **141**
 render() (*Atk4Data\Persistence\SqLExpression*
 method), **144**
 required (*Atk4Data\Field* property), **67**
 reset() (*Atk4Data\Persistence\SqLQuery* method), **158**
 rollBack() (*Atk4Data\Persistence\SqLConnection*
 method), **161**

S

save() (*Atk4Data\Model* method), **49**
 scope() (*Atk4Data\Model* method), **73**
 Scope__construct() (*Atk4Data\ModelScope*
 method), **74**
 Scopeclear() (*Atk4Data\ModelScope* method), **74**
 ScopeCondition__construct()
 (*Atk4Data\ModelScopeCondition* method), **74**

ScopeConditionnegate()
 (*Atk4Data\ModelScopeCondition* method), **75**
 ScopeConditiontoWords()
 (*Atk4Data\ModelScopeCondition* method), **75**
 ScopecreateAnd() (*Atk4Data\ModelScope* method),
74
 ScopecreateOr() (*Atk4Data\ModelScope* method), **74**
 ScopeisAnd() (*Atk4Data\ModelScope* method), **74**
 ScopeisOr() (*Atk4Data\ModelScope* method), **74**
 Scopenegate() (*Atk4Data\ModelScope* method), **74**
 Scopesimplify() (*Atk4Data\ModelScope* method), **74**
 set() (*Atk4Data\Field* method), **68**
 set() (*Atk4Data\Model* method), **40**
 set() (*Atk4Data\Persistence\SqLQuery* method), **157**
 setLimit() (*Atk4Data\Model* method), **43**
 setMulti() (*Atk4Data\Model* method), **40**
 setNull() (*Atk4Data\Field* method), **68**
 setNull() (*Atk4Data\Model* method), **40**
 setOnlyFields() (*Atk4Data\Model* method), **40**
 setOrder() (*Atk4Data\Model* method), **43**
 Slatomic() (*Atk4Data\PersistenceSql* method), **79**
 Slexpr() (*Atk4Data\PersistenceSql* method), **80**
 SqlExpressionFieldexpr
 (*Atk4Data\FieldSqlExpressionField* prop-
 erty), **79**
 SqlExpressionFieldgetDsqlExpression()
 (*Atk4Data\FieldSqlExpressionField* method),
79
 Static__construct() (*Atk4Data\PersistenceStatic_*
 method), **87**
 system (*Atk4Data\Field* property), **67**

T

table (*Atk4Data\Model* property), **39**
 table() (*Atk4Data\Persistence\SqLQuery* method), **149**
 template (*Atk4Data\Persistence\SqLExpression* prop-
 erty), **146**
 templateDelete (*Atk4Data\Persistence\SqLQuery*
 property), **160**
 templateInsert (*Atk4Data\Persistence\SqLQuery*
 property), **160**
 templateReplace (*Atk4Data\Persistence\SqLQuery*
 property), **160**
 templateSelect (*Atk4Data\Persistence\SqLQuery*
 property), **160**
 templateTruncate (*Atk4Data\Persistence\SqLQuery*
 property), **160**
 templateUpdate (*Atk4Data\Persistence\SqLQuery*
 property), **160**
 titleField (*Atk4Data\Model* property), **42**
 tryLoad() (*Atk4Data\Model* method), **50**
 type (*Atk4Data\Field* property), **65**
 typecastLoadField() (*Atk4Data\Model* method), **54**
 typecastLoadRow() (*Atk4Data\Model* method), **54**

`typecastSaveField()` (*Atk4Data\Model* method), **54**
`typecastSaveRow()` (*Atk4Data\Model* method), **54**

U

`ui` (*Atk4Data\Field* property), **67**
`unload()` (*Atk4Data\Model* method), **50**
`unset()` (*Atk4Data\Model* method), **40**

V

`values` (*Atk4Data\Field* property), **66**

W

`where()` (*Atk4Data\Persistence\Sq\Query* method), **151**
`with()` (*Atk4Data\Persistence\Sq\Query* method), **155**
`withPersistence()` (*Atk4Data\Model* method), **39, 57**