# Core Features

## Basic Concepts and Components

The core feature is the basis of Odysseus. It holds all stuff that is needed by Odysseus in any case, e.g. that the query processing needs a transformation process etc. Since Odysseus is a framework, its components can be extended and configured via services and interfaces. However, the core also includes the most common implementations for the components, so that there is at least one complete system configuration. In the following, we list some basic framework concepts of Odysseus and say how a framework concept is implemented by a dedicated technology/concept

### Component-based Architecture

Odysseus is based on OSGi so that it is a component-based architecture. Most of the functionalities in Odysseus can be extended or configured via services through the components. This also allows the adaption of the system during runtime. Since it is also possible to substitute components, default concepts can be replaced by other/own concepts.

### Independent Processing Objects

Most systems only have one processing type (e.g. relational, XML or just strings). Odysseus can handle arbitrary processing types. The default processing type is a "relational".

### Extendable Metadata (TimeIntevals and Latency...)

Each object that is processed by the system can be enriched with metadata. Besides fixed metadata, each object can be optionally annotated with key-value pairs. The fixed metadata is not optional, because it is for example used for the processing. Therefore, the default metadata is "TimeInterval" (or also referred to as "interval" or "interval approach"). The TimeInterval metadata provides two timestamps that indicate the start and the end of the validity of the processing object. Each operator that recognizes the time interval metadata uses this metadata to process only those processing objects that are valid at the same time. This allows the processing of a potentially unbounded data stream by windowing the stream through time intervals. Furthermore, there is the possibility to use a latency metadata - which is used for measuring the latency of a processing object.

### Arbitrary Schema and Data types

The processing objects can be specified by a schema and data type that is extendable and is called simple description framework (SDF). In the relational processing, for example, the schema describes the names (attributes) and the data types of the tuple (which is the processing object in the relational world). The schema can be seen as a list of attributes and each attribute has a name and a datatype. Although there are data types for integer, float or something else, it is also possible to introduce your own datatype. This could be everything and is only a marker that could be used by operators, but neither the schema nor the data types are normally used by the processing (Needless to say, that some relational operators during the relational processing, e.g. a projection, need the schema, so that it have to fit to the processing object).

### Transformation: Logical and Physical Operators

Odysseus distinguishes between logical operators that only say "what" this operator does but does not say "how". Thus, the logical layer is independent of any implementation and normally also from any processing object types (see above). The transformation converts the logical into a physical representation. This physical counterpart provides the real implementation. Thus, it is possible to have more than one implementation for one logical representation. A set of rules and a rule engine manages how a logical operator is transformed and which physical operator is used. Since these rules can be complemented and overloaded, it is also possible to announce own rules (e.g. to transform the logical operators into physical operators for a new processing type)

### Rewrite: Optimization of Logical Operators

A given logical plan (which is a graph based on logical operators) can be optimized via restructuring using a set of rewrite rules. For the relational processing for example, there is a rule that pushes a selection down to the source, so that the number of processing objects is reduced as early as possible. These rules can also be overloaded and completed by new ones.

### Creating New Operators

It is possible to create your own new operators. For this, the only thing is to create a logical and a physical operator and an appropriate transformation rules that says how this logical operator should be transformed into the physical operator.

### Build-in Math Expression Parser (MEP)

A math expression parser (MEP) provides the possibility to evaluation expressions like "(a + 5)*93/12". MEP can be extended by new functions (like round, floor, abs...)

### Several Scheduling Strategies

Odysseus provides several scheduling techniques which can be extended. The default scheduler is a single thread scheduler who schedules all operators by one thread. However, it is also possible to split up the plan into several thready. There are also several scheduler strategies available like "aurora min latency/memory/cost" or a round-robin.

## Extendable Parser

The parser is used for transforming a query (normally it's a string) into a logical query plan. It is possible to have multiple parser at one. There exists, e.g. PQL and CQL. PQL is a default parser where each operator can be expressed via a procedure. CQL is based on SQL and is similar to StreamSQL. It is also possible to integrate new languages. Furthermore, PQL can be easily extended for new operators by annotating the logical operator.

## Web service and Console Executor

The executor manages all things (installs and runs query or adds and removes sources). So the executor is the interface for external accesses. It can be used via code, a web service or a console. However, it is also possible to extend the executor to provide a new accessibility. The webservice interface, for example, can be used by other applications (even non-Java) to access to Odysseus.

## Odysseus Script

An own script language that is called "Odysseus Script" provides the possibility to run a set of queries or setting parameters through one script-file.

## Access Framework

The access framework of Odysseus is responsible for creating source and sink operators. For example, a source operator is used for connecting to a sensor to open a data socket where the sensor can push its data. The access framework provides several layers/parts which can be combined to build a suitable access operator. For example, there is a transport layer that describes how the data is provided (e.g. a TCP socket, as a file or through a serial port). Based on this, a protocol handler tells how the data is represented (e.g. as lines or text or byte buffers). This protocol handler uses a set of data handlers which say how the text or the bytes are transformed into the data. All these handlers can be extended as well.

## Buffer Placement

Sometimes it may be necessary that operators are not directly coupled. For this, it is possible to insert some buffer operators. Although they can be inserted by hand (e.g. via PQL), a build in buffer placement mechanism can be activated so that buffers are automatically placed into the query plan during the installation. There are also several buffer placement strategies like for each source or for each operator.

## Query Sharing

Query sharing is an optimization technique where Odysseus reuses existing partial plans when a new query plan is installed. For this, it only reuses a partial plan, if new and existing operators are semantically equal.

## User Management

Odysseus is a multi-user system. For this, each installed query or source is dedicated to a user. Since more than one user can access the system, it is also possible to grant or revoke special rights to other users.

## Punctuations / Heartbeat Mechanism

Odysseus has a built-in punctuation mechanism (in the relational processing). Since the window concept of the interval approach may cause a blocking of operators, because a processing step of an operator needs further elements to produce results. However, if there are no further elements, the processing blocks. At this point, punctuations indicate that the stream is still "alive" but there are currently no elements (thus, also called heartbeat). So, punctuation are used to unblock the operator earlier.

- Autostart
- Dashboard Feature
- Database Feature
- Interval Feature
- KeyValue Feature
- Key Value Store Feature
- Systemload Feature