



SPT Object Database  
System Manual  
Version 0.1.1

Prepared by  
Rakesh Vidyadharan

Prepared for  
Interested Developers/Deployers

Dated: July 12, 2008

## Contents

<b>1</b>	<b>Introduction</b>	<b>3</b>
1.1	Object Database . . . . .	3
1.2	Prevayler . . . . .	3
1.3	Source Repository . . . . .	4
<b>2</b>	<b>Project Goals</b>	<b>4</b>
<b>3</b>	<b>Database System</b>	<b>5</b>
3.1	PrevalentObject . . . . .	5
3.2	Prevalent System . . . . .	6
3.2.1	StorageSystem . . . . .	6
3.2.2	IndexSystem . . . . .	6
3.2.3	ConstraintSystem . . . . .	7
3.2.4	ObjectGraphSystem . . . . .	7
3.2.4.1	Decomposition . . . . .	8
3.2.4.2	Recomposition . . . . .	8
3.2.5	PrevalentSystem . . . . .	9
3.2.5.1	save . . . . .	9
3.2.5.2	fetch . . . . .	9
3.2.5.3	delete . . . . .	10
3.3	Transactions . . . . .	10
3.4	Queries . . . . .	10
3.5	PrevalentManager . . . . .	10
3.6	Configuration . . . . .	10

## List of Tables

1	Source file packages . . . . .	5
2	JVM System Properties . . . . .	11

## List of Figures

1	Storage Model Class Diagram . . . . .	7
---	---------------------------------------	---

## 1 Introduction

SPT Object Database is a simple Object Database developed around the [Prevayler](#) transactional object persistence framework. SPT Object Database attempts to address some of the most glaring deficiencies in [Prevayler](#) while making full use of its great features. See [1] for a presentation of object databases and [Prevayler](#).

### 1.1 Object Database

An object database is a persistent storage engine of objects<sup>1</sup> in much the same way that a relational database is a storage engine for raw data. The following are the basic attributes that we feel are required in an object database.

1. Persistent store for objects and object graphs.
  - 1.1. Persistence by reachability.
  - 1.2. Transparent handling of references to other persisted objects. Honour references when de-serialising a serialised persisted object.
2. Transactional semantics for object storage and management.
3. Query mechanism to retrieve objects.
4. Handling of indices and constraints declared on the objects being persisted.
5. Managed relationships - both uni-directional and bi-directional relationships between the persistent objects.
6. Fault tolerance.
7. Simple backup and recovery processes.
8. Simple to use without having to write a lot of code.

### 1.2 Prevayler

[Prevayler](#) is an open-source transactional object serialisation framework. Implementation of [Prevayler](#) are available for other technologies as well. As covered in [1], [Prevayler](#) is not a database. On the other hand, it makes users create something akin to an object database. The best part of [Prevayler](#) is its high degree of fault tolerance. In numerous tests, [Sans Pareil Technologies, Inc.](#) has been unable to cause any data corruption in the data persisted using [Prevayler](#).

---

<sup>1</sup>In our case only [Java](#) objects.

### 1.3 Source Repository

Full source code for the database engine and associated unit test suite are available in the [java.net subversion repository](#). The following instructions describe the process of checking out the source code and building the engine.

```
svn co https://sptodb.dev.java.net/svn/sptodb/trunk \  
    sptodb --username <username>  
cd sptodb  
cp ant.properties.sample ant.properties  
vi ant.properties # Edit as necessary to suit environment  
ant # Builds the jar file and cleans class files  
ant test # Runs the unit test suite
```

The **trunk** area of [java.net subversion repository](#) will contain the latest snapshot of code for the engine. The **tags** area will contain tagged versions with numbers reflecting the progress of the database engine.

## 2 Project Goals

The goal of this project is to create a simple object database for persisting and managing Java objects without requiring the user to have to write their own database engine. The following are the features that are planned for the database engine.

1. Transparently persist complex object graphs. Prevaler requires the developer to de-couple direct references to other persistent objects due to the inherent nature of [Java](#) object serialisation. This is obviously entirely contrary to all Object-Oriented principles and is an unjustifiable requirement for application development. This project will transparently handle the de-coupling and reconstitution of the object references thus preserving the purity of the object model.
2. Genericised transaction and query objects. [Prevaler](#) requires the developer to create distinct transaction and query objects for every type of operation the user wishes to perform against the store. This is a tedious and unnecessary process that may also be avoided by created generic transaction and query objects that may be used under a wide range of operating conditions.
3. Automatically index persistent objects by annotating the appropriate fields in the class. Annotations attempt to stay similar to [JDO](#) annotations.
4. Automatically index fields annotated as searchable using [Lucene](#) to support full-text search capability.
5. Enforce constraints declared through annotations on the persistent class while managing (adding/updating/deleting) persistent objects.

6. Managed relationships. Persistent objects maintain a relationship with other persistent objects when they maintain references to them either directly or in collections. The database engine will ensure that proper actions are taken when a persistent object that is in use by other persistent objects is to be removed from the store.

### 3 Database System

Source code for the SPT Object Database system is organised under the `com.sptci.prevayler` root package. Table 1 shows the packages used and their purpose. Note that the root `com.sptci` has been omitted from the package descriptions for succinctness.

Table 1: Source file packages

Pacakge	Description
<code>prevayler</code> <sup>a</sup>	The root package under which the core classes for the database engine reside.
<code>prevayler.annotations</code> <sup>b</sup>	The root package under which the annotations that are used to configure rules for the persisted objects reside.
<code>prevayler.query</code> <sup>c</sup>	The root package under which the general purpose queries implemented for the system reside.
<code>prevayler.transaction</code> <sup>d</sup>	The root package under which the standard transactions implemented for the system reside.

<sup>a</sup><http://www.sptci.com/products/odb/docs/com/sptci/prevayler/package-frame.html>

<sup>b</sup><http://www.sptci.com/products/odb/docs/com/sptci/prevayler/annotations/package-frame.html>

<sup>c</sup><http://www.sptci.com/products/odb/docs/com/sptci/prevayler/query/package-frame.html>

<sup>d</sup><http://www.sptci.com/products/odb/docs/com/sptci/prevayler/transaction/package-frame.html>

#### 3.1 PrevalentObject

SPT Object Database is designed to only persist instances of `PrevalentObject`<sup>2</sup>. `PrevalentObject` is an abstract class and as such cannot be used directly. The following reasons lead to Sans Pareil Technologies, Inc. deciding to make `PrevalentObject` an abstract class instead of a more flexible Interface.

- The `PrevalentSystem` needs to set the `objectId` for the prevalent object when not using application specified object id. This can be guaranteed only through the use of a super-class since interfaces cannot mandate the use of fields.
- The `PrevalentSystem` needs to maintain the metadata associated with the prevalent object that is maintained. This again requires the presence of a private field in the object.

<sup>2</sup><http://www.sptci.com/products/odb/docs/com/sptci/prevayler/PrevalentObject.html>

- Object id must be immutable. This is not possible to enforce by using an interface, since the implementation can chose to make the object id mutable via a publicly accessible mutator method.
- Interfaces have no implementation. An abstract base class allows the system to supply default implementations of critical methods such as `equals`<sup>3</sup> and `hashCode`<sup>4</sup> that sub-classes are discouraged from over-riding. Note that these methods have not been declared `final`, although that could change in future.

## 3.2 Prevalent System

The prevalent system is the primary class that is managed by an instance of `Prevayler`<sup>5</sup>. The system is broken up into an inheritance hierarchy based upon the facets managed by each class in the chain. The ultimate class in the chain is named `PrevalentSystem` and is used to initialise the `Prevayler` framework. Figure 1 shows the class diagram for the classes that comprise the prevalent system.

### 3.2.1 StorageSystem

The root of the inheritance hierarchy as depicted in figure 1. This `class`<sup>6</sup> is responsible for managing the various storage containers used to store and index the prevalent objects managed by the prevalent system. This class does not implement any rules for managing the persisted objects. All rules required for managing the persisted objects are implemented in the sub-classes as appropriate. Storage containers are usually `Map`<sup>7</sup> instances for efficient retrieval.

### 3.2.2 IndexSystem

The `class`<sup>8</sup> responsible for indexing the prevalent objects being persisted to the prevalent system. Indexing is performed by analysing the annotations defined on the prevalent class and populating the appropriate storage containers made available by `StorageSystem` (see section 3.2.1). This class is also responsible for enforcing any unique constraints defined on the prevalent objects.

---

<sup>3</sup>[http://www.sptci.com/products/odb/docs/com/sptci/prevayler/PrevalentObject.html#equals\(java.lang.Object\)](http://www.sptci.com/products/odb/docs/com/sptci/prevayler/PrevalentObject.html#equals(java.lang.Object))

<sup>4</sup>[http://www.sptci.com/products/odb/docs/com/sptci/prevayler/PrevalentObject.html#hashCode\(\)](http://www.sptci.com/products/odb/docs/com/sptci/prevayler/PrevalentObject.html#hashCode())

<sup>5</sup><http://docs.rakeshv.org/java/prevayler/org/prevayler/Prevayler.html>

<sup>6</sup><http://www.sptci.com/products/odb/docs/com/sptci/prevayler/StorageSystem.html>

<sup>7</sup><http://docs.rakeshv.org/java/j2sdk1.5/docs/api/java/util/Map.html>

<sup>8</sup><http://www.sptci.com/products/odb/docs/com/sptci/prevayler/IndexSystem.html>

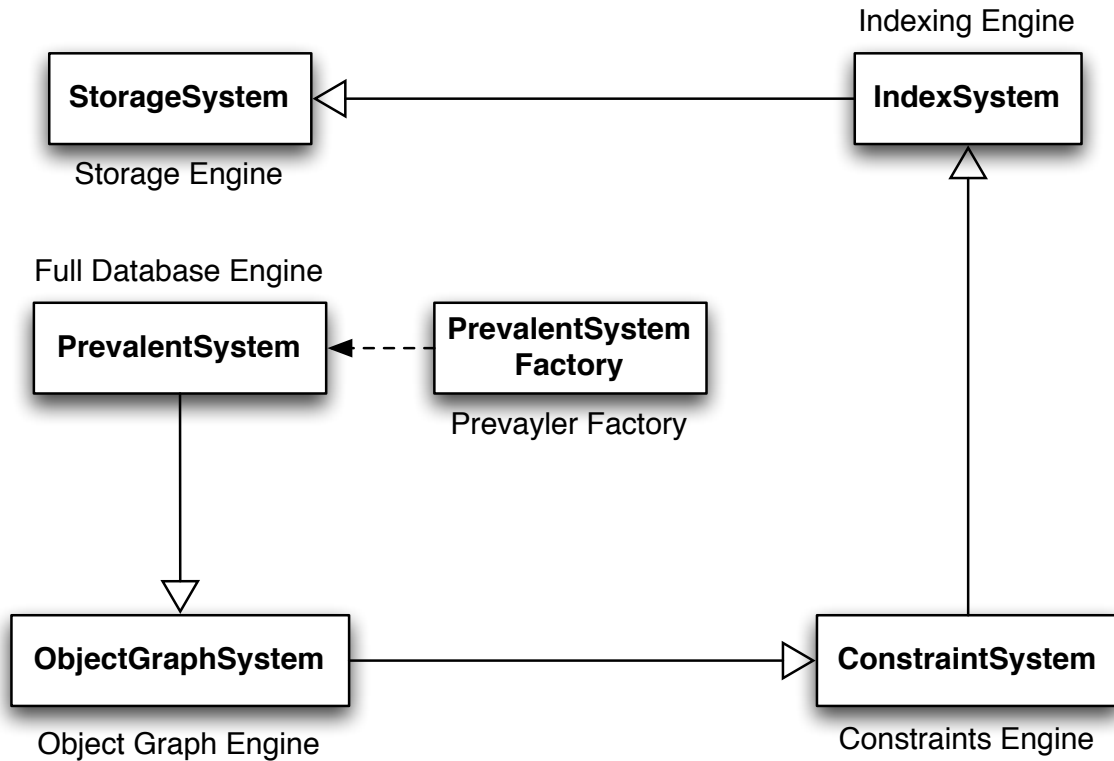


Figure 1: Storage Model Class Diagram

### 3.2.3 ConstraintSystem

The [class](#)<sup>9</sup> responsible for processing foreign key relations as annotated in the prevalent class. Foreign keys are implicitly indexed and hence fields annotated as foreign keys do not need to be additionally indexed. The operations defined in this class are complimentary to the operations defined in **IndexSystem** (see section 3.2.2). Similar to **IndexSystem** this class also enforces unique constraints defined for the foreign key fields.

### 3.2.4 ObjectGraphSystem

The [class](#)<sup>10</sup> responsible for decomposing and reconstituting the direct references from a prevalent object to other prevalent objects in the prevalent system.

<sup>9</sup><http://www.sptci.com/products/odb/docs/com/sptci/prevayler/ConstraintSystem.html>

<sup>10</sup><http://www.sptci.com/products/odb/docs/com/sptci/prevayler/ObjectGraphSystem.html>



**3.2.4.1 Decomposition** [Decomposition](#)<sup>11</sup> of a prevalent object is performed by applying the following rules:

1. Operations on direct references to other prevalent objects.
  - 1.1. Clone the prevalent object that is being persisted. This ensures that the objects persisted in the prevalent system are encapsulated from client code. Note that it is highly recommended that the [Object#clone](#)<sup>12</sup> method be implemented in all the prevalent classes in your system to ensure proper encapsulation of persisted objects.
  - 1.2. Persist the referenced object if it is a prevalent object and not already persisted.
  - 1.3. Add the reference key (the object id for the referenced object) along with the type information to a container managed specifically for instances of the object being persisted. This step is recursive ensuring that the entire object graph is decomposed.
  - 1.4. Set the field in the cloned prevalent object instance to `null`.
2. Operations on references to other prevalent objects stored as collections in the prevalent object being persisted.
  - 2.1. Clone the prevalent object that is being persisted. This ensures that the objects persisted in the prevalent system are encapsulated from client code.
  - 2.2. Clone the collection field as default cloning does not clone the objects contained in the object that was cloned.
  - 2.3. Persist the referenced objects in the collection if they are not already persisted.
  - 2.4. Add the reference keys (the object id for the referenced objects) along with the type information to a container managed specifically for instances of the object being persisted.
  - 2.5. Remove the processed prevalent object from the collection.

**3.2.4.2 Recomposition** [Recomposition](#)<sup>13</sup> of a prevalent object is performed when clients request instances of the object through a query. Recomposition is performed by applying the following rules:

1. Operations on direct references to other prevalent objects.

---

<sup>11</sup>[http://www.sptci.com/products/odb/docs/com/sptci/prevayler/ObjectGraphSystem.html#decompose\(com.sptci.prevayler.PrevalentObject\)](http://www.sptci.com/products/odb/docs/com/sptci/prevayler/ObjectGraphSystem.html#decompose(com.sptci.prevayler.PrevalentObject))

<sup>12</sup>[http://docs.rakeshv.org/java/j2sdk1.5/docs/api/java/lang/Object.html#clone\(\)](http://docs.rakeshv.org/java/j2sdk1.5/docs/api/java/lang/Object.html#clone())

<sup>13</sup>[http://www.sptci.com/products/odb/docs/com/sptci/prevayler/ObjectGraphSystem.html#compose\(com.sptci.prevayler.PrevalentObject\)](http://www.sptci.com/products/odb/docs/com/sptci/prevayler/ObjectGraphSystem.html#compose(com.sptci.prevayler.PrevalentObject))

- 1.1. Clone the prevalent object that was persisted. This ensures that the objects returned are encapsulated against unwitting modifications through client code. Note that it is highly recommended that the `Object#clone`<sup>14</sup> method be implemented in all the prevalent classes in your system to ensure proper encapsulation of persisted objects.
- 1.2. Reconstitute the prevalent object for the referencing field by fetching the recomposed instance based upon its object id as stored in the relevant container. This step is recursive and ensures that the entire object graph is recomposed.
- 1.3. Set the field in the cloned prevalent object instance to the recomposed prevalent object.
2. Operations on references to other prevalent objects stored as collections in the prevalent object being persisted.
  - 2.1. Clone the prevalent object that was persisted. This ensures that the objects returned are encapsulated against unwitting modifications through client code.
  - 2.2. Instantiate a new collection and populate it with the recomposed references based upon the object id values stored in the relevant container.
  - 2.3. Set the field in the cloned prevalent object instance to the initialised collection.

### 3.2.5 PrevalentSystem

This is the last `class`<sup>15</sup> in the inheritance hierarchy that constitutes a prevalent system that may be managed by a `Prevayler`<sup>16</sup> instance. This class defines the methods that are necessary to support object manipulation and query methods.

**3.2.5.1 save** This is the method that supports saving prevalent object instances to the prevalent system. Constraint checking as configured for the prevalent class through annotations are performed prior to persisting the object to the prevalent store. The prevalent object is decomposed (see section 3.2.4.1) prior to persisting. Note that this method may be used to add or update prevalent instances in the system.

**3.2.5.2 fetch** A few over-loaded `fetch` methods are provided for fetching prevalent instances from the prevalent system. Standard methods for fetching by `objectId`, by indexed fields, objects within a range ... are provided by default.

---

<sup>14</sup>[http://docs.rakeshv.org/java/j2sdk1.5/docs/api/java/lang/Object.html#clone\(\)](http://docs.rakeshv.org/java/j2sdk1.5/docs/api/java/lang/Object.html#clone())

<sup>15</sup><http://www.sptci.com/products/odb/docs/com/sptci/prevayler/PrevalentSystem.html>

<sup>16</sup><http://docs.rakeshv.org/java/prevayler/org/prevayler/Prevayler.html>

**3.2.5.3 delete** This is the method that supports deleting prevalent object instances from the prevalent system. Constraint checking is performed on objects that hold references (direct or via collections) to the object being deleted to determine if the object may be deleted.

### 3.3 Transactions

Genericised [Transaction](#)<sup>17</sup> objects are provided for invoking the `add` (see section 3.2.5.1) and `delete` (see section 3.2.5.3).

### 3.4 Queries

Genericised [Query](#)<sup>18</sup> objects are provided for the various standard queries supported by the [PrevalentSystem](#).

### 3.5 PrevalentManager

This is a façade around the [PrevalentSystem](#) that presents a more natural higher level API than that provided by [Prevayler](#). The interface is designed to be similar to [JDO PersistenceManager](#)<sup>19</sup> and abstracts operations on the [PrevalentSystem](#) through a simple direct API. Using the prevalent system through this façade is recommended over using the system through the [PrevalentSystem](#) and [Prevayler](#).

#### Note:

The [PrevalentManager](#) does not allow for customisation of the storage for the persisted objects. If you wish more control over the number of [PrevalentSystem](#) instances used to manager your object graph then you will need to use [PrevalentSystem](#) directly. Other aspects such as controlling the location of the persistent storage may be controlled as usual through JVM properties<sup>20</sup>.

### 3.6 Configuration

The prevalent system may be configured through JVM system properties. Table 2 shows the system properties that are supported and their purpose.

---

<sup>17</sup><http://docs.rakeshv.org/java/prevayler/org/prevayler/TransactionWithQuery.html>

<sup>18</sup><http://docs.rakeshv.org/java/prevayler/org/prevayler/Query.html>

<sup>19</sup><http://docs.rakeshv.org/java/jdo2/javax/jdo/PersistenceManager.html>

<sup>20</sup>See section 3.6 for details.

Table 2: JVM System Properties

Property	Default Value	Description
<code>sptodb.data.dir</code>	<code>/var/data/sptodb</code>	The root directory under which the database snapshot and journal files are stored. <a href="#">Lucene</a> search indices are also stored under this directory.
<code>sptodb.snapshot.interval</code>	86400	The interval in seconds at which snapshots of the prevalent system are to be taken. Note that snapshots reduce the start-up time of the prevalent system.
<code>sptodb.serializer.format</code>	java	<p>The format to use for taking snapshots of the prevalent system and creating transaction journals. The supported options are:</p> <p><b>java</b> - Indicates that regular Java object serialisation be used to take the snapshot and write journals.</p> <p><b>xml</b> - Indicates that the journals should be written and snapshot taken using <a href="#">XStream</a>. XML serialisation is slower (both serialisation and deserialisation), however gives you more options for processing the prevalent system for other purposes. Also useful if you need to restore the system after heavy modifications (refactoring) to the object model.</p>

## References

- [1] *Object Databases*<sup>21</sup> - *CJUG*<sup>22</sup> Presentation.

<sup>21</sup> <http://www.sptci.com/products/articles/ObjectDatabases.pdf>

<sup>22</sup> <http://www.cjug.org/Wiki.jsp?page=2008.5.20.downtown>