



SQL Replication

- *Replication Strategies*
- *Evaluated Technologies*
- *Enhydra Octopus:*
 - *Architecture*
 - *Setup*
 - *Experience in Atlas:*
 - *GeometryDB: Oracle -> MySQL*
 - *ConditionsDB: MySQL -> MySQL*



Replication Ways

- *Native Tools (Oracle, MySQL,...)*
- *Generic Converters (Octopus,...)*
- *Actual Applications (JDBC-based, Pool/RAL-based)*



Replication Use Cases

- Replication within technology (MySQL->MySQL,...):
 - *Native Tools are fast and sure*
 - *Generic Converters provide unified API*
- Replication between technologies (Oracle->MySQL,...):
 - *If same Schema in both technologies: Generic Tools do well*
 - *If almost same Schema in both technologies: Generic Tools can be configured with customized mapping (using configuration files or special plugins)*
 - *If different Schema in both technologies or data have to be modified: Actual Applications is usually the best choice*

Evaluated Technologies



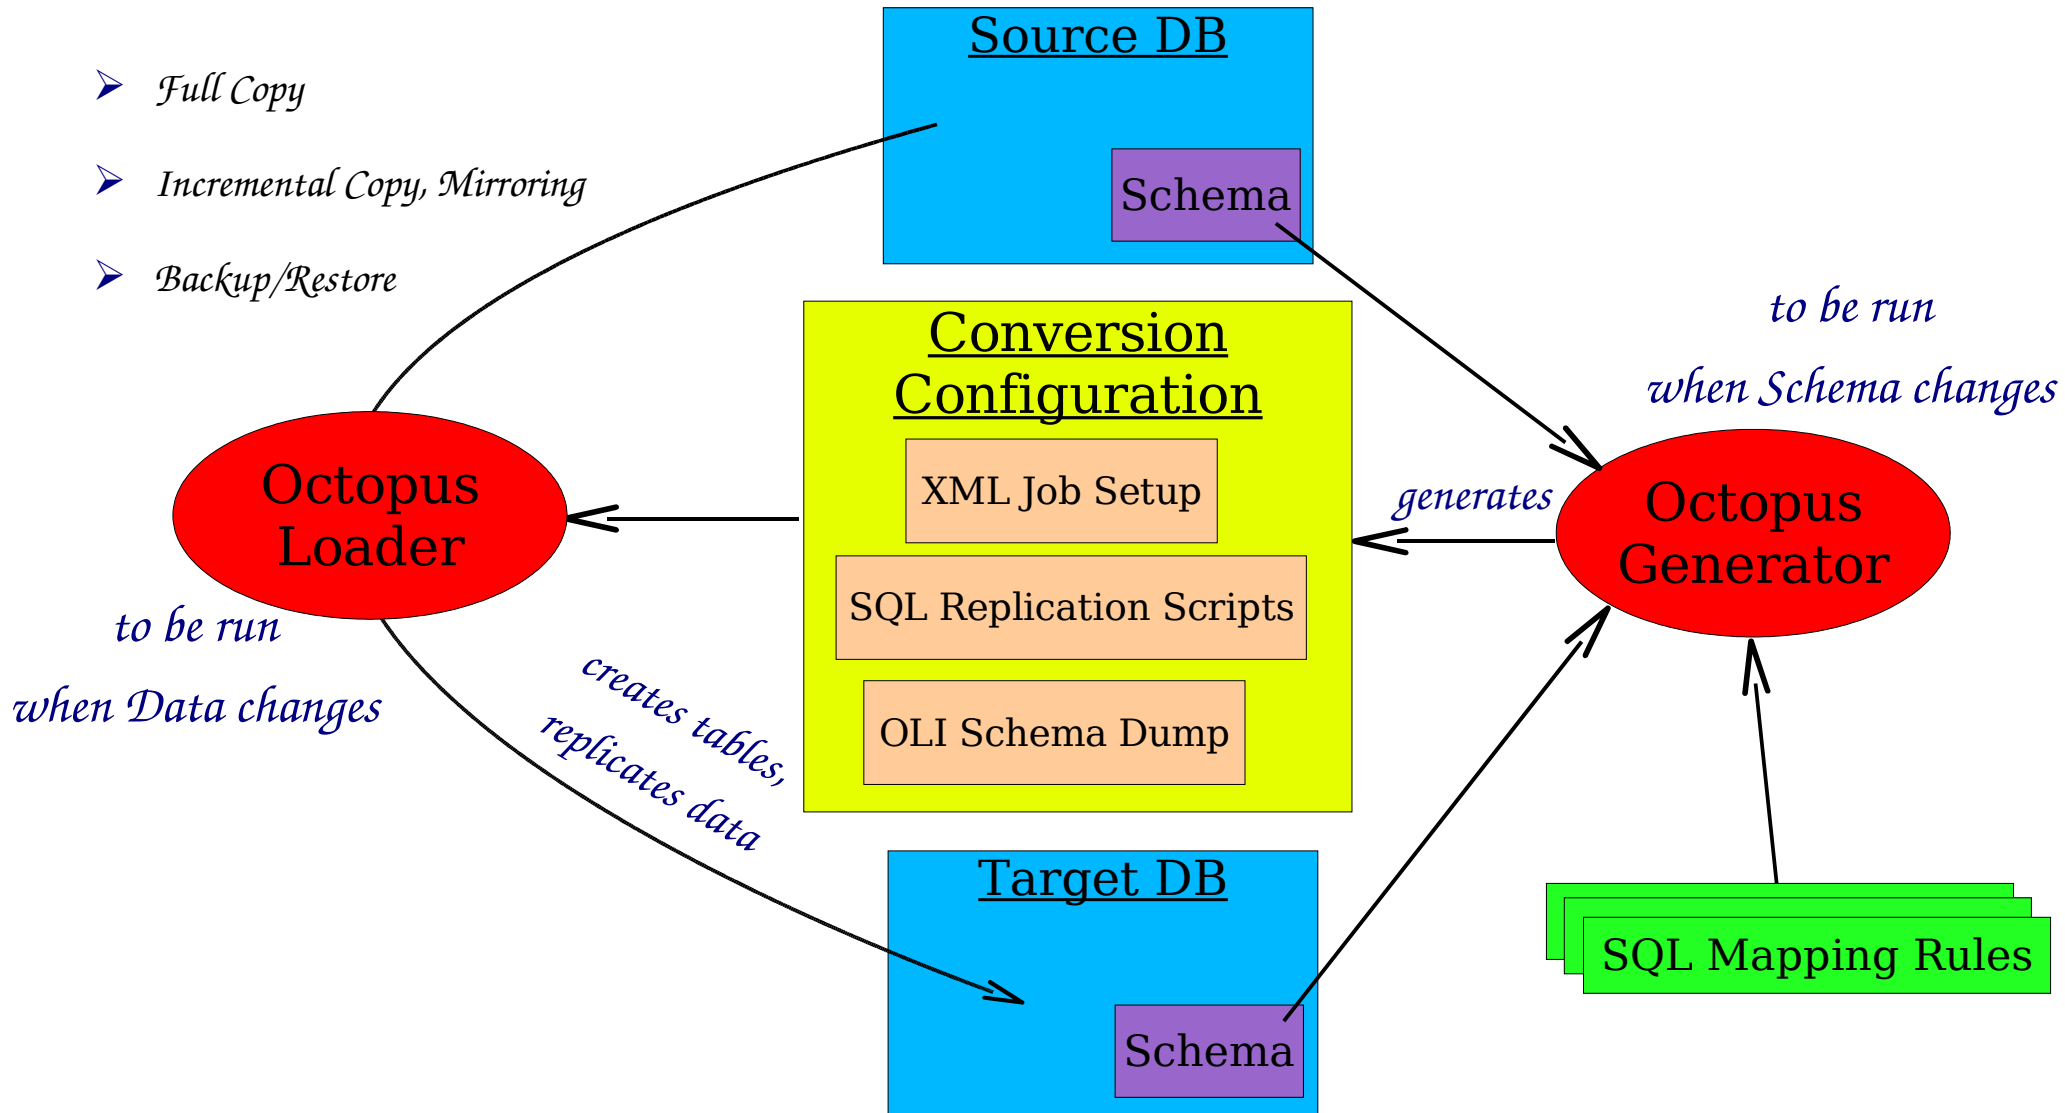
(for the heterogeneous replication)

- SQLTuple/ColMan (AIDA NTuple API for SQL databases): good for AttributeList/NTuple or similar mapping
 - Is able to filter and merge while replicating Pool AttributeList on MySQL, Oracle, PostgreSQL and McKoi
- Groovy-based Replicator: learning tool, simple, easy, basic
- JDBC Importer: a bit different mission, needs a lot of customizations
- Enhydra Octopus: seems to be the best match
 - Replicates (set of) tables keeping all constrains (keys, indexes,...)

Octopus Architecture

- Support for Access, DB2, MSQl, Paradox, CJDDBC, Excel, Informix, McKoi, PostgreSQL, Sybase, CSV, Hypersonic, Instantdb, MySQL, Oracle, QED and XML.
- Modes:

- Full Copy
- Incremental Copy, Mirroring
- Backup/Restore





Octopus Setup

- Ant build.xml file
- Run description build.properties file
- Customised mapping map.properties file
(designed to satisfy special Atlas requirements)

```
Src.Db=@sundb07.cern.ch:1521:pdb01-1
Src.User=atlasdd
Src.Passwd=bla
Src.Schema=ATLASDD

Dest.Db=atlasdbdev.cern.ch:3306/CollectionTest
Dest.User=CollTester
Dest.Passwd=bla

Log.Mode=full

Tables=ALIN_DATA;ALIN_DATA2TAG
```

```
# Clean everything
ant clean

# Generate replication scripts
ant generate

# Perform the replication
ant load

# Recompile Atlas customizations
ant patch

# Get help
ant -projecthelp
```

```
octopus.mysql.varchar2(4000).type=varchar
octopus.mysql.varchar2(4000).length=255

octopus.mysql.float(126).type=double

octopus.mysql.float(63).type=float

octopus.mysql.number(1).type=tinyint
octopus.mysql.number(1).length=1

octopus.mysql.number(10).type=integer
octopus.mysql.number(10).length=10
```



GeometryDB Replication

- Oracle->MySQL
- Fixed problems:
 - Support for Oracle/varchar2 to MySQL/varchar, etc.: The requirement is to convert Oracle/varchar2 into MySQL/varchar. The problem is that Atlas GeometryDB contains varchar2(4000), while MySQL allows the longest varchar(255). Standard Octopus doesn't support that (of course). This is now supported by Atlas Patch to Octopus via map.properties file. Note, that those user-supplied rules can be highly unsecure.
 - Support for replication of all tables belonging to an Oracle Schema: This was a missing feature in Octopus (Octopus could replicate either the whole server or just an explicit list of tables). It has been implemented and fed back to Octopus developers.
 - Unique Index on multiple columns: This was a bug of Octopus. It has been fixed.



Conditions DB Replication

- *MySQL->MySQL*
- *Fixed problems:*
 - *MySQL tables contain a Primary Index with the name "PRIMARY":* While some MySQL commands accept that, it is not generally supported as MySQL identifier because "PRIMARY" is a MySQL reserved word. Atlas Octopus Patch now creates commands which accept "PRIMARY" as a Primary Index name. Note, that using reserved words as Identifiers is a bug in a Database design.



Documentation

- *Enhydra Octopus is provided by ObjectWeb Consortium, released under GPL, it has active user base and responsive developers.*
- *Reusable Atlas-motivated fixes and modifications fed back to Octopus developers to will be included in the distribution.*
- *Octopus has already helped us to find bugs in our SQL Schemas.*
- *Documentation:*
 - *Home: <http://octopus.objectweb.org>*
 - *Atlas Wiki: <https://uimon.cern.ch/twiki/bin/view/Atlas/DatabaseReplication>*