















# Database Related Issues

## Backlog

Number	Notes	Relates To	Type	Assignee	Rank
 <b>POLICY-1824</b> - Persistence of run time policy state in Pdps <span>CLOSED</span>	<p>Investigate ways of persisting policy state information with different structures at run time. One approach could be to use Avro &amp; Apache Hive. [There is also a NoSQL/JSON option available in MariaDB and Postgres]</p> <p>Another option is the <a href="#">Apache Cassandra</a> database. Datastax provide a Java Driver for Apache Cassandra.</p> <p>Writing plugins towards state information stored in CPS might be another approach that could work across all the PDPs.</p>	POLICY-2898	Epic		5
 <b>POLICY-2898</b> - Policy should function in a multi-cluster environment <span>CLOSED</span>	<a href="#">Multi-cluster database support</a>	POLICY-1821	Story		1
 <b>POLICY-2745</b> - Allow underlying database to be configured: MariaDB or Postgres <span>CLOSED</span>	Allow DB to be changed based on user needs e.g Postgres for MariaDB	POLICY-1787	Story		3 (xacml & drools need to be tested with postgres)
 <b>POLICY-1787</b> - Support mariadb upgrade /rollback functionality <span>CLOSED</span>	<a href="#">Upgrade/rollback of database tables</a>	POLICY-2715	Story	Kevin Timoney	1 (follows backup /restore JIRA)
 <b>POLICY-2086</b> - Remove references to mariadb from resource files <span>CLOSED</span>	<p>References to DB should be moved from persistence.xml to properties file to facilitate the use of alternate databases.</p> <p>We can modify the current code to read in the location of a properties file</p> <div> <b>JDBC properties</b> <pre> -Djdbc.properties=/path/to/jdbc.properties String propertiesPath = System.getProperty( "jdbc.properties" );  if ( propertiesPath != null ) {     FileInputStream in = new FileInputStream ( propertiesPath );      try     {         jdbcProperties = Properties.load( in );     }     finally     {         in.close( );     } } </pre> </div> <p><b>Kubernetes</b></p> <p>There are four different ways that you can use a ConfigMap to configure a container inside a Pod:</p> <ul style="list-style-type: none"> <li>• Inside a container command and args</li> <li>• Environment variables for a container</li> <li>• Add a file in read-only volume, for the application to read</li> <li>• Write code to run inside the Pod that uses the Kubernetes API to read a ConfigMap</li> </ul> <p><a href="#">Kubernetes Configmap</a></p>		Story		3

 <b>POLICY-3059</b> - Fix name of target-database property in persistence.xml files <span>CLOSED</span>	<p>Invalid target-database property in persistence.xml in apex pdp.</p> <p>Persistence.xml should use eclipselink.target-database</p> <p>e.g. &lt;property name="eclipselink.target-database" value="MySQL" /&gt;</p>		Task	Ajjith Sreekumar <ul style="list-style-type: none"> <li>reassign to Kevin</li> </ul>	3 (bug fix)  <b>Done</b>
 <b>POLICY-4837</b> - Review transaction boundaries of models <span>CLOSED</span>	<p>Currently, the models Provider classes manage transactions. Transaction management should be moved to client for better performance and atomicity. This will also eliminate the need for caching on the client side.</p> <p>Note: The transactional annotation can be used to group transactions together.</p> <p>Ideally transactions should be executed serially in order to avoid dirty reads, non-repeatable reads and phantom reads.</p> <p>Note: DatabasePolicyModelsProviderImpl in the models-provider package provides rest APIs such as createServiceTemplate, updateServiceTemplate and deleteServiceTemplate which in turn call the corresponding APIs from AuthorativeToscaProvider in the models-tosca package..</p>		Story		5 (depends on Spring direction)
 <b>POLICY-2540</b> - Proper handling of data types in policy-models and policy-api <span>CLOSED</span>	<p>How best to deal with CRUD of data types in policy-api and policy-models. It is possible only to create data types indirectly in policy type create requests. Update and delete of data types is not possible.</p> <p>SimpleToscaProvider contains the following methods not available in AuthorativeToscaProvider : getDataTypes, getCascadedDataTypes, createDataTypes, updateDataTypes and deleteDataType.</p>		Story	Assign to Liam	5
 <b>POLICY-3234</b> - Implement backup and restore functionality for PolicyDB <span>CLOSED</span>  <b>POLICY-3000</b> - Create an ability to recover if policy database gets corrupted <span>CLOSED</span>	<p>Database backup and restore</p>	MDEV-23119	Story		1  rename to backup/restore
 <b>POLICY-2747</b> - Multi-Tenancy support for Policy Framework <span>CLOSED</span>	<p>What policy components needs to be centralized vs de-centralized and moved to tenant namespaces. (R8 ONAP to support multi-tenancy)</p> <p><a href="#">Multi-tenancy support in Policy Framework</a></p> <p>Note: EclipseLink supports <a href="#">Using Table-Per-Tenant Multi-Tenancy</a></p>		Story		8 (design is available)
 <b>POLICY-2997</b> - Clean up old records from various DB tables <span>CLOSED</span>	<p>Clean up/Roll up of old DB data. Purge/Archive job should be created and run at a scheduled interval.</p> <p><a href="#">MariaDB Event Scheduler</a>   <a href="#">Postgres pg_cron</a></p> <p>MariaDB allows for the partitioning of tables : <a href="#">Partitioning Overview</a></p>		Story		4 (for purge)  6 (roll-up)  create separate tickets for purge & roll-up, and separate tickets for operations history and statistics
 <b>POLICY-3453</b> - Fix Db connection issues in TOSCA control loop <span>CLOSED</span>	<p>Concurrent DB access issues in control loop POC.</p>		Bug	Ramesh Murugan Iyer	bug (related to transaction handling ?)
 <b>POLICY-3209</b> - CLAMP Component Lifecycle Management using Spring Framework <span>CLOSED</span>	<p>Investigate what's involved in switching to spring</p>		Story	Liam Fallon	4
	<p><a href="#">Externalizing ONAP DBs to a separate namespace</a></p> <p>We would like to improve the current situation by implementing a two staged deployment:</p> <ol style="list-style-type: none"> <li>1) Deploy all required DB engines (can be done using community charts or any user chart)</li> <li>2) Deploy ONAP components &amp; configure them to make use of those engines</li> </ol> <p>This would allow use to basically bring his own database for ONAP no matter if it is running in the same k8s cluster or it has been provided by some DBaaS solution. Additionally it makes our deployment more modular and configurable thus may result in significant footprint savings.</p>				6
 <b>POLICY-3456</b> - Review the design of storing of PDP statistics in the DB <span>CLOSED</span>	<p>How much statistics should be stored</p>		Task		

<div><div><input checked="" type="checkbox"/> <b>POLICY-3404</b> - Rolling DB errors in log output for API, PAP, and DB components</div><div>CLOSED</div></div>	JPA table-creation errors		Task		2

Jira relationships