

# Task partitioning

Similar to workers management, task partitioning can be done manually or automatically. Unlike `WorkersCreationTaskHandler`, there are (potentially) many handlers that provide automatic partitioning: *generic partitioning task handler* and handlers for specific tasks, currently represented by *partitioned reconciliation task handler*. The difference between generic and specific partitioning handlers is that the specific ones supply reasonable defaults for the partitioning. For example, reconciliation task handler always creates three partitions corresponding to three stages of reconciliation (operation completion, resource reconciliation, shadow reconciliation).

Here's a short description how to configure generic and specific partitioning handlers.

## An example

```
<task oid="8f8de5ad-e699-439e-8362-77cbb994117c"
  xmlns="http://midpoint.evolveum.com/xml/ns/public/common/common-3"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xmlns:xsd="http://www.w3.org/2001/XMLSchema"
  xmlns:noop="http://midpoint.evolveum.com/xml/ns/public/task/noop/handler-3">
  <name>Partitioned single-node task</name>
  <extension>
    <noop:delay xsi:type="xsd:int">1000</noop:delay>
    <noop:steps xsi:type="xsd:int">10</noop:steps>
  </extension>
  <ownerRef oid="00000000-0000-0000-0000-000000000002"/>
  <executionStatus>runnable</executionStatus>
  <category>Demo</category>
  <handlerUri>http://midpoint.evolveum.com/xml/ns/public/task/generic-partitioning/handler-3</handlerUri>
  <workManagement>
    <taskKind>partitionedMaster</taskKind>
    <partitions>
      <count>4</count>
      <handlerUri>http://midpoint.evolveum.com/xml/ns/public/task/noop/handler-3#{index}</handlerUri>
      <copyMasterExtension>true</copyMasterExtension>
    </partitions>
  </workManagement>
  <recurrence>single</recurrence>
</task>
```

The `workManagement/partitions` defines here the following:

1. There are 4 partitions.
2. Subtasks will execute handlers with the URIs of:
  - a. <http://midpoint.evolveum.com/xml/ns/public/task/noop/handler-3#1>
  - b. <http://midpoint.evolveum.com/xml/ns/public/task/noop/handler-3#2>
  - c. <http://midpoint.evolveum.com/xml/ns/public/task/noop/handler-3#3>
  - d. <http://midpoint.evolveum.com/xml/ns/public/task/noop/handler-3#4>
3. Subtasks will get task extension from the master (so they receive settings of delay = 1000, steps = 10).

## Definition options

Here we describe `partitions` configuration section.

Item	Description	Default value
count	Number of partitions to create. Optional.	Number of explicitly defined partitions.
sequentialExecution	Whether the partition tasks should be executed sequentially (true) or in parallel (false).	true
taskName	How to derive names for partition tasks. It is possible to use substitution strings of <code>{masterTaskName}</code> , <code>{masterTaskHandlerUri}</code> and <code>{index}</code> here.	<code>{masterTaskName} ({index})</code>
handlerUri	Worker tasks handler URI. It is possible to use the above substitution strings here as well.	<code>{masterTaskHandlerUri}#{index}</code> or <a href="http://midpoint.evolveum.com/xml/ns/public/task/workers-creation/handler-3">http://midpoint.evolveum.com/xml/ns/public/task/workers-creation/handler-3</a> (for subtasks that are coordinators)

workManagement	Work management configuration to use for partition task.	-
copyMasterExtension	Whether to copy task extension from the master task into partition tasks.	false
otherDeltas	Other deltas to be applied to newly created partition tasks.	-
partition (multivalued)	Definition of the partition(s). May be omitted if there are no specific configuration elements for individual partitions.	-

Individual partitions are configured using the following items. They override items in the main configuration (unless noted otherwise).

Item	Description
index	The number of the partition, starting at 1. It might be omitted if all partitions are explicitly defined, and all of them are unnumbered.
taskName	How to derive the name for this partition task.
handlerUri	How to derive the handler URI for this partition task.
workManagement	Work management configuration for this partition task.
copyMasterExtension	Whether to copy task extension from the master task to this partition task.
otherDeltas	Other deltas to be applied to this partition task. Applied along with deltas from the parent element.
dependents (multivalued)	Dependents of this subtask, i.e. subtasks that should be started only after this subtask has finished. Provided as indices starting at 1. If the sequential execution is chosen, this setting is applied along with the default dependencies of 1 -> 2, 2 -> 3, ..., N-1 -> N.

## Two examples of partitioned multi-node reconciliation

The following task will start a partitioned reconciliation, whose second stage is separated into buckets and executes in 4 worker tasks on each node.

```

<task oid="604866a6-de5e-4cfe-abb4-7aed95505deb"
  xmlns="http://midpoint.evolveum.com/xml/ns/public/common/common-3"
  xmlns:syncext="http://midpoint.evolveum.com/xml/ns/public/model/extension-3"
  xmlns:ri="http://midpoint.evolveum.com/xml/ns/public/resource/instance-3"
  xmlns:icfs="http://midpoint.evolveum.com/xml/ns/public/connector/icf-1/resource-schema-3"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xmlns:xsd="http://www.w3.org/2001/XMLSchema">

  <name>Partitioned multi-node reconciliation</name>

  <extension>
    <syncext:objectclass>ri:AccountObjectClass</syncext:objectclass>
  </extension>

  <ownerRef oid="00000000-0000-0000-0000-000000000002"/>
  <executionStatus>runnable</executionStatus>

  <handlerUri>http://midpoint.evolveum.com/xml/ns/public/model/synchronization/task/partitioned-reconciliation
/handler-3</handlerUri>
  <objectRef oid="ef2bc95b-76e0-48e2-86d6-3d4f02d3fafe" type="ResourceType"/>
  <workManagement>
    <partitions>
      <partition>
        <index>2</index>
        <workManagement>
          <taskKind>coordinator</taskKind>
          <buckets>
            <stringSegmentation>
              <discriminator>attributes/icfs:name</discriminator>
              <matchingRule>stringIgnoreCase</matchingRule>
              <!-- buckets are like: (start) -> aa, aa -> ab, ab -> ac, ac -> ad, ..., zy -> zz,
zz -> (end) -->
              <boundaryCharacters>abcdefghijklmnopqrstuvwxyz</boundaryCharacters>
              <boundaryCharacters>abcdefghijklmnopqrstuvwxyz</boundaryCharacters>
            </stringSegmentation>
          </buckets>
          <workers>
            <workersPerNode>
              <count>4</count>
            </workersPerNode>
          </workers>
        </workManagement>
      </partition>
    </partitions>
  </workManagement>
  <recurrence>single</recurrence>
</task>

```

And here is analogous situation, this time segmenting numeric space of 64-bit number attribute (i.e. from 0 to  $2^{64}-1$  inclusive) into 128 buckets.

```

<task oid="9fa6d045-06d8-4549-9b7c-599b92f202f3"
  xmlns="http://midpoint.evolveum.com/xml/ns/public/common/common-3"
  xmlns:syncext="http://midpoint.evolveum.com/xml/ns/public/model/extension-3"
  xmlns:ri="http://midpoint.evolveum.com/xml/ns/public/resource/instance-3"
  xmlns:icfs="http://midpoint.evolveum.com/xml/ns/public/connector/icf-1/resource-schema-3"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xmlns:xsd="http://www.w3.org/2001/XMLSchema">

  <name>Partitioned multi-node reconciliation [numeric]</name>

  <extension>
    <syncext:objectclass>ri:AccountObjectClass</syncext:objectclass>
  </extension>

  <ownerRef oid="00000000-0000-0000-0000-000000000002"/>
  <executionStatus>runnable</executionStatus>

  <handlerUri>http://midpoint.evolveum.com/xml/ns/public/model/synchronization/task/partitioned-reconciliation
/handler-3</handlerUri>
  <objectRef oid="ef2bc95b-76e0-48e2-86d6-3d4f02d3fafa" type="ResourceType"/>
  <workManagement>
    <partitions>
      <partition>
        <index>2</index>
        <workManagement>
          <taskKind>coordinator</taskKind>
          <buckets>
            <numericSegmentation>
              <discriminator>attributes/number</discriminator>
              <to>18446744073709551616</to>      <!-- 2^64 -->
              <numberOfBuckets>128</numberOfBuckets>
            </numericSegmentation>
          </buckets>
          <workers>
            <workersPerNode>
              <count>4</count>
            </workersPerNode>
          </workers>
        </workManagement>
      </partition>
    </partitions>
  </workManagement>
  <recurrence>single</recurrence>
</task>

```

Note that the connector must correctly support comparison of the number attribute typed as BigInteger.