

MidPoint Community Meetup 2025

Performance Tuning



Anton Tkáčik, May 2025 Backend Technical Leader

Agenda

- Performance Tuning of New Features
 - Provenance Metadata
 - Shadow Caching
 - Shadow Partitioning
- General Performance Tuning
 - Tasks, Threads & Heap
 - Operation Execution Recording
 - PostgreSQL Vacuuming





Provenance Metadata

- Introduced in 4.2, extended in 4.9
- By default enabled for multivalue properties, assignments
- Simplifies mapping of multivalue properties
- Performance Impact
 - More data per value more heap, more processing, larger DB size

```
"p-multi-0000": [
    "@value": "source-000:ad1TENquy6",
    "@metadata": {
        "@id": 48,
        "provenance": {
        "mappingSpecification": {
            "definitionObjectRef": {
                "oid": "23aa2a5a-0d9d-44fc-b6f9-28d16c5ca5ce",
                "relation": "org:default",
                "type": "c:ResourceType"
            },
            "objectType": {
                "kind": "account",
                "intent": "default"
```



Provenance Metadata - Mitigation

```
<internals>
                                                                                                                                                                                                                                                                                                                  <shadowCaching>
                                                                                                                                                                                                                                                                                                                                        <defaultPolicy>
• Can be turned off for multivalue properties
                                                                                                                                                                                                                                                                                                                                                              <cachingStrategy>passive</cachingStrategy>
                                                                                                                                                                                                                                                                                                                                        </defaultPolicy>
              No multivalue property mappings
                                                                                                                                                                                                                                                                                                                  </shadowCaching>
                                                                                                                                                                                                                                                                                                                  <tracing>...</tracing>
Changes can be done in System Configuration
                                                                                                                                                                                                                                                                                                                   <repository>...</repository>
              internals/valueMetadata/
                                                                                                                                                                                                                                                                                                                  <valueMetadata>
                           disableDefaultMultivalueProvenance
                                                                                                                                                                                                                                                                                                                                        <disableDefaultMultivalueProvenance>true</disableDefaultMultivalueProvenance>true</disableDefaultMultivalueProvenance>true</disableDefaultMultivalueProvenance>true</disableDefaultMultivalueProvenance>true</disableDefaultMultivalueProvenance>true</disableDefaultMultivalueProvenance>true</disableDefaultMultivalueProvenance>true</disableDefaultMultivalueProvenance>true</disableDefaultMultivalueProvenance>true</disableDefaultMultivalueProvenance>true</disableDefaultMultivalueProvenance>true</disableDefaultMultivalueProvenance>true</disableDefaultMultivalueProvenance>true</disableDefaultMultivalueProvenance>true</disableDefaultMultivalueProvenance>true</disableDefaultMultivalueProvenance>true</disableDefaultMultivalueProvenance>true</disableDefaultMultivalueProvenance>true</disableDefaultMultivalueProvenance>true</disableDefaultMultivalueProvenance>true</disableDefaultMultivalueProvenance>true</disableDefaultMultivalueProvenance>true</disableDefaultMultivalueProvenance>true</disableDefaultMultivalueProvenance>true</disableDefaultMultivalueProvenance>true</disableDefaultMultivalueProvenance>true</disableDefaultMultivalueProvenance>true</disableDefaultMultivalueProvenance>true</disableDefaultMultivalueProvenance>true</disableDefaultMultivalueProvenance>true</disableDefaultMultivalueProvenance>true</disableDefaultMultivalueProvenance>true</disableDefaultMultivalueProvenance>true</disableDefaultMultivalueProvenance>true</disableDefaultMultivalueProvenance>true</disableDefaultMultivalueProvenance>true</disableDefaultMultivalueProvenance>true</disableDefaultMultivalueProvenance>true</disableDefaultMultivalueProvenance>true</disableDefaultMultivalueProvenance>true</disableDefaultMultivalueProvenance>true</disableDefaultMultivalueProvenance>true</disableDefaultMultivalueProvenance>true</disableDefaultMultivalueProvenance>true</disableDefaultMultivalueProvenance>true</disableDefaultMultivalueProvenance>true</disableDefaultMultivalueProvenance>true</disableDefaultMultivalueProvenance>true</disableDefaultMultivalueProvenance>t
                                                                                                                                                                                                                                                                                                                  </valueMetadata>
                                                                                                                                                                                                                                                                                            </internals>
```



Shadow Caching

- Introduced in 4.9
- Enabled by default in useFresh mode for all resources
- Caches shadow attribute, activation and password in repository
 - Password usually only mark that it is present
 - Allows for repository search using attributes
- Performance Impact
 - More data per shadow more processing (additional DB writes), larger audit records

```
"attributes" : {
    "@ns" : "http://midpoint.evolveum.com/
    "name" : "michelangelo",
    "ri:first_name" : {
        "@type" : "xsd:string",
        "@value" : "Michelangelo"
    },
    "ri:last_name" : {
        "@type" : "xsd:string",
        "@type" : "xsd:string",
        "@value" : "di Lodovico Buonarroti S
    },
    "uid" : "michelangelo"
},
```



Shadow Caching - Mitigation

- Can be configured / turned off globally or per resource bases
- Slow resources: useCachedOrFresh
- Fast resources turn off maitenance of retrievalTimestamp
 - Can be done on System Configuration, Resource or Object Type level





Shadow Partitioning

- Introduced in 4.9
 - Enabled by default for new deployments
- Uses separate partitions for shadows per resource / object class
 - Faster queries, updates
- Can be enabled usign Admin GUI → About → Enable Shadow Partitioning
 - This may take long if you have lot of shadows

<systemConfiguration> <internals> <repository> <autoCreatePartitionsOnAdd>true< </repository> </internals> </systemConfiguration>



Tasks & Threads

- Most tasks have worker threads (configurable) and 1 coordinator thread
- Worker threads rule of thumb (for start)
 - Thread number should not exceed 2-times available CPU cores
- Increasing number of threads impacts required heap size



Java Heap Sizing

- Heap size depends on multiple factors
 - Number of threads
 - Number of resource accounts per user
 - Number of role memberships per user
- Careful with heap sizes over 32GB
 - By default JVM will switch to 8 byte pointers almost doubling size of some objects
 - Still possible to use 4 byte pointers by changing JVM object allignment



Operation Execution Recording

- Captures recent operations executed on object
 - In addition to audit, intended to be technical
- Stored 10 records per object
 - 5 simple records usually success
 - 5 complex records usually errors
- Changes can be done in System Configuration → Internals → operationExecutionRecording → skipWhenSuccess

```
<operationExecution id="2">
   <recordType>simple</recordType>
   <timestamp>2025-03-31T00:01:29.765Z</timestamp>
   <operation>
       <objectDelta>
           <t:changeType>modify</t:changeType>
           <t:objectType>c:ShadowType</t:objectType>
       </objectDelta>
       <executionResult>
           <operation>com.evolveum.midpoint.model.impl.lens.Ch
           <status>success</status>
           <importance>normal</importance>
           <token>10000000000002495</token>
       </executionResult>
       <objectName>uid=leonardo.ou=people.dc=example.dc=com</or>
       <resourceOid>ebd0bf7b-7e80-4175-ba5e-4fd5de2ecd62</reso
       <resourceName>LDAP Server (OpenLDAP)</resourceName>
       <shadowKind>account</shadowKind>
       <shadowIntent>default</shadowIntent>
   </operation>
   <status>success</status>
   <!-- administrator -->
   </initiatorRef>
   <taskRef oid="486f0af0-990b-4f15-b81e-17f9a8e90d1a" relation
       <!-- HR import -->
   </taskRef>
   <activityPath/>
   <channel>http://midpoint.evolveum.com/xml/ns/public/common/
</operationExecution>
```



PostgreSQL Vacuuming

- Tasks may do lots of inserts / updates
 - PostgreSQL may after large set of insert / updates opt to use full scan instead of indexes
 - Tables may need to be analyzed / vacuumed
- PostgreSQL does automatic scanning and analyze
- Changes can be done PostgreSQL configuration
 - autovacuum_naptime = 10s
 - Changes periodicity of vacuuming scans





Where to learn more

- Longer detailed version at our documentation
 - Midpoint / Reference / Diagnostics / Performance
- Search our documentation for Performance and Performance Tuning
- PostgreSQL vacuum analyze
 - https://www.postgresql.org/docs/current/runtime-config-au tovacuum.html
- Java Heap Sizing
 - https://www.baeldung.com/jvm-compressed-oops





Evolveum Thank you for your attention

Feel free to ask your questions now!



MidPoint Community Meetup 2025