



MidPoint 4.4 - Upgrade and migration

Anton Tkáčik / January 2022



- Short overview of changes
- How to upgrade to midPoint 4.4
- How to migrate to PostgreSQL native repository

| Changes since 4.0 affecting upgrade

- Java 8 platform is no longer supported
 - Java 11 or Java 17 is recommended
- MySQL and MariaDB are no longer supported,
 - there is upgrade path to 4.4 and then migration to native repository
- Many deprecated elements were removed from schema
- Jasper-based reports are no longer supported
- Tasks should be reimported / updated

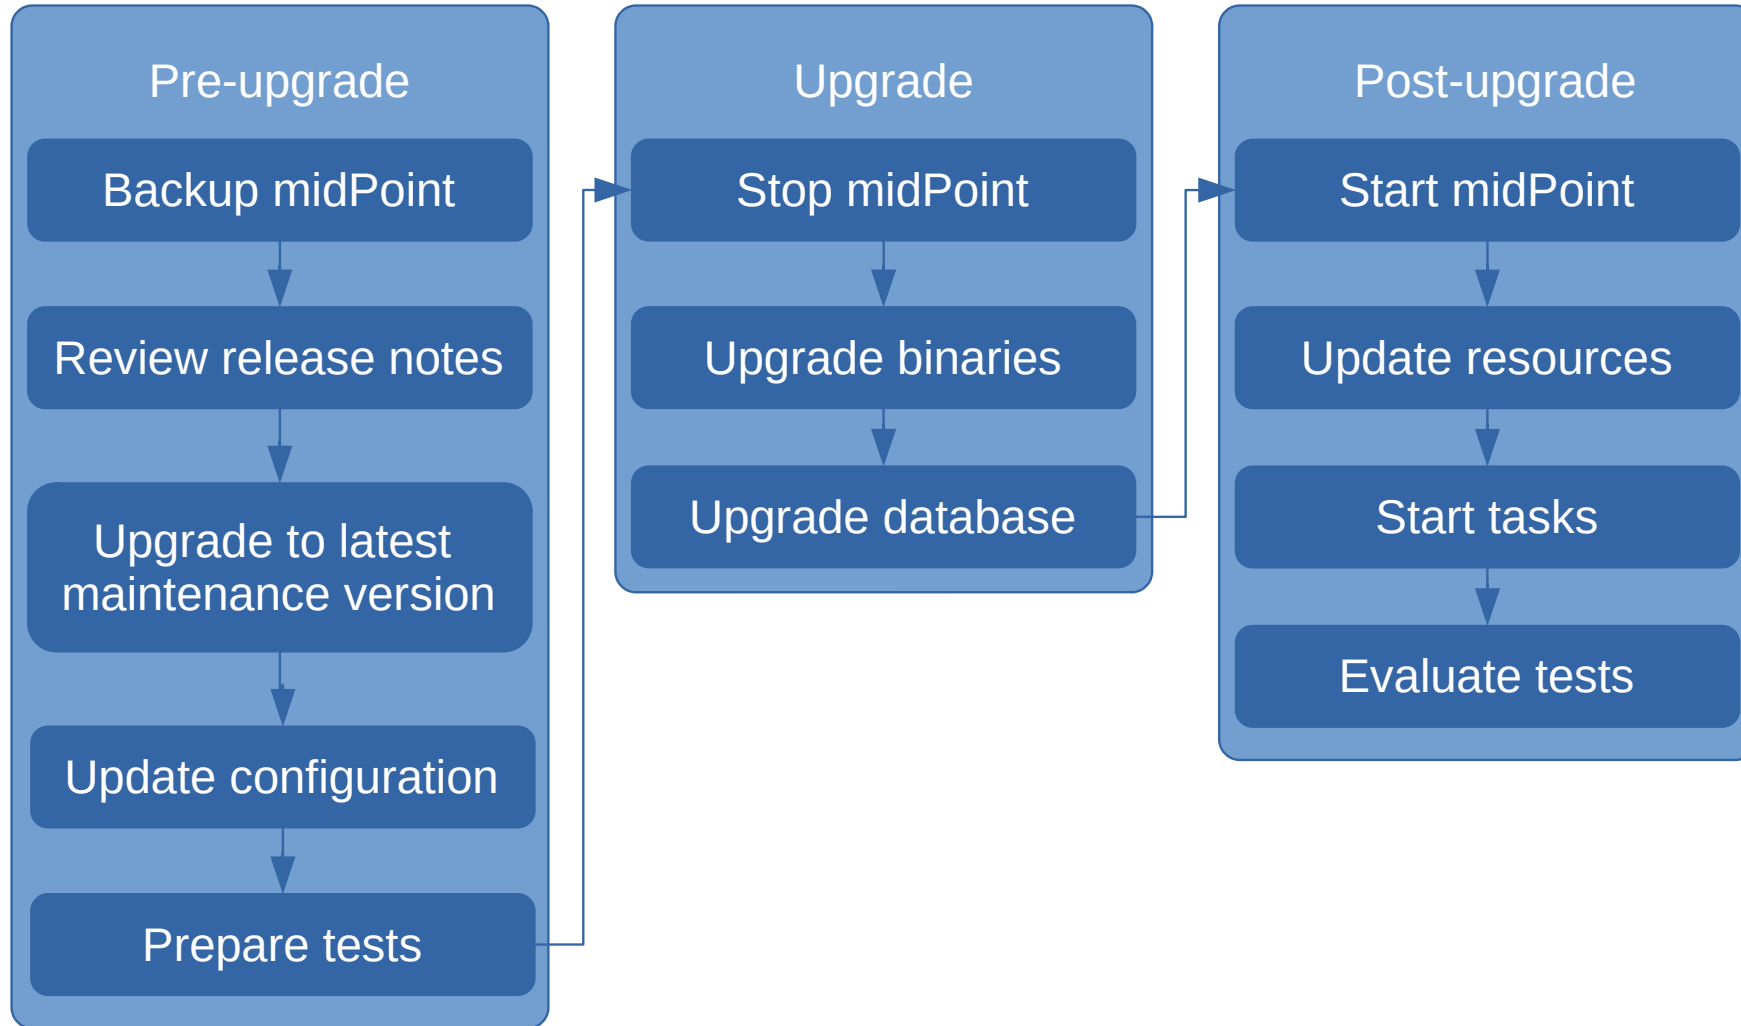


How to upgrade to midPoint 4.4

| How to upgrade to midPoint 4.4

- Supported upgrade paths are:
 - midPoint 4.0.4 to midPoint 4.4.1
 - midPoint 4.3.2 to midPoint 4.4.1
 - midPoint 4.4.1 with generic repository to midPoint 4.4.1 PostgreSQL Native repository
- Upgrade process may require changes in your tasks, mappings, etc. so it is strongly recommended to perform upgrade in testing / staging environment first to get a good estimate of changes required and time needed.

Upgrade Process Overview



| Pre-upgrade steps

- Download 4.4.1 release

<https://evolveum.com/download/4-4-lts-tesla/>

- Create fresh backup of your midPoint and database
- Review release notes for 4.4.1
 - Pay attention to changes, upgrade sections and deprecated / removed features
- Upgrade to latest maintenance version
 - Latest maintenance version contains pre-upgrade checks (e.g. deprecations)

| Pre-upgrade steps

- All paths are relative to midPoint installation (home) directory
 - Usually this is /opt/midpoint
- Use ninja tool to check for deprecated / removed functionality

```
bin/ninja.sh verify -w plannedRemoval
```
- If you are using functionality planned for removal adjust system configuration to not use that functionality (if possible)
 - Data in removed elements are not accessible after upgrade, if you need to remap them to other properties, this needs to happen before upgrade.


```
tony@tdm-box: /opt/midpoint
tony@tdm-box:/opt/midpoint$ bin/ninja.sh verify -w plannedRemoval
Initializing repository using midpoint home
Starting verify
Processed: 0, error: 0, skipped: 0, avg: ,0obj/s
WARNING user:5c61ff17-5a99-412f-b389-05ce09da6054(tonyadm) employeeType planned for removal in version 4.4
WARNING task:00000000-0000-0000-0000-000000000006(Validity Scanner) workState planned for removal in version 4.4
WARNING task:00000000-0000-0000-0000-000000000007(Trigger Scanner) workState planned for removal in version 4.4
WARNING node:e8a977c5-6dd8-4c33-9505-653ffb5e3002(DefaultNode) jmxPort planned for removal in version 4.4
WARNING node:e8a977c5-6dd8-4c33-9505-653ffb5e3002(DefaultNode) operationalStatus planned for removal in version 4.4
WARNING node:e8a977c5-6dd8-4c33-9505-653ffb5e3002(DefaultNode) running planned for removal in version 4.4
WARNING report:00000000-0000-0000-0000-000000000009(Audit logs report) reportEngine planned for removal in version 4.4
WARNING report:00000000-0000-0000-0000-000000000100(Reconciliation report) reportEngine planned for removal in version 4.4
WARNING report:00000000-0000-0000-0000-000000000110(Users in MidPoint) reportEngine planned for removal in version 4.4
WARNING report:00000000-0000-0000-0000-000000000130(Certification definitions report) reportEngine planned for removal in version 4.4
WARNING report:00000000-0000-0000-0000-000000000140(Certification campaigns report) reportEngine planned for removal in version 4.4
WARNING report:00000000-0000-0000-0000-000000000150(Certification cases report) reportEngine planned for removal in version 4.4
WARNING report:00000000-0000-0000-0000-000000000160(Certification decisions report) reportEngine planned for removal in version 4.4
WARNING objectCollection:00000000-0000-0000-0001-000000000002(Modification audit records) auditSearch planned for removal in version 4.4
WARNING objectCollection:00000000-0000-0000-0001-000000000004(Error audit records) auditSearch planned for removal in version 4.4
Finished verify in 2,77s. Processed: 399, error: 0, skipped: 0, avg: 99,75obj/s
tony@tdm-box:/opt/midpoint$
```

- Some of removed schema elements
 - **objectType, employeeType, roleType, orgType and serviceType**. The functionality was replaced by **subtype** property, which was later replaced by archetypes.
 - **subtype** property still remains, however it is deprecated. It is strongly recommended migrating all object subtyping functionality to archetypes.
- More exhaustive list is present in release notes

| Stop the system

- Suspend all scheduled and automatic tasks
 - this prevents any unexpected behavior which may happen after upgrade
- Shutdown midPoint server processes
 - in cluster deployment all nodes must be shut down

| Upgrade binaries

- Usual location of midPoint installation is **/opt/midpoint**
- Replace binaries with new binaries from distribution package
 - **bin** and **lib** directories must be replaced
 - it is strongly recommended to upgrade **doc** and **README**
- **var** directory should be left untouched, since it contains deployment specific files
- Make sure all cluster nodes contains same upgraded binaries

| Upgrade binaries

```
tony@tdm-box: /opt/midpoint
tony@tdm-box:/opt/midpoint$ mkdir backup
tony@tdm-box:/opt/midpoint$ ls
backup/  doc/      lib/      LICENSE.AL  NEWS      README    tmp/
bin/     INSTALL  LICENSE   LICENSE.EUPL  NOTICE   RELEASE-NOTES  var/
tony@tdm-box:/opt/midpoint$ mv -t backup bin lib doc README RELEASE-NOTES
tony@tdm-box:/opt/midpoint$ ls
backup  INSTALL  LICENSE  LICENSE.AL  LICENSE.EUPL  NEWS  NOTICE  tmp  var
tony@tdm-box:/opt/midpoint$ mv -t . tmp/midpoint-4.4.1/{bin,lib,doc,README,RELEASE-NOTES}
tony@tdm-box:/opt/midpoint$ ls
backup  bin  doc  INSTALL  lib  LICENSE  LICENSE.AL  LICENSE.EUPL  NEWS  NOTICE  README  RELEASE-NOTES  tmp  var
tony@tdm-box:/opt/midpoint$
```

- Upgrade scripts are located in **doc/config/sql/generic-old**
 - Naming pattern is database type, original version – target version
(eg. postgresql-upgrade-4.0-4.4.sql)
- Use appropriate tool for your database to run these scripts

PostgreSQL:

```
psql --host=localhost --username=midpoint  
midpoint -f doc/config/sql/postgresql-upgrade-  
4.0-4.4.sql
```

| Upgrade database structure

```
tony@tdm-box: /opt/midpoint/doc/config/sql/generic-old
tony@tdm-box:/opt/midpoint$ cd doc/config/sql/generic-old/
tony@tdm-box:/opt/midpoint/doc/config/sql/generic-old$ ls
h2-upgrade-4.0-4.4.sql      mysql-upgrade-4.0-4.4-utf8mb4.sql  oracle-upgrade-4.3-4.4.sql      postgresql-4.4-all.sql
h2-upgrade-4.3-4.4.sql      mysql-upgrade-4.3-4.4.sql          oracle-4.4-all.sql            sqlserver-upgrade-4.0-4.4.sql
h2-4.4-all.sql            mysql-upgrade-4.3-4.4-utf8mb4.sql  postgresql-upgrade-4.0-4.4.sql  sqlserver-upgrade-4.3-4.4.sql
mysql-upgrade-4.0-4.4.sql  oracle-upgrade-4.0-4.4.sql         postgresql-upgrade-4.3-4.4.sql  sqlserver-4.4-all.sql
tony@tdm-box:/opt/midpoint/doc/config/sql/generic-old$ psql --host=localhost --username=midpoint webinar40 -f postgresql-upgrade-4.0-4.4.sql
Password for user midpoint: █
```

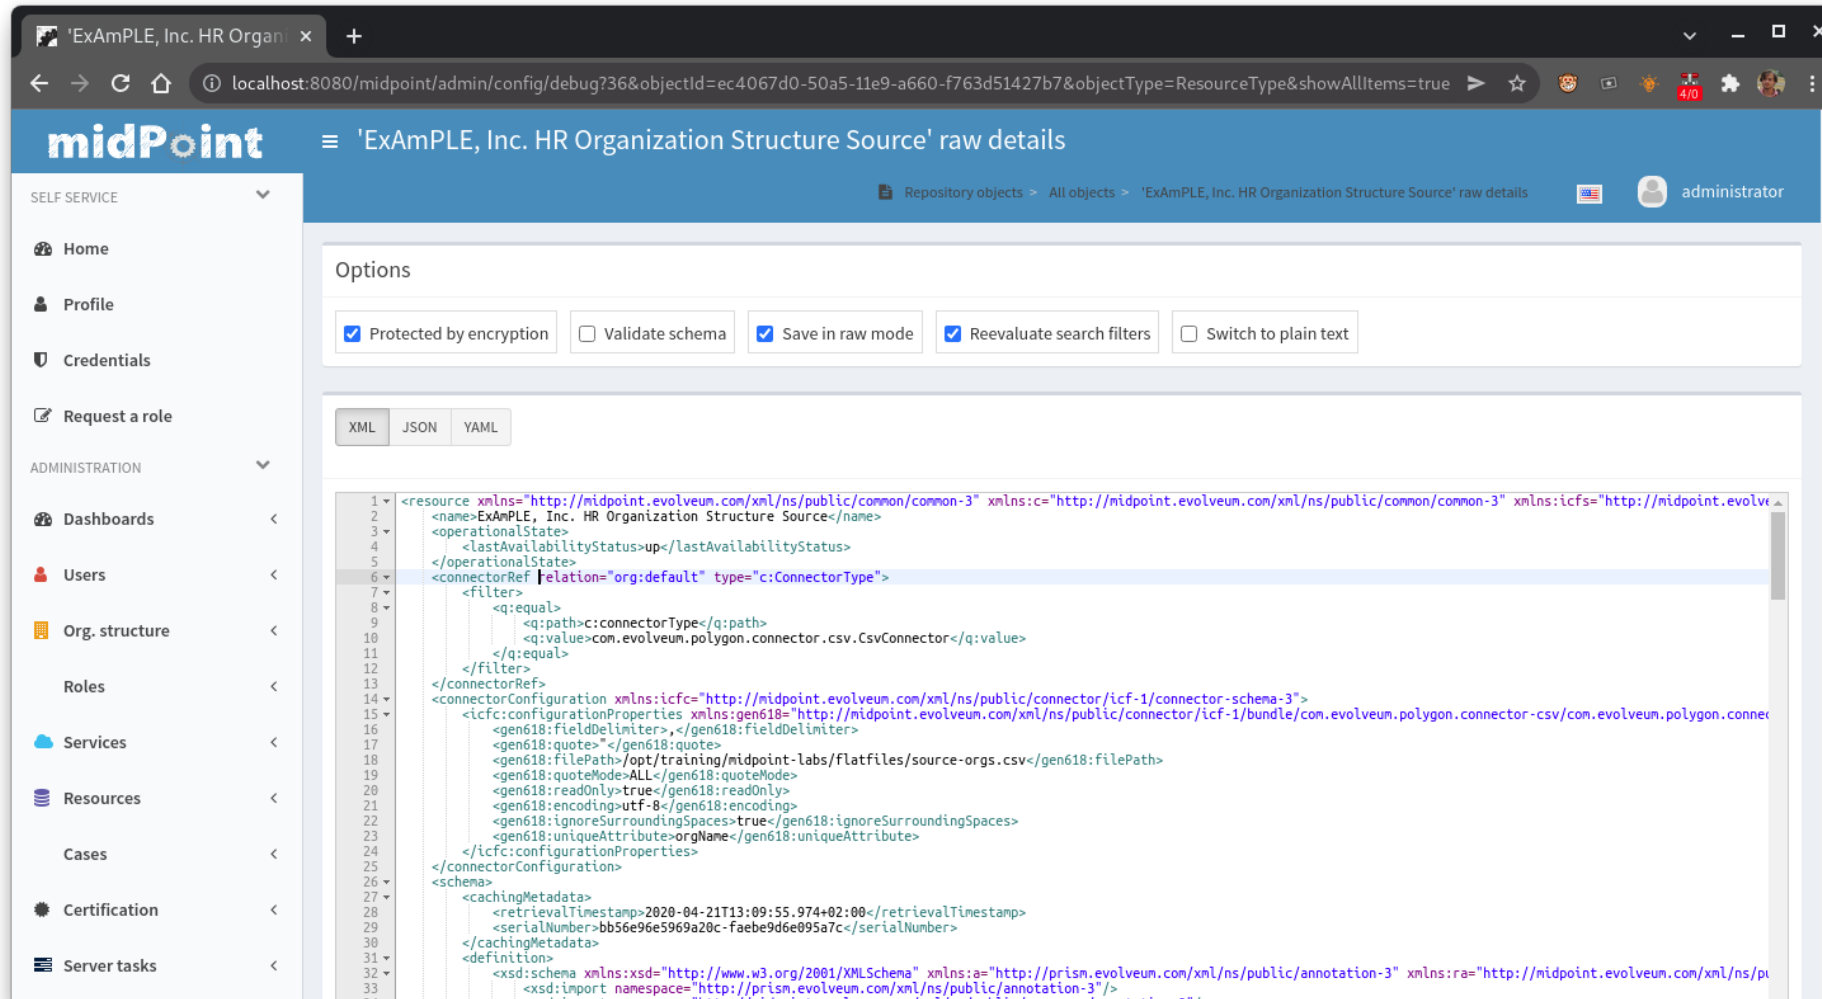
| Post-upgrade steps

- Start upgraded midPoint and monitor midPoint logs for any errors
- 4.4.1 upgrade provides newer versions of connector, so resources are probably not working
- In 4.4.1 some of the initial objects changed, but these changes are not applied automatically if you have an older version

| Post-upgrade steps

- Upgrade Resources to point to new connector versions
 - delete old connector objects in midPoint repository
- Review changes in initial objects and default system configuration
 - Initial objects are located in doc/config/initial-objects
- If your upgrade path was upgrade to 4.4, switch to native repository, then upgrade to 4.4.1 with Native repository, you **need to run reindex task** to ensure shadow attributes indexes are constructed properly.

Upgrade connectors to point to new connector versions



The screenshot shows the midPoint web interface. The browser address bar indicates the URL: `localhost:8080/midpoint/admin/config/debug?36&objectId=ec4067d0-50a5-11e9-a660-f763d51427b7&objectType=ResourceType&showAllItems=true`. The page title is "'ExAmPLE, Inc. HR Organization Structure Source' raw details".

The interface includes a sidebar with navigation options under 'SELF SERVICE' (Home, Profile, Credentials, Request a role) and 'ADMINISTRATION' (Dashboards, Users, Org. structure, Roles, Services, Resources, Cases, Certification, Server tasks). The main content area is titled 'Options' and contains several checkboxes: 'Protected by encryption' (checked), 'Validate schema' (unchecked), 'Save in raw mode' (checked), 'Reevaluate search filters' (checked), and 'Switch to plain text' (unchecked).

Below the options, there are tabs for 'XML', 'JSON', and 'YAML'. The XML view is selected, showing the raw configuration for a connector. The XML snippet includes:

```
1 <resource xmlns="http://midpoint.evolveum.com/xml/ns/public/common/common-3" xmlns:c="http://midpoint.evolveum.com/xml/ns/public/common/common-3" xmlns:icfs="http://midpoint.evolveum.com/xml/ns/public/common/common-3">
2   <name>ExAmPLE, Inc. HR Organization Structure Source</name>
3   <operationalState>
4     <lastAvailabilityStatus>up</lastAvailabilityStatus>
5   </operationalState>
6   <connectorRef relation="org:default" type="c:ConnectorType">
7     <filter>
8       <q:equal>
9         <q:path>c:connectorType</q:path>
10        <q:value>com.evolveum.polygon.connector.csv.CsvConnector</q:value>
11      </q:equal>
12    </filter>
13  </connectorRef>
14  <connectorConfiguration xmlns:icfc="http://midpoint.evolveum.com/xml/ns/public/connector/icf-1/connector-schema-3">
15    <icfc:configurationProperties xmlns:gen618="http://midpoint.evolveum.com/xml/ns/public/connector/icf-1/bundle/com.evolveum.polygon.connector-csv/com.evolveum.polygon.connector-csv">
16      <gen618:fieldDelimiter></gen618:fieldDelimiter>
17      <gen618:quote>"</gen618:quote>
18      <gen618:filePath>/opt/training/midpoint-labs/flatfiles/source-orgs.csv</gen618:filePath>
19      <gen618:quoteMode>ALL</gen618:quoteMode>
20      <gen618:readOnly>true</gen618:readOnly>
21      <gen618:encoding>utf-8</gen618:encoding>
22      <gen618:ignoreSurroundingSpaces>true</gen618:ignoreSurroundingSpaces>
23      <gen618:uniqueAttribute>orgName</gen618:uniqueAttribute>
24    </icfc:configurationProperties>
25  </connectorConfiguration>
26  <schema>
27    < cachingMetadata>
28      < retrievalTimestamp>2020-04-21T13:09:55.974+02:00</retrievalTimestamp>
29      < serialNumber>bb56e96e5969a20c-faeb9d6e095a7c</serialNumber>
30    </ cachingMetadata>
31    < definition>
32      < xsd:schema xmlns:xsd="http://www.w3.org/2001/XMLSchema" xmlns:a="http://prism.evolveum.com/xml/ns/public/annotation-3" xmlns:ra="http://midpoint.evolveum.com/xml/ns/public/annotation-3">
33        < xsd:import namespace="http://prism.evolveum.com/xml/ns/public/annotation-3"/>
34        < xsd:import namespace="http://midpoint.evolveum.com/xml/ns/public/resource/annotation-3"/>
```

Review changes in existing imported initial objects

```
tony@tdm-box: /opt/midpoint
tony@tdm-box: /opt/midpoint$ ls doc/
config/ samples/ schema/
tony@tdm-box: /opt/midpoint$ ls doc/config/
config-native.xml checkstyle/ initial-objects/ sql/
tony@tdm-box: /opt/midpoint$ ls doc/config/initial-objects/
000-system-configuration.xml          130-report-certification-definitions.xml
010-value-policy.xml                  140-report-certification-campaigns.xml
015-security-policy.xml                150-report-certification-cases.xml
020-archetype-system-user.xml          160-report-certification-decisions.xml
021-archetype-system-role.xml          200-lookup-languages.xml
022-archetype-business-role.xml        210-lookup-locales.xml
023-archetype-manual-provisioning-case.xml 220-lookup-timezones.xml
024-archetype-operation-request.xml     230-lookup-lifecycle-state.xml
025-archetype-approval-case.xml        240-lookup-state.xml
026-archetype-trace.xml                 250-object-collection-resource-all.xml
030-role-superuser.xml                 251-object-collection-resource-up.xml
040-role-enduser.xml                   260-object-collection-task-all.xml
041-role-approver.xml                  261-object-collection-task-active.xml
042-role-reviewer.xml                  262-object-collection-task-report.xml
043-role-delegator.xml                 270-object-collection-audit.xml
050-user-administrator.xml             271-object-collection-audit-24h.xml
059-archetype-report.xml               272-object-collection-audit-errors-24h.xml
060-archetype-report-dashboard.xml     273-object-collection-audit-modifications-24h.xml
061-archetype-report-collection.xml    280-object-collection-certification-campaign-all.xml
090-report-audit.xml                   290-object-collection-shadow-all.xml
100-report-reconciliation.xml           310-dashboard-admin.xml
110-report-user-list.xml               330-object-collection-my-cases.xml
501-archetype-task-reconciliation.xml  502-archetype-task-recomputation.xml
503-archetype-task-import.xml          504-archetype-task-live-sync.xml
505-archetype-task-async-update.xml    506-archetype-task-cleanup.xml
507-archetype-task-report.xml          508-archetype-task-single-bulk-action.xml
509-archetype-task-iterative-bulk-action.xml 510-archetype-task-report-import-classic.xml
511-archetype-task-report-export-classic.xml 511-archetype-task-report-export-distributed.xml
520-archetype-task-certification.xml    521-archetype-task-approval.xml
528-archetype-task-util.xml            529-archetype-task-system.xml
550-task-cleanup.xml                  560-task-validity.xml
570-task-trigger.xml                  600-origin-internal.xml
tony@tdm-box: /opt/midpoint$
```

| Update of server tasks

- Deleting working state of server tasks is recommended
- Update server tasks definitions if necessary
 - Apply server task archetypes if missing
 - Configuration of advanced distribution features was changed and needs to be migrated
 - e.g. buckets, worker tasks, and task partitions
 - See <https://docs.evolveum.com/midpoint/reference/tasks/activities/migration/> and midPoint Studio, which provides tooling for tasks conversion.

| Post-upgrade testing

- Before testing we recommend to restart system
- Evaluate updated system and run your post-upgrade tests
 - Monitor logs for possible errors
- If there are no critical errors, your system is ready for use.

| Follow-up steps

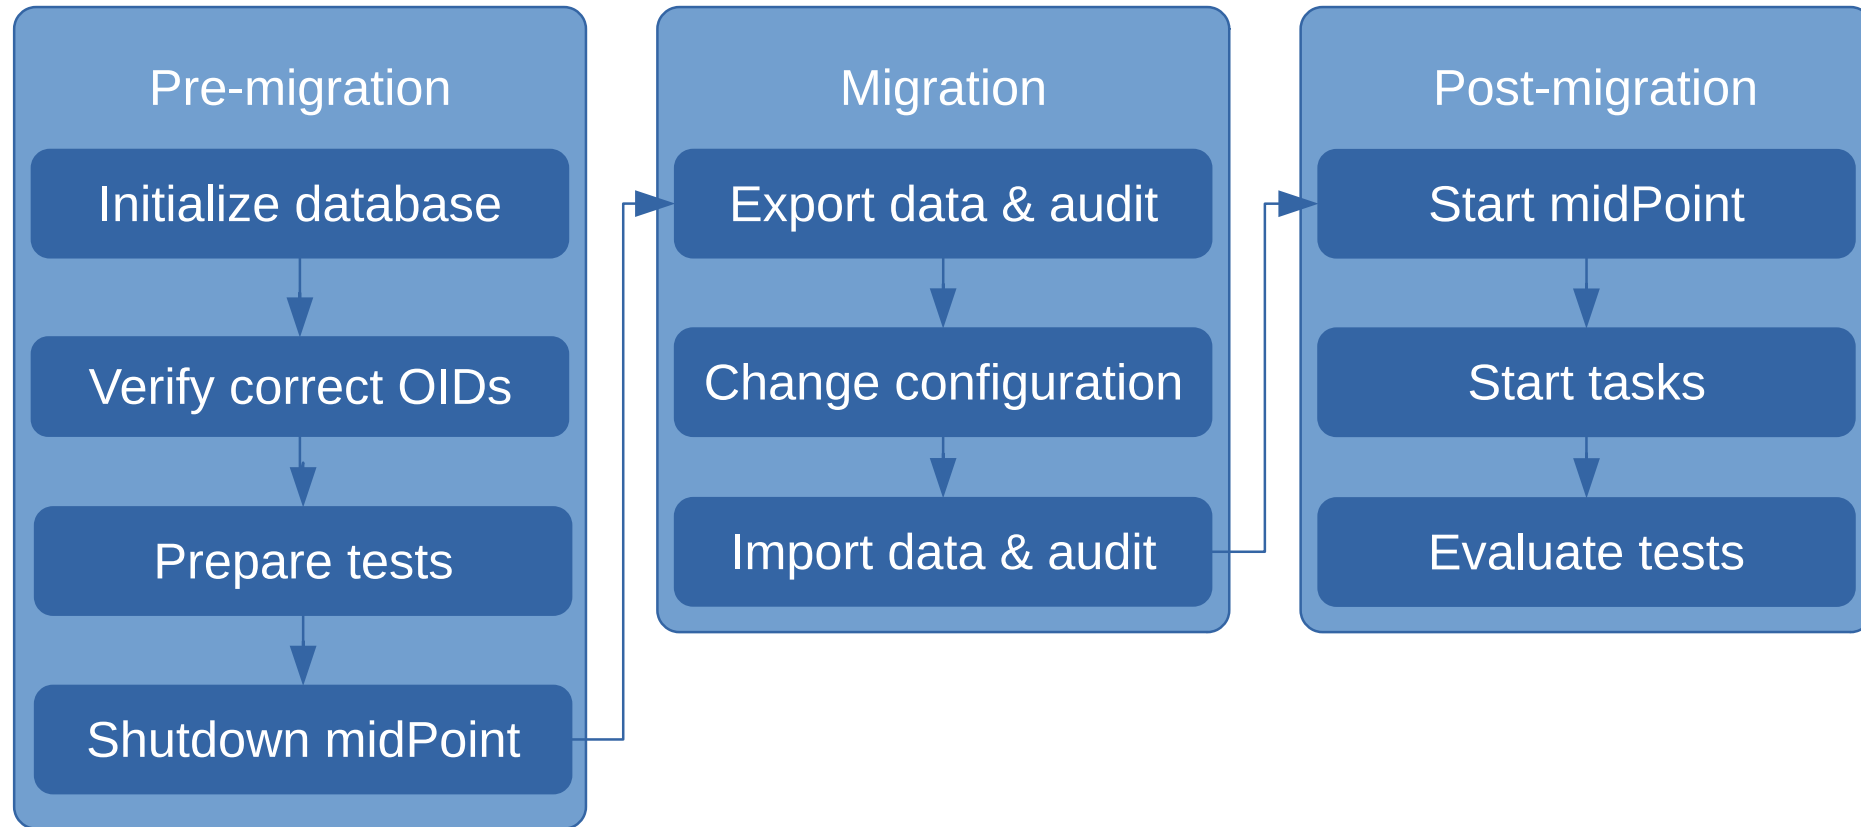
- Review release notes again and use **ninja verify** to determine which functionality is deprecated and you may need to plan migration to newer mechanisms.
- Consider and plan for migration to PostgreSQL native repository.

How to migrate to PostgreSQL Native repository

| Migration to PostgreSQL Native Repository

- Prerequisites
 - midPoint updated to 4.4.1
 - PostgreSQL database of version 13 and later
- Familiarize yourself with Native Repository design
- Decide if midPoint and midPoint audit database is same or will be two different databases
- Repository migration process should be treated similar as upgrade process with separate testing in test environment, planned downtime

Migration to PostgreSQL Native Repository



| Initialize Native PostgreSQL database

- Database schema and creation scripts are present in **doc/config/sql/native-new/**
- For midPoint database use following scripts:
 - postgres-new.sql
 - postgres-new-quartz.sql
- For midPoint audit database (could be same as midPoint DB) use following scripts
 - postgres-new-audit.sql

| Initialize Native PostgreSQL database

```
tony@tdm-box: /opt/midpoint/doc/config/sql/native-new
tony@tdm-box: /opt/midpoint$ cd doc/config/sql/native-new/
tony@tdm-box: /opt/midpoint/doc/config/sql/native-new$ psql --host=localhost --username=midpoint webinar44 -f postgres-new.sql -f postgres-new-quartz.sql -f
postgres-new-upgrade.sql
Password for user midpoint:
tony@tdm-box: /opt/midpoint/doc/config/sql/native-new$ psql --host=localhost --username=midpoint webinar44 -f postgres-new-audit.sql -f postgres-new-upgrade-
audit.sql
```

| Verify that only UUIDs are used

- Native repository requires oids to be valid UUIDs
 - objects generated by midPoint always use UUIDs, only manually created objects may use invalid UUIDs
- Use ninja verify for detecting incorrect oids
 - `bin/ninja.sh verify -w uuid`**
- midPoint does not allow for oid modifications, you need to export objects with broken oids, remove them from midPoint, fix oids to be UUIDs and reimport them
- You also need to update references to use new correct oids

| Verify that only UUIDs are used

```
tony@tdm-box: /opt/midpoint
tony@tdm-box: /opt/midpoint$ bin/ninja.sh verify -w uuid
Initializing using midpoint home; with repository connection
Starting verify
Processed: 0, error: 0, skipped: 0, avg processed: 0obj/s
WARNING user:tony(tony-0ID)  OID 'tony' is not valid UUID
Finished verify in 2,84s. Processed: 422, error: 0, skipped: 0, avg processed: 148,64obj/s
tony@tdm-box: /opt/midpoint$
```

| Shutting down midPoint

- It's recommended to suspend all server tasks and shutdown midPoint instances before migration / exporting data to simplify process
- It is possible to migrate midPoint to new repository with old instances running, but you need to carefully plan migration and its execution.
 - Exporting data could be done while midPoint is running, but you may need reexport any changes, that happened after initial export
 - midPoint with old repository and midPoint with new repository using same data should never be running at same time

| Exporting existing data

- Migration does not support direct database operations and must be done using `ninja`.
- MidPoint data needs to be exported using **ninja** to file system
- Depending on your deployment complete export may take few minutes to few hours and may result in file size of several GB
- Audit data needs to be exported using **ninja** to file system

```
bin/ninja.sh -v export -r -z -0 export.zip
```

```
bin/ninja.sh -v exportAudit -z -0 audit-export.zip
```

| Exporting existing data

```
tony@tdm-box: /opt/midpoint
tony@tdm-box:/opt/midpoint$ bin/ninja.sh -v export -r -z -0 export.zip
2022-01-18 19:23:59,563 [main] INFO (SYSOUT): Initializing using midpoint home; with repository connection
2022-01-18 19:24:05,429 [main] INFO (SYSOUT): Starting export
2022-01-18 19:24:05,429 [main] DEBUG (SYSOUT): Operation: started
2022-01-18 19:24:07,452 [pool-1-thread-2] INFO (SYSOUT): Processed: 0, error: 0, skipped: 0, avg processed: 0obj/s
2022-01-18 19:24:08,397 [pool-1-thread-1] DEBUG (SYSOUT): Operation: producer finished
2022-01-18 19:24:08,475 [pool-1-thread-3] DEBUG (SYSOUT): Operation: finished
2022-01-18 19:24:09,455 [main] INFO (SYSOUT): Finished export in 3,05s. Processed: 421, error: 0, skipped: 0, avg processed: 138,21obj/s
tony@tdm-box:/opt/midpoint$
tony@tdm-box:/opt/midpoint$ bin/ninja.sh -v exportAudit -z -0 audit-export.zip
2022-01-18 19:24:26,333 [main] INFO (SYSOUT): Initializing using midpoint home; with repository connection
2022-01-18 19:24:32,119 [main] INFO (SYSOUT): Starting exportAudit
2022-01-18 19:24:32,120 [main] DEBUG (SYSOUT): Operation: started
2022-01-18 19:24:34,149 [pool-1-thread-2] INFO (SYSOUT): Processed: 0, error: 0, skipped: 0, avg processed: 0obj/s
2022-01-18 19:24:34,384 [pool-1-thread-1] DEBUG (SYSOUT): Operation: producer finished
2022-01-18 19:24:34,539 [pool-1-thread-3] DEBUG (SYSOUT): Operation: finished
2022-01-18 19:24:36,152 [main] INFO (SYSOUT): Finished exportAudit in 2,42s. Processed: 217, error: 0, skipped: 0, avg processed: 89,71obj/s
tony@tdm-box:/opt/midpoint$
```


| Changing midPoint configuration

- Shutdown your midPoint instance
- Modify var/config.xml to enable new native repository and configure its connection to PostgreSQL
 - Replace repository and audit sections with configuration for new native repository
 - See **doc/config/config-native.xml** for quick reference and <https://docs.evolveum.com/midpoint/reference/repository/native-postgresql/> for more details
- **Do not start midPoint after configuration change** – it would initialize as fresh install

| Example native repository section

```
<?xml version="1.0"?>
<configuration>
  <midpoint>
    <repository>
      <type>native</type>
      <jdbcUrl>jdbc:postgresql://localhost:5432/midpoint</jdbcUrl>
      <jdbcUsername>midpoint</jdbcUsername>
      <jdbcPassword>password</jdbcPassword>
    </repository>
    <audit>
      <auditService>
        <auditServiceFactoryClass>com.evolveum.midpoint.audit.impl.LoggerAuditServiceFactory</
auditServiceFactoryClass>
      </auditService>
      <auditService>
        <auditServiceFactoryClass>com.evolveum.midpoint.repo.sqlale.audit.SqlaleAuditServiceFactory</
auditServiceFactoryClass>
      </auditService>
    </audit>
  </midpoint>
</configuration>
```

| Importing existing data

- Use `ninja import` to import production data to new repository

```
bin/ninja.sh -v import -l 4 -z -i export.zip
```

- Import will fail fast on schema problems (missing definitions for extensions)

- Observe / review output of command for potential problems

- Objects with invalid UUIDs will not be imported and will be reported

- Use `ninja importAudit` to import audit data to new repository

```
bin/ninja.sh -v importAudit -l 4 -z -i audit-export.zip
```

- Run `ninja import` multithreaded using `-l` option.

Importing existing data (with invalid OIDs)

```
tony@tdm-box: /opt/midpoint
tony@tdm-box:/opt/midpoint$ bin/ninja.sh -v import -z -i export.zip
2022-01-18 19:00:02,853 [main] INFO (SYSOUT): Initializing using midpoint home; with repository connection
2022-01-18 19:00:08,797 [pool-1-thread-1] INFO (SYSOUT): Starting import
2022-01-18 19:00:08,798 [pool-1-thread-1] DEBUG (SYSOUT): Operation: started
2022-01-18 19:00:08,798 [pool-1-thread-1] INFO (SYSOUT): Processing file export.zip.xml
2022-01-18 19:00:10,813 [pool-1-thread-2] INFO (SYSOUT): Processed: 0, error: 0, skipped: 0, avg processed: 0obj/s
2022-01-18 19:00:12,601 [pool-1-thread-3] ERROR (SYSERR): Couldn't add object user:tony(tony-OID), reason: Cannot convert oid 'tony' to UUID
2022-01-18 19:00:12,603 [pool-1-thread-3] ERROR (SYSERR): Exception details
com.evolveum.midpoint.util.exception.SystemException: Cannot convert oid 'tony' to UUID
    at com.evolveum.midpoint.repo.sqlale.SqaleServiceBase.handledGeneralException(SqaleServiceBase.java:88)
    at com.evolveum.midpoint.repo.sqlale.SqaleRepositoryService.addObject(SqaleRepositoryService.java:334)
    at com.evolveum.midpoint.ninja.action.worker.ImportRepositoryConsumerWorker.run(ImportRepositoryConsumerWorker.java:59)
    at java.base/java.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.java:1128)
    at java.base/java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:628)
    at java.base/java.lang.Thread.run(Thread.java:829)
Caused by: java.lang.IllegalArgumentException: Cannot convert oid 'tony' to UUID
    at com.evolveum.midpoint.repo.sqlale.SqaleUtils.oidToUuid(SqaleUtils.java:62)
    at com.evolveum.midpoint.repo.sqlale.qmodel.object.QObjectMapping.toRowObjectWithoutFullObject(QObjectMapping.java:210)
    at com.evolveum.midpoint.repo.sqlale.qmodel.focus.QFocusMapping.toRowObjectWithoutFullObject(QFocusMapping.java:143)
    at com.evolveum.midpoint.repo.sqlale.qmodel.focus.QUserMapping.toRowObjectWithoutFullObject(QUserMapping.java:79)
    at com.evolveum.midpoint.repo.sqlale.qmodel.focus.QUserMapping.toRowObjectWithoutFullObject(QUserMapping.java:23)
    at com.evolveum.midpoint.repo.sqlale.update.AddObjectContext.addObjectWithOid(AddObjectContext.java:106)
    at com.evolveum.midpoint.repo.sqlale.update.AddObjectContext.execute(AddObjectContext.java:78)
    at com.evolveum.midpoint.repo.sqlale.update.AddObjectContext.execute(AddObjectContext.java:56)
    at com.evolveum.midpoint.repo.sqlale.SqaleRepositoryService.executeAddObject(SqaleRepositoryService.java:351)
    at com.evolveum.midpoint.repo.sqlale.SqaleRepositoryService.addObject(SqaleRepositoryService.java:331)
    ... 4 common frames omitted
Caused by: java.lang.IllegalArgumentException: Invalid UUID string: tony
    at java.base/java.util.UUID.fromString(UUID.java:215)
    at com.evolveum.midpoint.repo.sqlale.SqaleUtils.oidToUuid(SqaleUtils.java:60)
    ... 13 common frames omitted
2022-01-18 19:00:12,814 [pool-1-thread-2] INFO (SYSOUT): Processed: 71, error: 1, skipped: 0, avg processed: 17,680obj/s
2022-01-18 19:00:14,814 [pool-1-thread-2] INFO (SYSOUT): Processed: 240, error: 1, skipped: 0, avg processed: 39,890obj/s
2022-01-18 19:00:16,814 [pool-1-thread-2] INFO (SYSOUT): Processed: 293, error: 1, skipped: 0, avg processed: 36,550obj/s
2022-01-18 19:00:17,496 [pool-1-thread-1] DEBUG (SYSOUT): Operation: producer finished
2022-01-18 19:00:18,814 [pool-1-thread-2] INFO (SYSOUT): Processed: 406, error: 1, skipped: 0, avg processed: 40,540obj/s
2022-01-18 19:00:18,977 [pool-1-thread-3] DEBUG (SYSOUT): Operation: finished
2022-01-18 19:00:20,815 [main] INFO (SYSOUT): Import finished in 10,18s. Processed: 421, error: 1, skipped: 0, avg processed: 41,360obj/s
tony@tdm-box:/opt/midpoint$
tony@tdm-box:/opt/midpoint$
tony@tdm-box:/opt/midpoint$
```

| Post-migration steps

- Migration of data to native PostgreSQL repository is finished
- Make sure same configuration changes are applied to all nodes in cluster deployment
- You may start your migrated instance
- You should observe logs and resume all suspended tasks
- Verify that migration happened successfully.

- Audit data may be too large to be migrated at single planned downtime window
- It's possible to migrate audit at later date, but you need to plan for that from beginning of migration
- See Native Audit section in documentation for more information about postponed audit migration

<https://docs.evolveum.com/midpoint/reference/repository/native-audit/#audit-migration-from-other-database>

| Migration performance numbers

- midPoint with 20k users, 40k shadows, 100k audit records
- Test machine
 - Intel i8700 6-core
 - 32GB RAM
 - HDD WD Blue 5400rpm
 - PostgreSQL running on same machine

Operation	Average objects / sec	Total time
Export data	859	73s
Import data (12 threads)	196	320s
Export audit	380	286s
Import audit (12 threads)	309	350s

- midPoint 4.4 Release notes

<https://docs.evolveum.com/midpoint/release/4.4/>

- midPoint Upgrade Guide

<https://docs.evolveum.com/midpoint/reference/upgrade/upgrade-guide/>

- PostgreSQL native repository

<https://docs.evolveum.com/midpoint/reference/repository/native-postgresql/>

- Native PostgreSQL Audit Trail

<https://docs.evolveum.com/midpoint/reference/repository/native-audit/>

- Target:
 - Tens of millions of identities
- Key results:
 - Improved scalability of midPoint
 - Improved visibility, diagnostic and reliability of midPoint
 - Improved performance and user experience of midPoint user interface



This project has received funding from the European Union's Horizon 2020 research and innovation programme under the NGI_TRUST grant agreement no 825618.

- Tasks, February 3, 2022 by Pavol Mederly
- Customizing GUI, February 10, 2022 by Katarína Bolemant
- Native reports, February 16, 2022 by Lukáš Škublík

Thank you for your time!

See other talks at <https://docs.evolveum.com>

Also **follow us** on our social media for further information!



/Evolveum



/Evolveum



/Evolveum



@Evolveum



/Evolveum

Evolveum