MidPoint 4.4 – Tasks

Pavol Mederly / February 2022

Agenda

Evolveum

- Activities (big redesign)
- Simplified task definition
- Improvements
- New features
- Migration

Redesign of Tasks

Evolveum

- Tasks are here since first versions of midPoint (2011)
  - 3.1 (2015): multiple threads
  - **3.8 (2018): multiple nodes (clustering)**
  - 4.0 (2019): thresholds on a single node (lightweight partitioning)

- Not all the features were easy to use, though.
- After gathering enough experience, we were able to redesign the task execution from scratch.

Activity-based Conceptual Model

Evolveum

Activity = work that needs to be done

Simple activity = action + item set

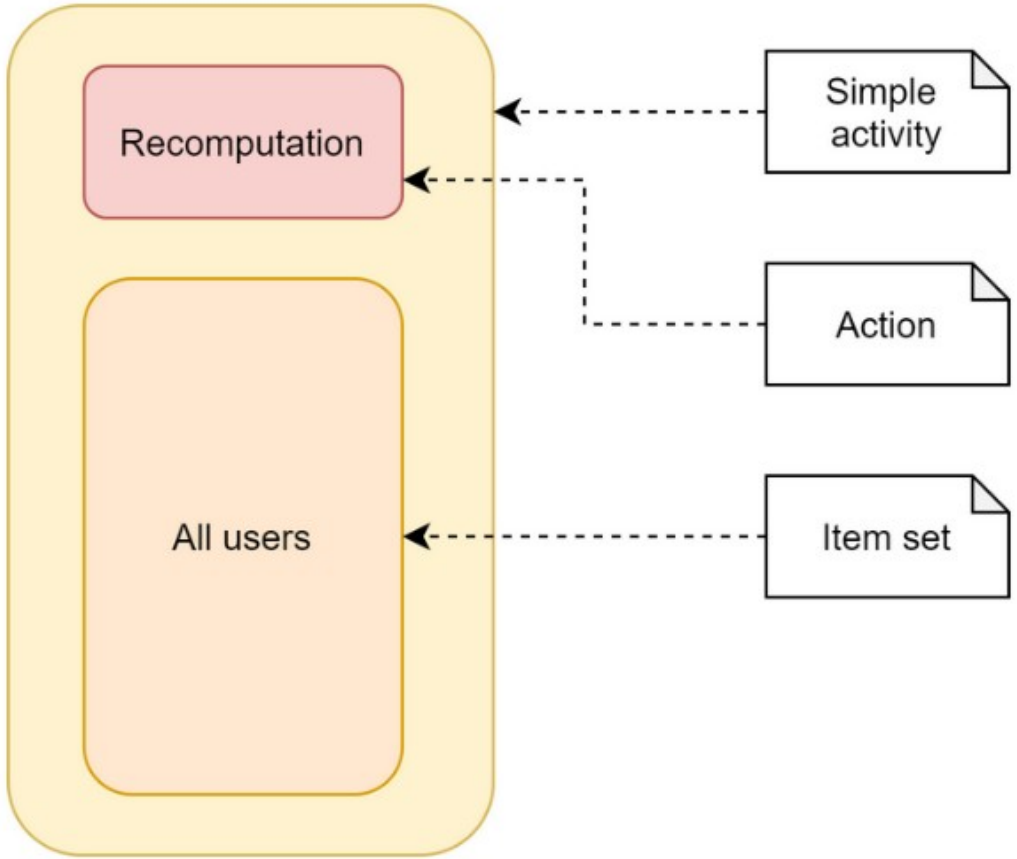

Actions for simple activities

recomputation	triggerScan
import	shadowRefresh
asynchronousUpdate	shadowCleanup
liveSynchronization	reindexing
reportExport	objectIntegrityCheck
reportImport	shadowIntegrityCheck
iterativeScripting	activityAutoScaling
nonIterativeScripting	propagation
changeExecution	multiPropagation
noOp	

Operating on items:

Objects (repo or resource)
Audit records
Live sync events
Async update events
File lines
Integers
...

Composite Activities

Composite activity: contains sub-activities

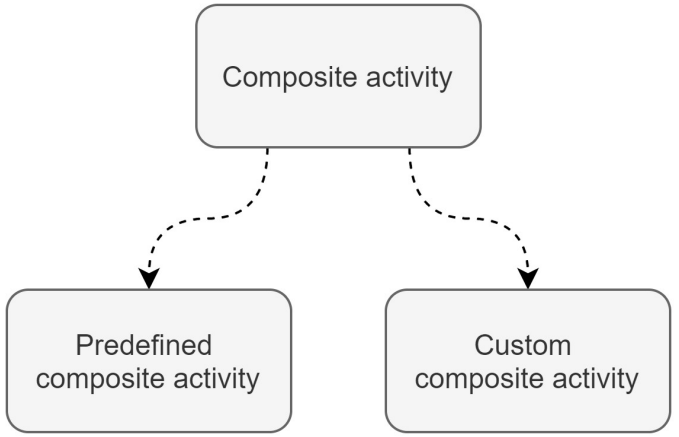

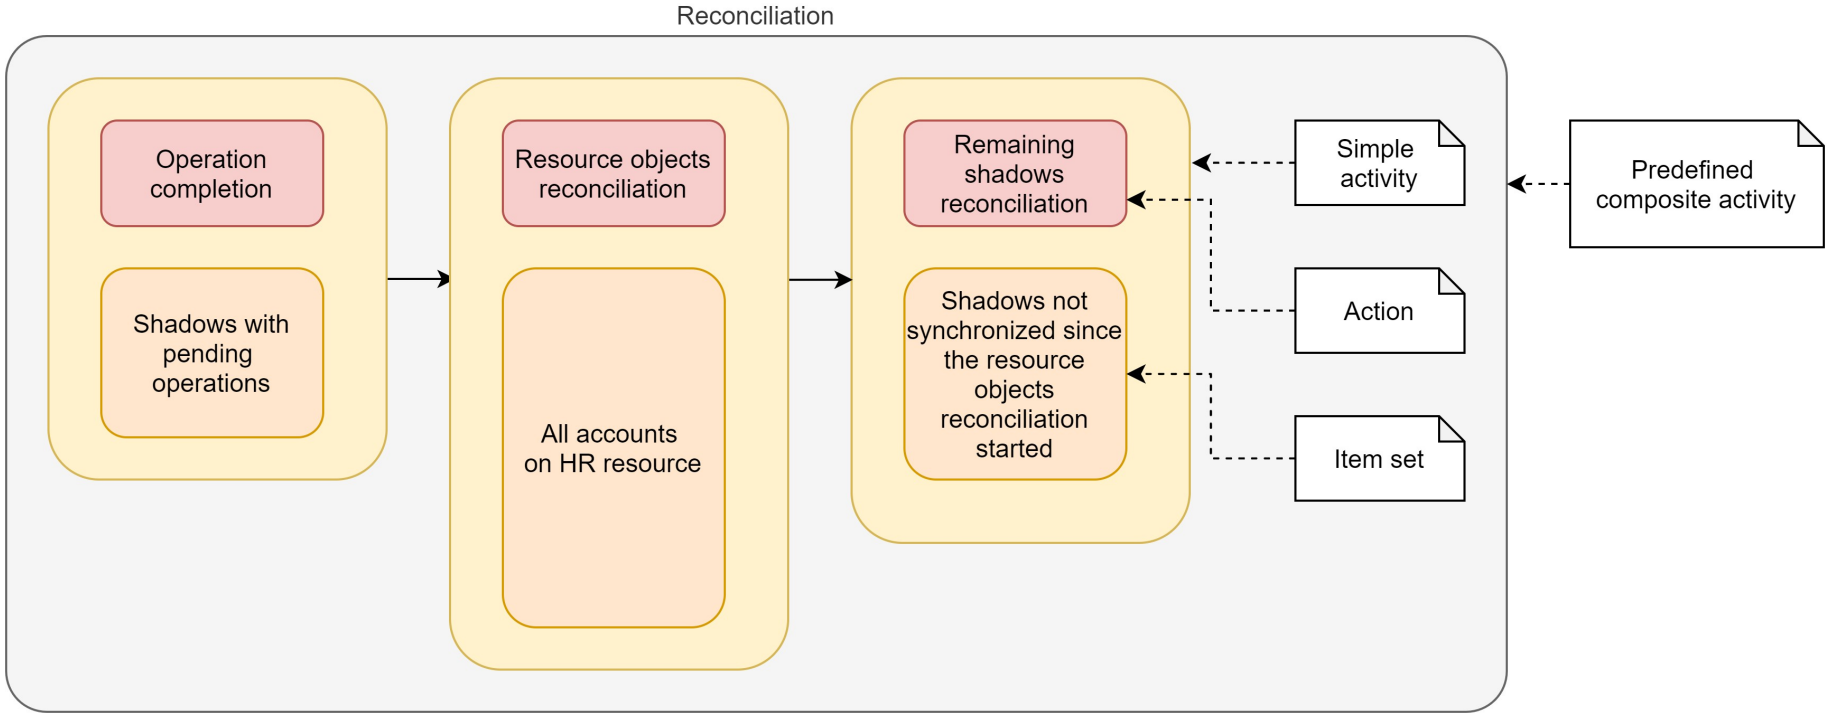

Types of predefined composite activities:

reconciliation
cleanup
distributedReportExport
focusValidityScan

Custom Composite Activities

Evolveum

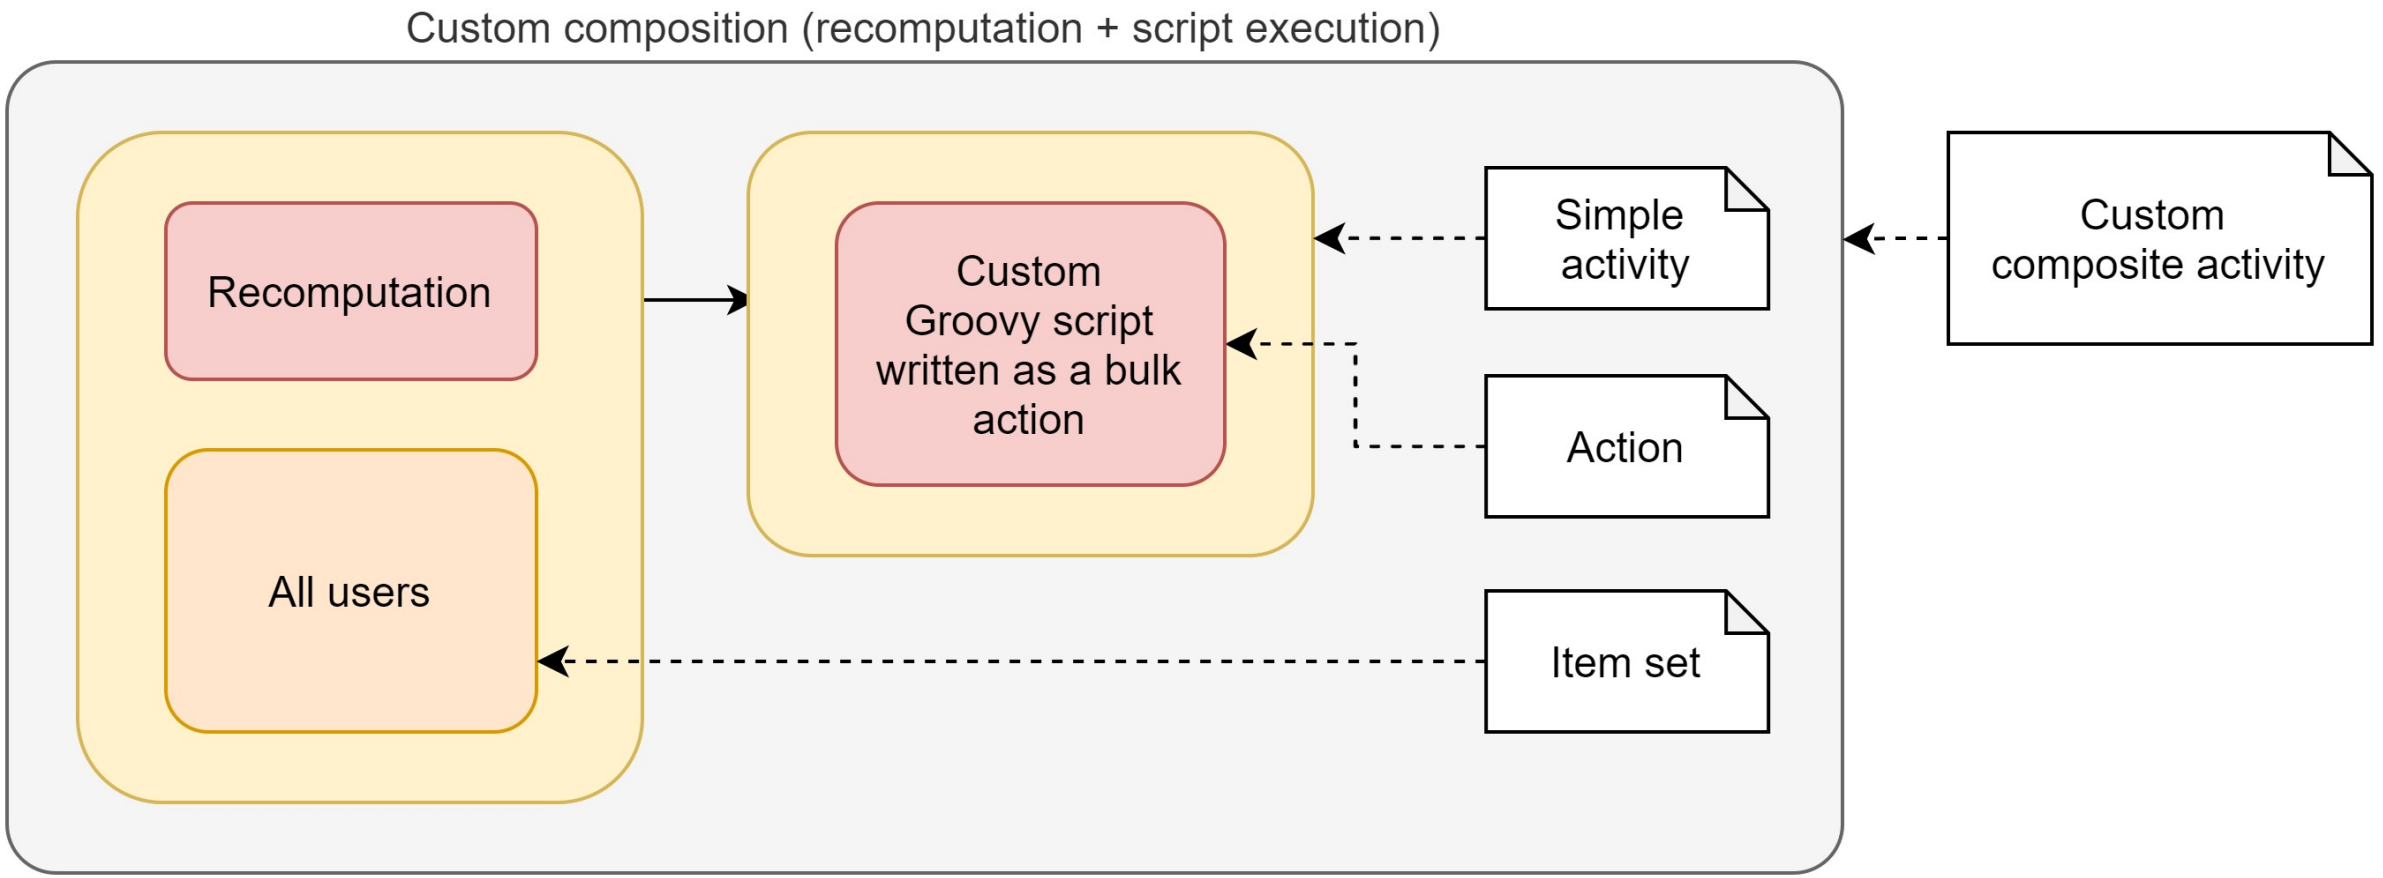

Evolveum

Simplified Task Definition

The Legacy Style

Evolveum

```
<task
    xmlns="http://midpoint.evolveum.com/xml/ns/public/common/common-3"
    xmlns:ext="http://midpoint.evolveum.com/xml/ns/public/model/extension-3">

    <name>Reconciliation</name>
    <extension>
        <ext:kind>account</ext:kind>
        <ext:intent>default</ext:intent>
    </extension>
    <ownerRef oid="a59e4b45-12e7-44e8-bf8c-b51d3a1bbc3e" type="UserType"/>
    <executionState>runnable</executionState>
    <handlerUri>http://midpoint.evolveum.com/xml/ns/public/model/synchronization/task/reconciliation/handler-3</handlerUri>
    <objectRef oid="9125fb57-9f62-4c82-9122-839fa12cf74b" type="ResourceType"/>
</task>
```

- Generic structures: extension, handler URI, object reference
- Dynamic & flexible, but hard to write
  (e.g. what extension items are supported?)

The Clustered Reconciliation Example

Evolveum

```
<task oid="..." xmlns="...">
    ...
    <extension>
        <mext:objectclass>ri:AccountObjectClass</mext:objectclass>
        <mext:workerThreads>4</mext:workerThreads>
    </extension>

    <handlerUri>http://.../partitioned-reconciliation/handler-3</handlerUri>
    <objectRef oid="c845f703-b3c9-..." type="ResourceType" />

    ...
    <workManagement>
        <partitions>
            <partition>
                <index>2</index> <!-- 2 = resource reconciliation -->
                <handlerUri> ... </handlerUri> <!-- optional -->
                <workManagement>
                    <taskKind>coordinator</taskKind>
                    <buckets>
                        <stringSegmentation> ... </stringSegmentation>
                    </buckets>
                    <workers>
                        <workersPerNode> ... </workersPerNode>
                    </workers>
                </workManagement>
            </partition>
        </partitions>
    </workManagement>
</task>
```

Red = what to do
Dark red = on which objects
Blue = work distribution

The New Style

Evolveum

```
<task
    xmlns="http://midpoint.evolveum.com/xml/ns/public/common/common-3">

    <name>Reconciliation</name>
    <ownerRef oid="a59e4b45-12e7-44e8-bf8c-b51d3a1bbc3e" type="UserType"/>
    <executionState>runnable</executionState>
    <activity>
        <work>
            <reconciliation>
                <resourceObjects>
                    <resourceRef oid="9125fb57-9f62-4c82-9122-839fa12cf74b"/>
                    <kind>account</kind>
                    <intent>default</intent>
                </resourceObjects>
            </reconciliation>
        </work>
    </activity>
</task>
```

- Everything in one place
- Statically-typed, yet customer-extensible
- Easy to read, easy to write

Even more important when advanced features (like clustering) are to be configured!

Activity Definition Overview

Evolveum

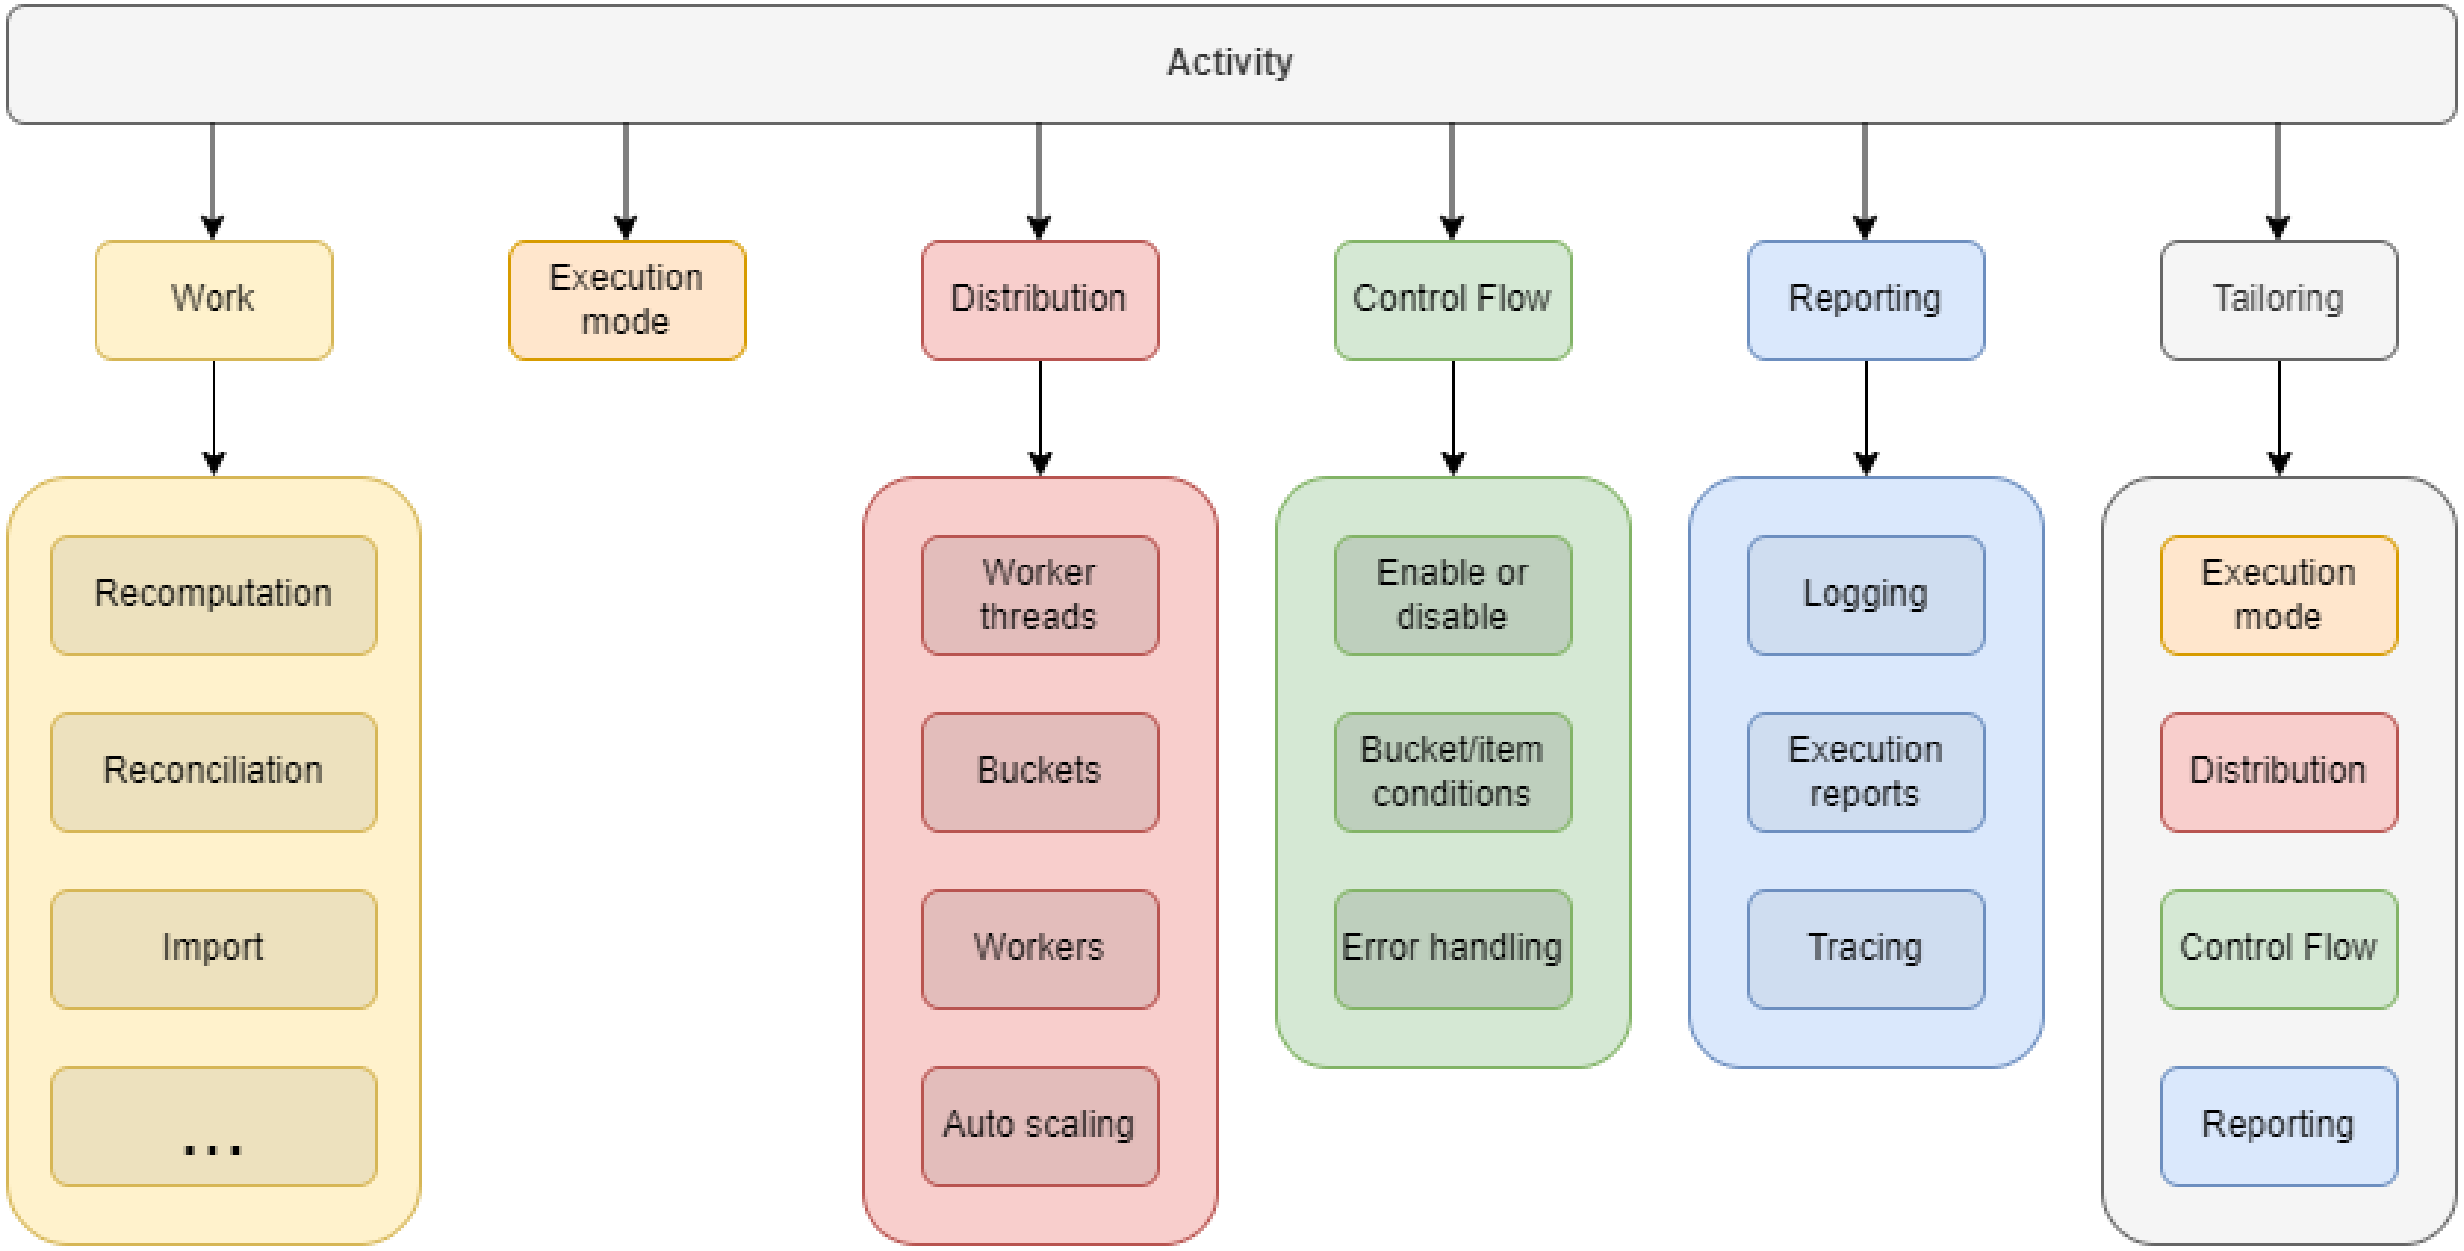

Short Demo

Evolveum

What are Buckets, Anyway?

Evolveum

- A task obtains objects typically through a query.
  - (They are then processed in single or multiple worker threads.)
- If we want to have multiple worker tasks, we need to split the objects set to smaller chunks, called buckets.
- Worker tasks then allocate and process individual buckets.

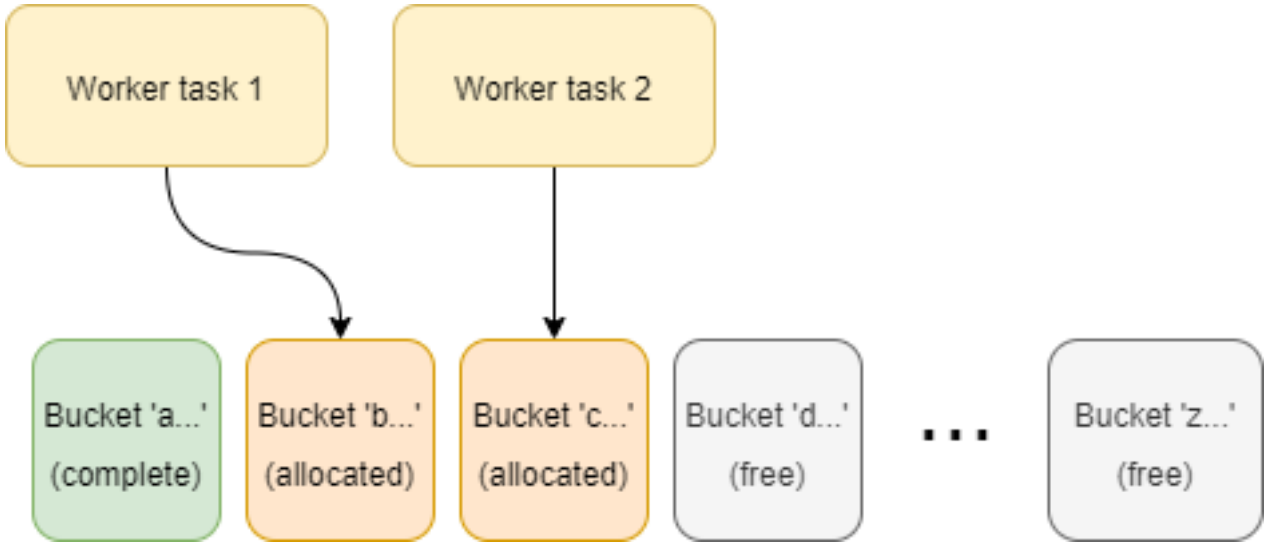


```
<distribution>
    <buckets>
        <stringSegmentation>
            <discriminator>attributes/icfs:name</discriminator>
            <boundary>
                <characters>a-z</characters>
            </boundary>
        </stringSegmentation>
    </buckets>
    <workers>
        <workersPerNode>
            <count>1</count>
        </workersPerNode>
    </workers>
</distribution>
```

Evolveum

Visibility Improvements

Progress, Items Processed, and Throughput

Evolveum

- Progress = how far we have got up to now
- Items processed = how many items we have processed
- Throughput = items processed per time unit

Progress + throughput

⇩

Estimated time to complete (future)

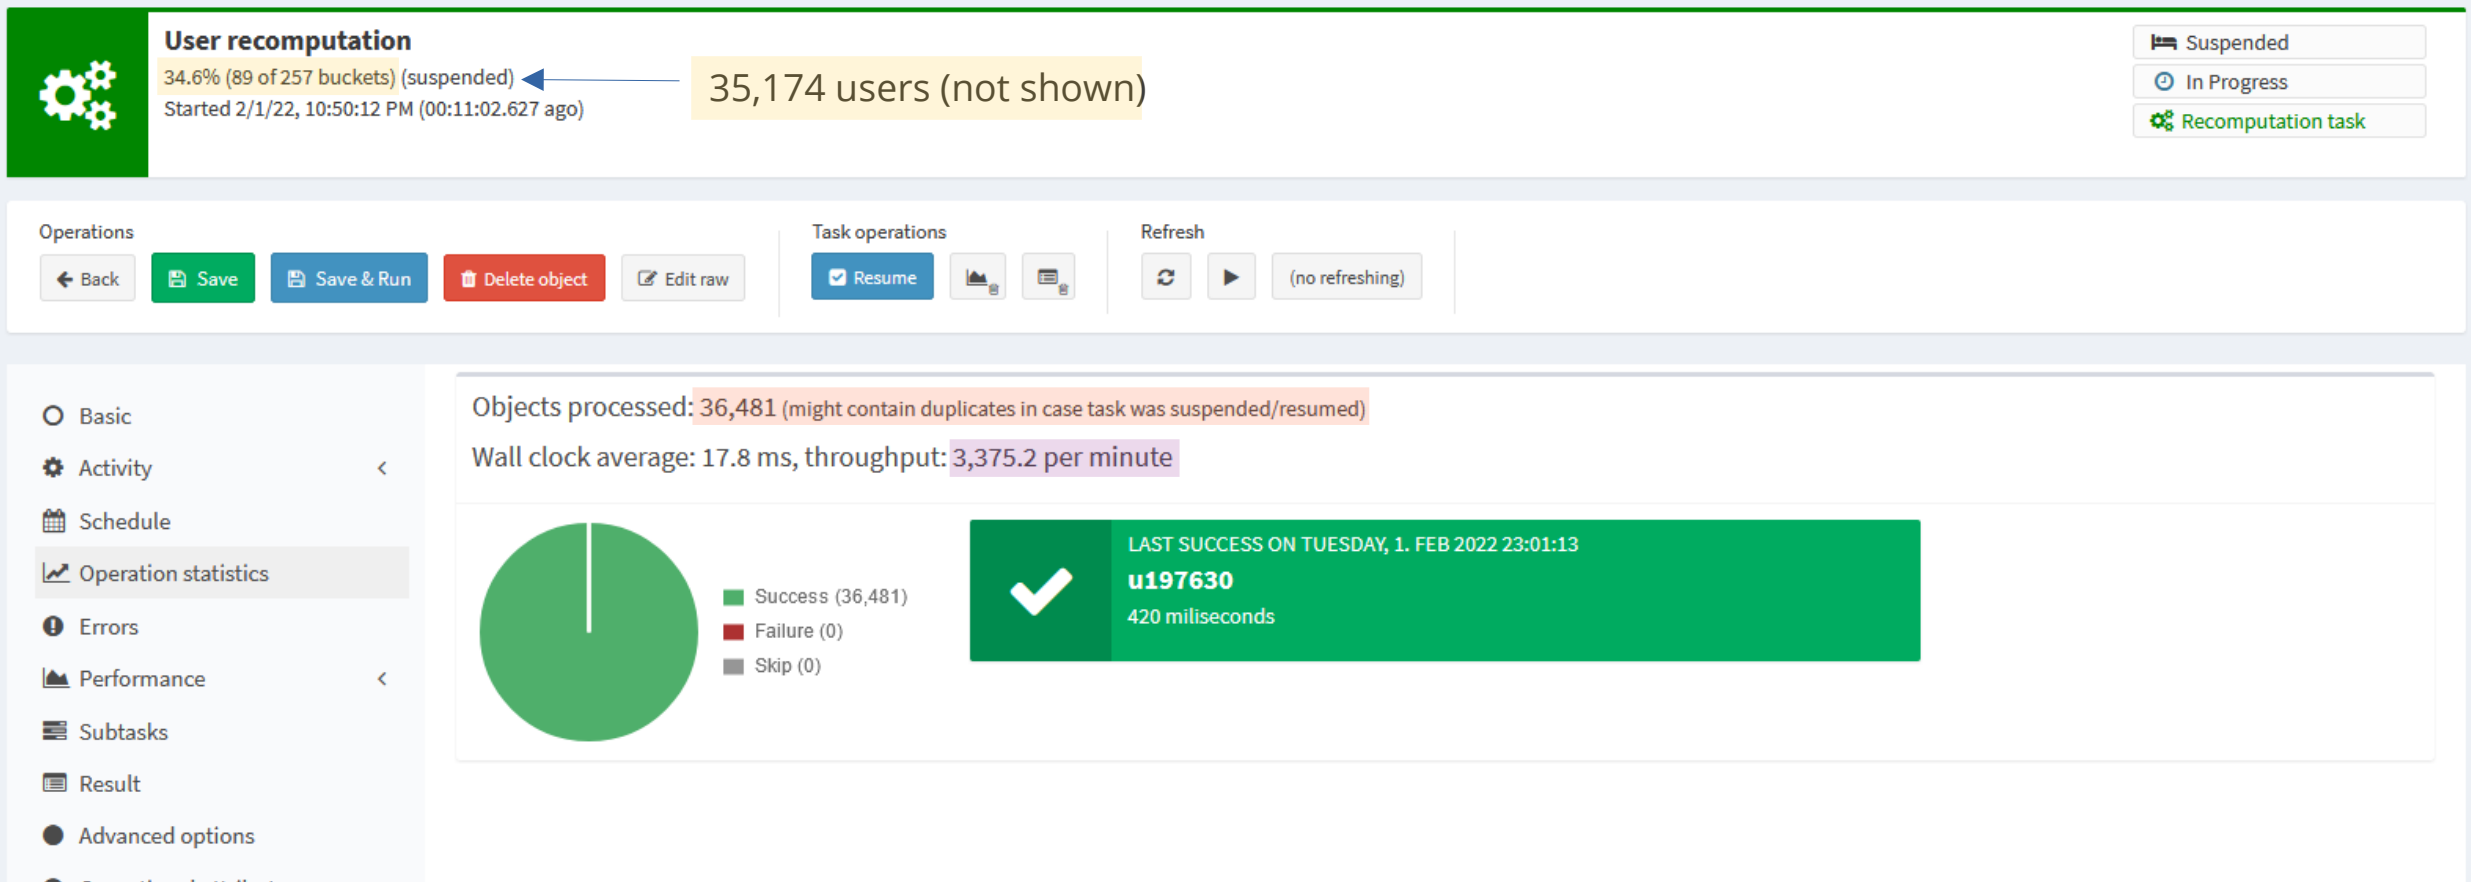

Activity Tree Overview

Evolveum

Key information in the root task – no more need to “click through” to see it

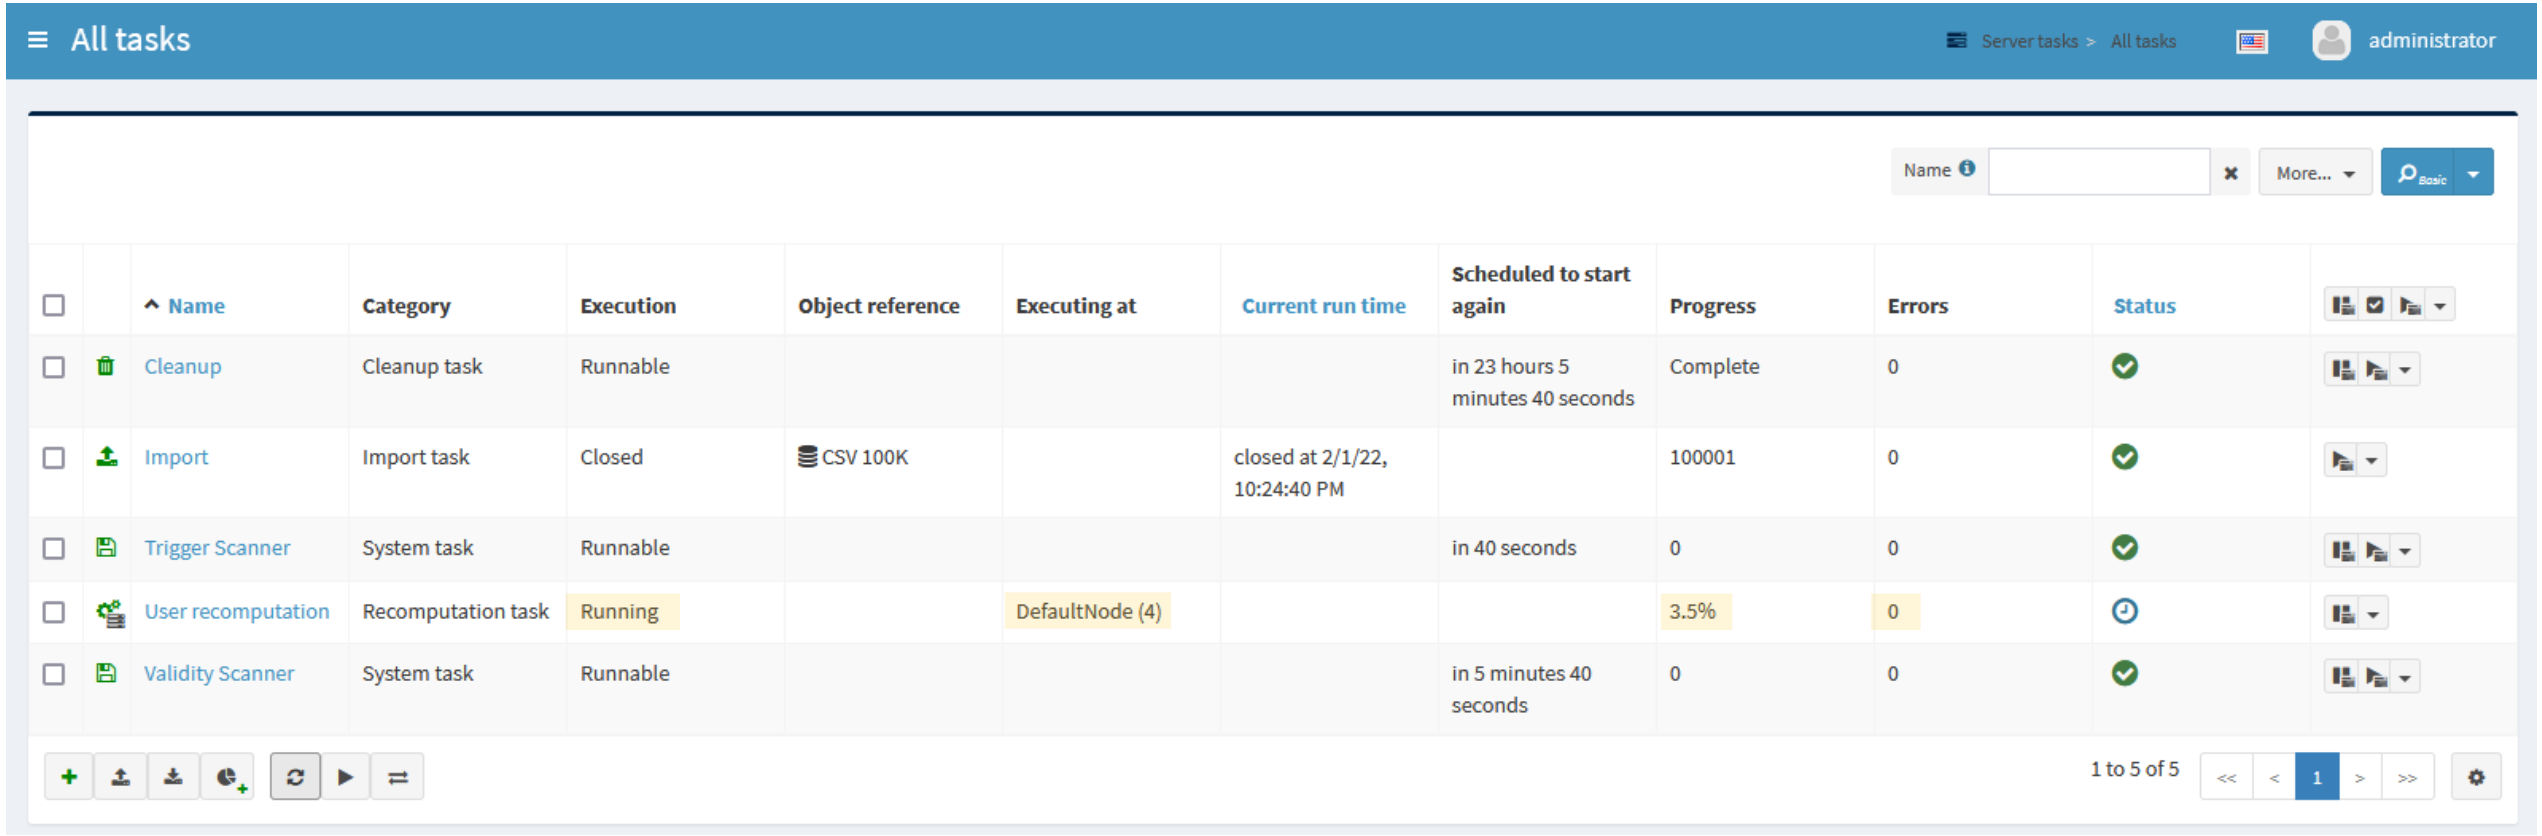

Name	Category	Execution	Object reference	Executing at	Current run time	Scheduled to start again	Progress	Errors	Status
Cleanup	Cleanup task	Runnable				in 23 hours 5 minutes 40 seconds	Complete	0	
Import	Import task	Closed	CSV 100K		closed at 2/1/22, 10:24:40 PM		100001	0	
Trigger Scanner	System task	Runnable				in 40 seconds	0	0	
User recomputation	Recomputation task	Running		DefaultNode (4)			3.5%	0	
Validity Scanner	System task	Runnable				in 5 minutes 40 seconds	0	0	

Evolveum

New Features

Auto-Scaling

Evolveum

- Starting or suspending worker tasks when adding/removing nodes
- Manually or automatically

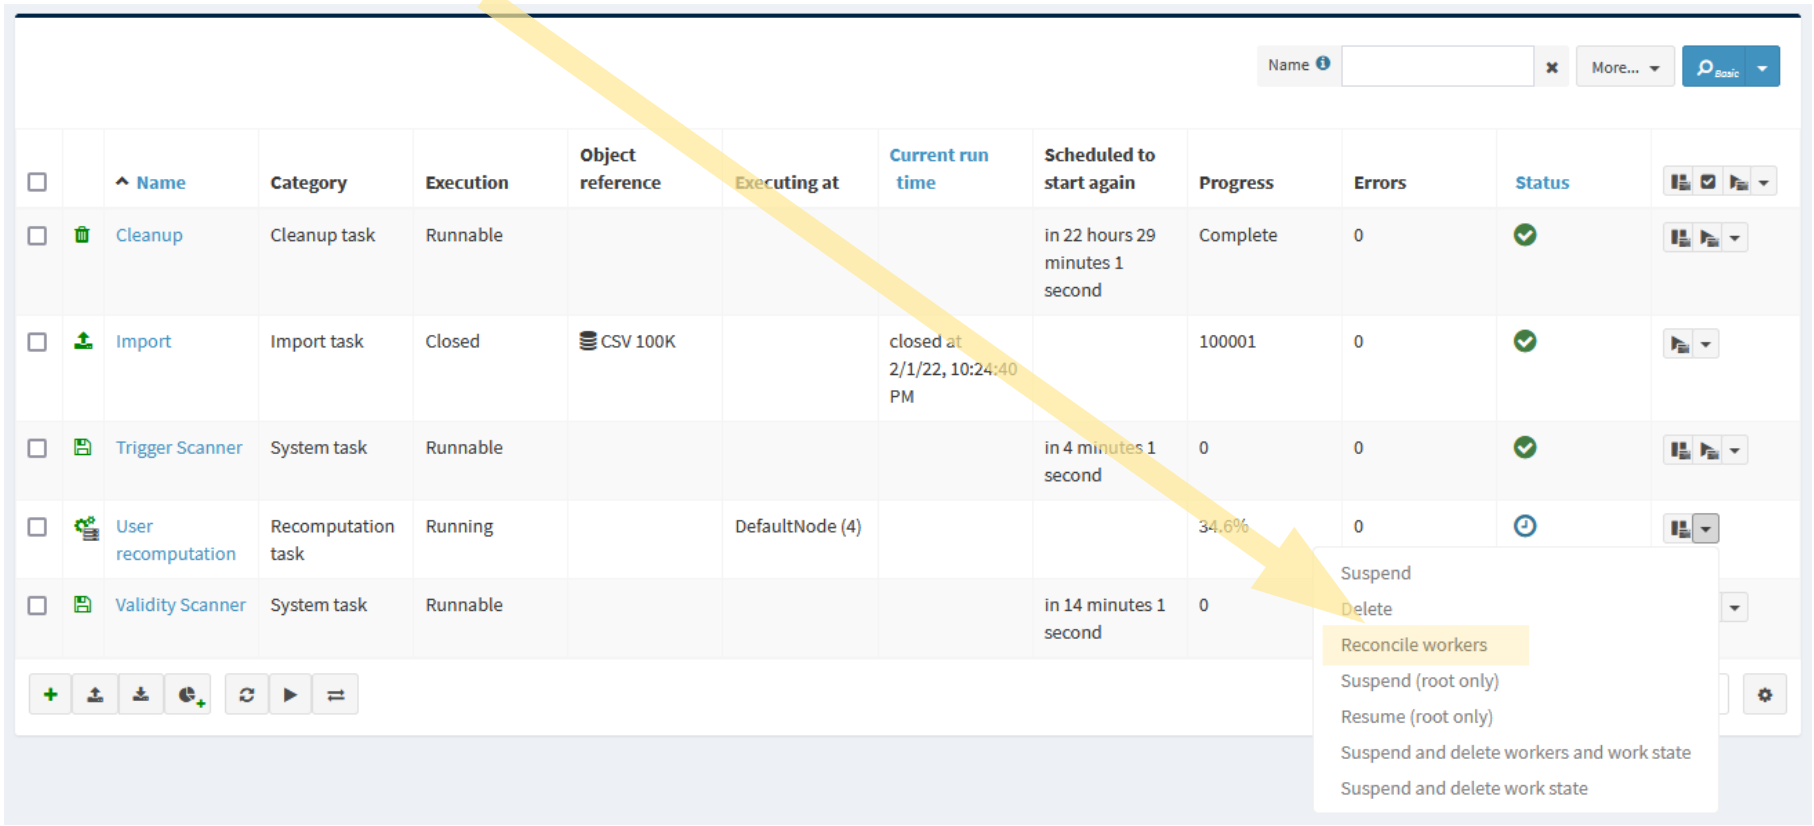


```
<task>
    <name>Auto scaling</name>
    <!— ownerRef and execution state —>
    <schedule>
        <interval>60</interval>
    </schedule>
    <activity>
        <work>
            <activityAutoScaling/>
        </work>
    </activity>
</task>
```

Cluster-wide Thresholds

Evolveum

- Sometimes we want to block unexpected actions before they cause a catastrophe.
- MidPoint allows to declare limits on actions in the form of policy rules.

```xml
<role>
    <name>Stop on delete 10 users</name>
    <inducement>
        <policyRule>
            <policyConstraints>
                <modification>
                    <operation>delete</operation>
                </modification>
            </policyConstraints>
            <policyThreshold>
                <lowWaterMark>
                    <count>10</count>
                </lowWaterMark>
            </policyThreshold>
            <policyActions>
                <suspendTask/>
            </policyActions>
        </policyRule>
        <order>2</order>
    </inducement>
</role>
```

- Roles with such rules are then assigned to tasks.
- After the threshold is reached, a task is suspended.
  - Optionally, a simulation activity can be run right before the execution takes place.
- In midPoint 4.0 the counters were held in memory, so this feature was limited to single node only.
- In 4.4 we maintain the counters directly in the tasks. Therefore, cluster-wide operation is possible.

Execution Reports

Evolveum

- Allows reporting on the execution at the levels of:
  - buckets,
  - items (~ objects),
  - individual ConnId operations,
  - individual internal midPoint operations.
- Easy default configuration, but with great level of flexibility

Buckets Report

Evolveum

#	content	content-from	content-to	size	itemsSuccessfullyProcessed	itemsFailed	itemsSkipped	startTimestamp	startTimestamp-millis	endTimestamp	endTimestamp-millis	duration
1	[5-230004)	5	230004	2935	2927	8	0	2021-10-08T11:00:19.015Z	1633690819015	2021-10-08T11:02:38.528Z	1633690958528	139513
6	[1150000-1379999)	1150000	1379999	2297	2291	6	0	2021-10-08T11:02:38.734Z	1633690958734	2021-10-08T11:04:35.088Z	1633691075088	116354
10	[2069996-2299995)	2069996	2299995	2061	2055	6	0	2021-10-08T11:04:35.239Z	1633691075239	2021-10-08T11:06:17.209Z	1633691177209	101970
14	[2989992-3219991)	2989992	3219991	3747	3734	13	0	2021-10-08T11:06:17.307Z	1633691177307	2021-10-08T11:09:25.689Z	1633691365689	188382
18	[3909988-4139987)	3909988	4139987	1462	1460	2	0	2021-10-08T11:09:25.791Z	1633691365791	2021-10-08T11:10:40.895Z	1633691440895	75104
20	[4369986-4599985)	4369986	4599985	1920	1919	1	0	2021-10-08T11:10:41.018Z	1633691441018	2021-10-08T11:12:15.892Z	1633691535892	94874
23	[5059983-5289982)	5059983	5289982	2426	2419	7	0	2021-10-08T11:12:15.987Z	1633691535987	2021-10-08T11:14:13.690Z	1633691653690	117703
35	[7819971-8049970)	7819971	8049970	505	499	6	0	2021-10-08T11:14:13.805Z	1633691653805	2021-10-08T11:14:38.003Z	1633691678003	24198
38	[8509968-8739967)	8509968	8739967	1468	1461	7	0	2021-10-08T11:14:38.093Z	1633691678093	2021-10-08T11:15:51.241Z	1633691751241	73148
47	[10579959-10809958)	10579959	10809958	17	15	2	0	2021-10-08T11:15:51.406Z	1633691751406	2021-10-08T11:15:52.899Z	1633691752899	1493

```
<reporting>
    <executionReports>
        <buckets/>
    </executionReports>
</reporting>
```

Configurability:

- Row selection (using filter and/or expression)

Items (Objects) Report

Evolveum

#	name	displayName	type	oid	bucket#	outcome	outcome-qualifier	startTimestamp	startTimestamp-millis	endTimestamp	endTimestamp-millis	duration	errorMessage
0	ht	ht (ACCOUNT - default - AccountObjectClass)	{http://midpoint.evolveum.com/xml/ns/public/common/common-3}ShadowType	e010c326-071a-4666-b791-4be5e8480527	1	SUCCESS		2021-10-19T11:50:18.744+02:00	1634637018744	2021-10-19T11:50:18.832+02:00	1634637018832	88.09	
1	guybrush	guybrush (ACCOUNT - default - AccountObjectClass)	{http://midpoint.evolveum.com/xml/ns/public/common/common-3}ShadowType	22226666-2200-6666-6666-444400004444	1	SUCCESS		2021-10-19T11:50:18.917+02:00	1634637018917	2021-10-19T11:50:19.028+02:00	1634637019028	110.6	
2	daviejones	daviejones (ACCOUNT - default - AccountObjectClass)	{http://midpoint.evolveum.com/xml/ns/public/common/common-3}ShadowType	bb6a968d-3a0c-4753-9b58-6c932b1a5245	1	SKIP		2021-10-19T11:50:19.050+02:00	1634637019050	2021-10-19T11:50:19.051+02:00	1634637019051	0.29	
3	calypso	calypso (ACCOUNT - default - AccountObjectClass)	{http://midpoint.evolveum.com/xml/ns/public/common/common-3}ShadowType	a4793302-0624-42bd-9527-98a39ede2621	1	SKIP		2021-10-19T11:50:19.055+02:00	1634637019055	2021-10-19T11:50:19.056+02:00	1634637019056	0.28	
4	elaine	elaine (ACCOUNT - default - AccountObjectClass)	{http://midpoint.evolveum.com/xml/ns/public/common/common-3}ShadowType	c0c010c0-d34d-b33f-f00d-22220004000e	1	SUCCESS		2021-10-19T11:50:19.061+02:00	1634637019061	2021-10-19T11:50:19.191+02:00	1634637019191	129.55	
5	rapp	rapp (ACCOUNT - default - AccountObjectClass)	{http://midpoint.evolveum.com/xml/ns/public/common/common-3}ShadowType	c70312bb-a6fd-48d7-b2ed-303007a2e190	1	SUCCESS		2021-10-19T11:50:19.206+02:00	1634637019206	2021-10-19T11:50:19.331+02:00	1634637019331	125.03	
6	stan	stan (ACCOUNT - default - AccountObjectClass)	{http://midpoint.evolveum.com/xml/ns/public/common/common-3}ShadowType	22220000-2200-0000-0000-444400004455	1	SUCCESS		2021-10-19T11:50:19.348+02:00	1634637019348	2021-10-19T11:50:19.404+02:00	1634637019404	55.81	

ConnId Operations Report

Evolveum

#	itemName	itemOid	bucket#	identifier	resourceRef	resourceRef-name	objectClass	operation	status	message	size	startTimestamp	startTimestamp-millis	endTimestamp	endTimestamp-millis	duration
1	guybrush	22226666-2200-6666-6666-444400004444	1	1634637018962-0-1	10000000-0000-0000-0000-000000000004	Dummy Resource	{http://midpoint.evolveum.com/xml/ns/public/resource/instance-3}AccountObjectClass	ICF_GET				2021-10-19T11:50:18.962+02:00	1634637018962	2021-10-19T11:50:18.962+02:00	1634637018962	0.0
1	guybrush	22226666-2200-6666-6666-444400004444	1	1634637018965-0-1	10000000-0000-0000-0000-000000000004	Dummy Resource	{http://midpoint.evolveum.com/xml/ns/public/resource/instance-3}AccountObjectClass	ICF_UPDATE				2021-10-19T11:50:18.965+02:00	1634637018965	2021-10-19T11:50:18.965+02:00	1634637018965	0.0
1	guybrush	22226666-2200-6666-6666-444400004444	1	1634637019003-0-1	10000000-0000-0000-0000-000000000004	Dummy Resource	{http://midpoint.evolveum.com/xml/ns/public/resource/instance-3}AccountObjectClass	ICF_GET				2021-10-19T11:50:19.003+02:00	1634637019003	2021-10-19T11:50:19.004+02:00	1634637019004	1.0
4	elaine	c0c010c0-d34d-b33f-f00d-22220004000e	1	1634637019094-0-1	10000000-0000-0000-0000-000000000104	Dummy Resource Red	{http://midpoint.evolveum.com/xml/ns/public/resource/instance-3}AccountObjectClass	ICF_GET				2021-10-19T11:50:19.094+02:00	1634637019094	2021-10-19T11:50:19.095+02:00	1634637019095	1.0
4	elaine	c0c010c0-d34d-b33f-f00d-22220004000e	1	1634637019099-0-1	10000000-0000-0000-0000-000000000204	Dummy Resource Blue	{http://midpoint.evolveum.com/xml/ns/public/resource/instance-3}AccountObjectClass	ICF_GET				2021-10-19T11:50:19.099+02:00	1634637019099	2021-10-19T11:50:19.099+02:00	1634637019099	0.0
			1	1634637018739-0-1	10000000-0000-0000-0000-000000000004	Dummy Resource	{http://midpoint.evolveum.com/xml/ns/public/resource/instance-3}AccountObjectClass	ICF_SEARCH				2021-10-19T11:50:18.739+02:00	1634637018739	2021-10-19T11:50:19.414+02:00	1634637019414	0.0

Internal midPoint Operations Report

Evolveum

itemSequentialNumber	itemName	itemOid	operation	qualifiers	status	start	microsecond	params
0	ht	e010c326-071a-4666-b791-4be5e8480527	com.evolveum.midpoint.repo.cache.RepositoryCache.searchObjects	FocusType	SUCCESS	2021-10-19T11:50:18.772+02:00	1939	query=[Q{REF: linkRef, PRV(oid=e010c326-071a-4666-b791-4be5e8480527, targetType=null, relation={http://prism.evolveum.com/
0	ht	e010c326-071a-4666-b791-4be5e8480527	com.evolveum.midpoint.repo.api.RepositoryService.searchObjects	FocusType	SUCCESS	2021-10-19T11:50:18.772+02:00	1704	query=[Q{REF: linkRef, PRV(oid=e010c326-071a-4666-b791-4be5e8480527, targetType=null, relation={http://prism.evolveum.com/
0	ht	e010c326-071a-4666-b791-4be5e8480527	com.evolveum.midpoint.repo.cache.RepositoryCache.modifyObject	ShadowType	SUCCESS	2021-10-19T11:50:18.774+02:00	5267	oid=[e010c326-071a-4666-b791-4be5e8480527]; type=[http://midpoint.evolveum.com/xml/ns/public/common/common-3#ShadowTy
0	ht	e010c326-071a-4666-b791-4be5e8480527	com.evolveum.midpoint.repo.api.RepositoryService.modifyObject	ShadowType	SUCCESS	2021-10-19T11:50:18.774+02:00	5045	oid=[e010c326-071a-4666-b791-4be5e8480527]; type=[com.evolveum.midpoint.xml.ns._public.common.common_3.ShadowType];
0	ht	e010c326-071a-4666-b791-4be5e8480527	com.evolveum.midpoint.repo.cache.RepositoryCache.invalidateCacheEntries		SUCCESS	2021-10-19T11:50:18.779+02:00	95	additionalInfo=[ModifyObjectResult]; oid=[e010c326-071a-4666-b791-4be5e8480527]; type=[http://midpoint.evolveum.com/xml/ns/
0	ht	e010c326-071a-4666-b791-4be5e8480527	com.evolveum.midpoint.model.impl.lens.projector.Projector.focusActivation	INITIAL.e0p0	SUCCESS	2021-10-19T11:50:18.790+02:00	26	
0	ht	e010c326-071a-4666-b791-4be5e8480527	com.evolveum.midpoint.model.impl.lens.projector.Projector.focusActivation	INITIAL.e0p0	SUCCESS	2021-10-19T11:50:18.790+02:00	6	
0	ht	e010c326-071a-4666-b791-4be5e8480527	com.evolveum.midpoint.model.impl.lens.projector.Projector.focusActivation	INITIAL.e0p0	SUCCESS	2021-10-19T11:50:18.791+02:00	7	
0	ht	e010c326-071a-4666-b791-4be5e8480527	com.evolveum.midpoint.model.impl.lens.projector.Projector.focusCredentials	INITIAL.e0p0	SUCCESS	2021-10-19T11:50:18.791+02:00	11	
0	ht	e010c326-071a-4666-b791-4be5e8480527	com.evolveum.midpoint.repo.cache.RepositoryCache.getVersion	ResourceType	SUCCESS	2021-10-19T11:50:18.810+02:00	11	oid=[10000000-0000-0000-0000-000000000004]; type=[http://midpoint.evolveum.com/xml/ns/public/common/common-3#Resource
0	ht	e010c326-071a-4666-b791-4be5e8480527	com.evolveum.midpoint.repo.cache.RepositoryCache.getVersion	ResourceType	SUCCESS	2021-10-19T11:50:18.820+02:00	9	oid=[10000000-0000-0000-0000-000000000004]; type=[http://midpoint.evolveum.com/xml/ns/public/common/common-3#Resource
0	ht	e010c326-071a-4666-b791-4be5e8480527	com.evolveum.midpoint.repo.api.RepositoryService.hasConflict		SUCCESS	2021-10-19T11:50:18.831+02:00	767	watcherClass=[com.evolveum.midpoint.repo.sqlbase.ConflictWatcherImpl]; oid=[7bef09cb-8ee4-4746-8694-0d71eba86489]
0	ht	e010c326-071a-4666-b791-4be5e8480527	com.evolveum.midpoint.repo.api.RepositoryService.getVersion	ObjectType	SUCCESS	2021-10-19T11:50:18.831+02:00	744	oid=[7bef09cb-8ee4-4746-8694-0d71eba86489]; type=[com.evolveum.midpoint.xml.ns._public.common.common_3.ObjectType]
4	elaine	c0c010c0-d34d-b33f-f00d-22220004000e	com.evolveum.midpoint.repo.cache.RepositoryCache.searchObjects	FocusType	SUCCESS	2021-10-19T11:50:19.063+02:00	2277	query=[Q{REF: linkRef, PRV(oid=c0c010c0-d34d-b33f-f00d-22220004000e, targetType=null, relation={http://prism.evolveum.com/x
4	elaine	c0c010c0-d34d-b33f-f00d-22220004000e	com.evolveum.midpoint.repo.api.RepositoryService.searchObjects	FocusType	SUCCESS	2021-10-19T11:50:19.063+02:00	1936	query=[Q{REF: linkRef, PRV(oid=c0c010c0-d34d-b33f-f00d-22220004000e, targetType=null, relation={http://prism.evolveum.com/x
4	elaine	c0c010c0-d34d-b33f-f00d-22220004000e	com.evolveum.midpoint.repo.cache.RepositoryCache.modifyObject	ShadowType	SUCCESS	2021-10-19T11:50:19.066+02:00	4180	oid=[c0c010c0-d34d-b33f-f00d-22220004000e]; type=[http://midpoint.evolveum.com/xml/ns/public/common/common-3#ShadowTy
4	elaine	c0c010c0-d34d-b33f-f00d-22220004000e	com.evolveum.midpoint.repo.api.RepositoryService.modifyObject	ShadowType	SUCCESS	2021-10-19T11:50:19.066+02:00	4001	oid=[c0c010c0-d34d-b33f-f00d-22220004000e]; type=[com.evolveum.midpoint.xml.ns._public.common.common_3.ShadowType]; t
4	elaine	c0c010c0-d34d-b33f-f00d-22220004000e	com.evolveum.midpoint.repo.cache.RepositoryCache.invalidateCacheEntries		SUCCESS	2021-10-19T11:50:19.070+02:00	78	additionalInfo=[ModifyObjectResult]; oid=[c0c010c0-d34d-b33f-f00d-22220004000e]; type=[http://midpoint.evolveum.com/xml/ns/p
4	elaine	c0c010c0-d34d-b33f-f00d-22220004000e	com.evolveum.midpoint.model.impl.lens.projector.Projector.focusActivation	INITIAL.e0p0	SUCCESS	2021-10-19T11:50:19.085+02:00	65	
4	elaine	c0c010c0-d34d-b33f-f00d-22220004000e	com.evolveum.midpoint.model.impl.lens.projector.Projector.focusActivation	INITIAL.e0p0	SUCCESS	2021-10-19T11:50:19.086+02:00	8	
4	elaine	c0c010c0-d34d-b33f-f00d-22220004000e	com.evolveum.midpoint.model.impl.lens.projector.Projector.focusActivation	INITIAL.e0p0	SUCCESS	2021-10-19T11:50:19.087+02:00	37	
4	elaine	c0c010c0-d34d-b33f-f00d-22220004000e	com.evolveum.midpoint.model.impl.lens.projector.Projector.focusCredentials	INITIAL.e0p0	SUCCESS	2021-10-19T11:50:19.087+02:00	12	
4	elaine	c0c010c0-d34d-b33f-f00d-22220004000e	com.evolveum.midpoint.repo.cache.RepositoryCache.getVersion	ResourceType	SUCCESS	2021-10-19T11:50:19.151+02:00	15	oid=[10000000-0000-0000-0000-000000000004]; type=[http://midpoint.evolveum.com/xml/ns/public/common/common-3#Resource
4	elaine	c0c010c0-d34d-b33f-f00d-22220004000e	com.evolveum.midpoint.repo.cache.RepositoryCache.getVersion	ResourceType	SUCCESS	2021-10-19T11:50:19.159+02:00	13	oid=[10000000-0000-0000-0000-000000000004]; type=[http://midpoint.evolveum.com/xml/ns/public/common/common-3#Resource
4	elaine	c0c010c0-d34d-b33f-f00d-22220004000e	com.evolveum.midpoint.repo.api.RepositoryService.hasConflict		SUCCESS	2021-10-19T11:50:19.189+02:00	1017	watcherClass=[com.evolveum.midpoint.repo.sqlbase.ConflictWatcherImpl]; oid=[c0c010c0-d34d-b33f-f00d-11111111111e]
4	elaine	c0c010c0-d34d-b33f-f00d-22220004000e	com.evolveum.midpoint.repo.api.RepositoryService.getVersion	ObjectType	SUCCESS	2021-10-19T11:50:19.190+02:00	968	oid=[c0c010c0-d34d-b33f-f00d-11111111111e]; type=[com.evolveum.midpoint.xml.ns._public.common.common_3.ObjectType]
5	rapp	c70312bb-a6fd-48d7-b2ed-303007a2e190	com.evolveum.midpoint.repo.cache.RepositoryCache.searchObjects	FocusType	SUCCESS	2021-10-19T11:50:19.208+02:00	2157	query=[Q{REF: linkRef, PRV(oid=c70312bb-a6fd-48d7-b2ed-303007a2e190, targetType=null, relation={http://prism.evolveum.com/:
5	rapp	c70312bb-a6fd-48d7-b2ed-303007a2e190	com.evolveum.midpoint.repo.api.RepositoryService.searchObjects	FocusType	SUCCESS	2021-10-19T11:50:19.208+02:00	1918	query=[Q{REF: linkRef, PRV(oid=c70312bb-a6fd-48d7-b2ed-303007a2e190, targetType=null, relation={http://prism.evolveum.com/:
5	rapp	c70312bb-a6fd-48d7-b2ed-303007a2e190	com.evolveum.midpoint.repo.cache.RepositoryCache.modifyObject	ShadowType	SUCCESS	2021-10-19T11:50:19.210+02:00	8219	oid=[c70312bb-a6fd-48d7-b2ed-303007a2e190]; type=[http://midpoint.evolveum.com/xml/ns/public/common/common-3#ShadowTy
5	rapp	c70312bb-a6fd-48d7-b2ed-303007a2e190	com.evolveum.midpoint.repo.api.RepositoryService.modifyObject	ShadowType	SUCCESS	2021-10-19T11:50:19.210+02:00	8032	oid=[c70312bb-a6fd-48d7-b2ed-303007a2e190]; type=[com.evolveum.midpoint.xml.ns._public.common.common_3.ShadowType];
5	rapp	c70312bb-a6fd-48d7-b2ed-303007a2e190	com.evolveum.midpoint.repo.cache.RepositoryCache.invalidateCacheEntries		SUCCESS	2021-10-19T11:50:19.219+02:00	74	additionalInfo=[ModifyObjectResult]; oid=[c70312bb-a6fd-48d7-b2ed-303007a2e190]; type=[http://midpoint.evolveum.com/xml/ns/p
5	rapp	c70312bb-a6fd-48d7-b2ed-303007a2e190	com.evolveum.midpoint.model.impl.lens.projector.Projector.focusActivation	INITIAL.e0p0	SUCCESS	2021-10-19T11:50:19.248+02:00	32	
5	rapp	c70312bb-a6fd-48d7-b2ed-303007a2e190	com.evolveum.midpoint.model.impl.lens.projector.Projector.focusActivation	INITIAL.e0p0	SUCCESS	2021-10-19T11:50:19.248+02:00	6	
5	rapp	c70312bb-a6fd-48d7-b2ed-303007a2e190	com.evolveum.midpoint.model.impl.lens.projector.Projector.focusActivation	INITIAL.e0p0	SUCCESS	2021-10-19T11:50:19.249+02:00	8	
5	rapp	c70312bb-a6fd-48d7-b2ed-303007a2e190	com.evolveum.midpoint.model.impl.lens.projector.Projector.focusCredentials	INITIAL.e0p0	SUCCESS	2021-10-19T11:50:19.249+02:00	14	
5	rapp	c70312bb-a6fd-48d7-b2ed-303007a2e190	com.evolveum.midpoint.repo.cache.RepositoryCache.getVersion	ResourceType	SUCCESS	2021-10-19T11:50:19.292+02:00	14	oid=[10000000-0000-0000-0000-000000000004]; type=[http://midpoint.evolveum.com/xml/ns/public/common/common-3#Resource
5	rapp	c70312bb-a6fd-48d7-b2ed-303007a2e190	com.evolveum.midpoint.repo.cache.RepositoryCache.getVersion	ResourceType	SUCCESS	2021-10-19T11:50:19.301+02:00	12	oid=[10000000-0000-0000-0000-000000000004]; type=[http://midpoint.evolveum.com/xml/ns/public/common/common-3#Resource
5	rapp	c70312bb-a6fd-48d7-b2ed-303007a2e190	com.evolveum.midpoint.repo.api.RepositoryService.hasConflict		SUCCESS	2021-10-19T11:50:19.330+02:00	873	watcherClass=[com.evolveum.midpoint.repo.sqlbase.ConflictWatcherImpl]; oid=[c0c010c0-d34d-b33f-f00d-11111111c008]
5	rapp	c70312bb-a6fd-48d7-b2ed-303007a2e190	com.evolveum.midpoint.repo.api.RepositoryService.getVersion	ObjectType	SUCCESS	2021-10-19T11:50:19.330+02:00	845	oid=[c0c010c0-d34d-b33f-f00d-11111111c008]; type=[com.evolveum.midpoint.xml.ns._public.common.common_3.ObjectType]
6	stan	22220000-2200-0000-0000-444400004455	com.evolveum.midpoint.repo.cache.RepositoryCache.searchObjects	FocusType	SUCCESS	2021-10-19T11:50:19.350+02:00	1651	query=[Q{REF: linkRef, PRV(oid=22220000-2200-0000-0000-444400004455, targetType=null, relation={http://prism.evolveum.com/
6	stan	22220000-2200-0000-0000-444400004455	com.evolveum.midpoint.repo.api.RepositoryService.searchObjects	FocusType	SUCCESS	2021-10-19T11:50:19.350+02:00	1445	query=[Q{REF: linkRef, PRV(oid=22220000-2200-0000-0000-444400004455, targetType=null, relation={http://prism.evolveum.com/
6	stan	22220000-2200-0000-0000-444400004455	com.evolveum.midpoint.repo.cache.RepositoryCache.modifyObject	ShadowType	SUCCESS	2021-10-19T11:50:19.352+02:00	4134	oid=[22220000-2200-0000-0000-444400004455]; type=[http://midpoint.evolveum.com/xml/ns/public/common/common-3#ShadowTy
6	stan	22220000-2200-0000-0000-444400004455	com.evolveum.midpoint.repo.api.RepositoryService.modifyObject	ShadowType	SUCCESS	2021-10-19T11:50:19.352+02:00	3970	oid=[22220000-2200-0000-0000-444400004455]; type=[com.evolveum.midpoint.xml.ns._public.common.common_3.ShadowType];
6	stan	22220000-2200-0000-0000-444400004455	com.evolveum.midpoint.repo.cache.RepositoryCache.invalidateCacheEntries		SUCCESS	2021-10-19T11:50:19.356+02:00	69	additionalInfo=[ModifyObjectResult]; oid=[22220000-2200-0000-0000-444400004455]; type=[http://midpoint.evolveum.com/xml/ns/p
6	stan	22220000-2200-0000-0000-444400004455	com.evolveum.midpoint.model.impl.lens.projector.Projector.focusActivation	INITIAL.e0p0	SUCCESS	2021-10-19T11:50:19.366+02:00	23	
6	stan	22220000-2200-0000-0000-444400004455	com.evolveum.midpoint.model.impl.lens.projector.Projector.focusActivation	INITIAL.e0p0	SUCCESS	2021-10-19T11:50:19.366+02:00	6	
6	stan	22220000-2200-0000-0000-444400004455	com.evolveum.midpoint.model.impl.lens.projector.Projector.focusActivation	INITIAL.e0p0	SUCCESS	2021-10-19T11:50:19.367+02:00	7	
6	stan	22220000-2200-0000-0000-444400004455	com.evolveum.midpoint.model.impl.lens.projector.Projector.focusCredentials	INITIAL.e0p0	SUCCESS	2021-10-19T11:50:19.367+02:00	10	
6	stan	22220000-2200-0000-0000-444400004455	com.evolveum.midpoint.repo.cache.RepositoryCache.getVersion	ResourceType	SUCCESS	2021-10-19T11:50:19.384+02:00	20	oid=[10000000-0000-0000-0000-000000000004]; type=[http://midpoint.evolveum.com/xml/ns/public/common/common-3#Resource
6	stan	22220000-2200-0000-0000-444400004455	com.evolveum.midpoint.repo.cache.RepositoryCache.getVersion	ResourceType	SUCCESS	2021-10-19T11:50:19.392+02:00	13	oid=[10000000-0000-0000-0000-000000000004]; type=[http://midpoint.evolveum.com/xml/ns/public/common/common-3#Resource
6	stan	22220000-2200-0000-0000-444400004455	com.evolveum.midpoint.repo.api.RepositoryService.hasConflict		SUCCESS	2021-10-19T11:50:19.403+02:00	800	watcherClass=[com.evolveum.midpoint.repo.sqlbase.ConflictWatcherImpl]; oid=[0aea041e-2369-46fc-a78e-8d13a5bb47ab]
6	stan	22220000-2200-0000-0000-444400004455	com.evolveum.midpoint.repo.api.RepositoryService.getVersion	ObjectType	SUCCESS	2021-10-19T11:50:19.403+02:00	777	oid=[0aea041e-2369-46fc-a78e-8d13a5bb47ab]; type=[com.evolveum.midpoint.xml.ns._public.common.common_3.ObjectType]

Bucket Size Analysis

Evolveum

- Important for attaining top performance
  - Large/few buckets
    - inefficient distribution, coarse %, wasteful restarts
  - Small/many buckets
    - large overhead (including the contention)
  - Uneven distribution
    - short & long runners
- Options: ex-post, ex-ante

Bucket Sampling

```
<activity>
    <work>
        <noOp>
            <steps>10000</steps>
            <delay>10</delay>
            <stepInterruptibility>full</stepInterruptibility>
        </noOp>
    </work>
    <executionMode>bucketAnalysis</executionMode>
    <distribution>
        <buckets>
            <implicitSegmentation>
                <!-- 10000 steps divided into 200 buckets -->
                <numberOfBuckets>200</numberOfBuckets>
            </implicitSegmentation>
            <sampling>
                <regular>
                    <sampleSize>10</sampleSize>
                </regular>
            </sampling>
        </buckets>
    </distribution>
    <reporting>
        <executionReports>
            <buckets/>
        </executionReports>
    </reporting>
</activity>
```

sequentialNu	content	content-from	content-to	size
1	[-A)		A	100
2	[A-B)	A	B	450
3	[B-C)	B	C	200
4	[C-D)	C	D	200
5	[D-E)	D	E	200
6	[E-F)	E	F	50
7	[F-G)	F	G	200
8	[G-H)	G	H	350
9	[H-I)	H	I	100
10	[I-J)	I	J	150
11	[J-K)	J	K	0
12	[K-L)	K	L	50
13	[L-M)	L	M	150
14	[M-N)	M	N	250
15	[N-O)	N	O	50
16	[O-P)	O	P	100
17	[P-Q)	P		
18	[Q-R)	Q		
19	[R-S)	R		
20	[S-T)	S		
21	[T-U)	T		
22	[U-V)	U		
23	[V-W)	V		
24	[W-X)	W		
25	[X-Y)	X		
26	[Y-Z)	Y		
27	[Z-)	Z		

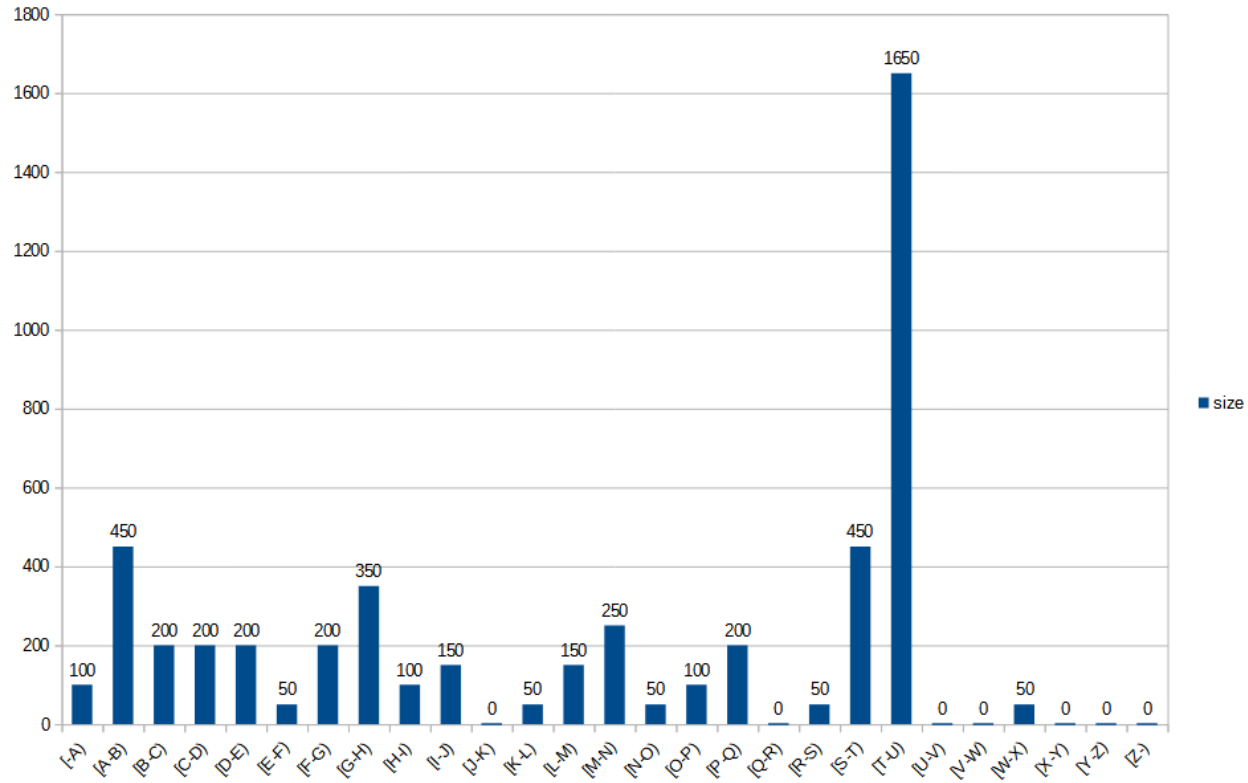

Other options: counting all buckets, using random sampling (by sample size or probability), or using custom expression.

Beware: Counting may not be supported by all resources!

Other New Features

Evolveum

- Failed objects re-processing (since 4.3)
  - Re-process failed objects during separate task execution
  - Re-synchronize failed accounts automatically after given time
- Buckets and items logging configuration
  - E.g. logging full information for each bucket completion, and brief one for each item
- Flexible tracing (before/after conditions), e.g. 🧪
  - Trace each 100-th item (fully) and each 10-th item (briefly)
  - Trace all failed items
  - Trace all items that took longer than 5,000 milliseconds
- Buckets and items processing conditions 🧪
  - Choosing what buckets and/or items to process, and what to skip.

Evolveum

Migration from 4.0.x/4.3.x

Migration & GUI Support

- Legacy-style task execution is supported for simple tasks
  - but no GUI support for these tasks,
  - and no support for advanced features (buckets, workers, partitions) at all.
- Task state data structures has changed
  - so it is strongly advisable to re-import even simple tasks,
  - an exception: live sync token and last scan timestamp are migrated automatically.
- The guide to tasks migration is in docs
  - https://docs.evolveum.com/midpoint/reference/tasks/activities/migration/
- But midPoint Studio can help with the migration! 🧪

Migration: An Example

Evolveum

- Assuming you are running 4.0.4 or 4.3.2
- Before stopping midPoint:
  - suspend the tasks,
  - export the tasks (unless kept externally e.g. in Studio).
- Do the upgrade according to the documentation
- Start midPoint and either:
  - migrate the tasks in Studio and re-import them,
  - or simply resume them – but there will be warnings in logs and some pieces missing in GUI.

Migration Demo

Evolveum

Evolveum

Conclusion

Conclusion

Evolveum

- Tasks were substantially changed in midPoint 4.4.
- Besides quality improvements, there are a lot of new features.
- Have a look at them, try them, and let us know.
  - (We are interested especially in the experimental ones.)
- Everything is in the docs!
  - https://docs.evolveum.com/midpoint/reference/tasks/activities/

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

Webinars

- 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