Performance

September 06, 2023

Michael Klug
Principal Architect, Architecture Services
SailPoint Technologies, Inc.
mike.klug@sailpoint.com

Performance Overview

- Platform
  - Cloud
  - On-Prem
- Hardware
  - CPU
  - Memory
- Network
- Software
  - Database
  - Application Servers
  - Java
  - IdentityIQ
- Performance Tuning
- Custom Coding
- Troubleshooting

Platform

Cloud

- AWS
- Azure
- GCP
- Cloud Native Database can offer benefits
  - Be wary of MySQL/Aurora due to size limitations
  - 8.4 supports PostgreSQL
- How to connect to end points
  - Direct Connect
  - Cloud Gateway

On-Prem

- Still common
- Remote Data Centers with end points
  - Cloud gateway

Hardware

Defined footprints should be considered a minimum and expanded as needed

https://community.sailpoint.com/t5/Other-Documents/IdentityIQ-Hardware-Sizing-Guide/ta-p/79521#toc-hId-479922737

Database

- **Monitoring and Alerting should be enabled**
- **CPU**
  - <40% is healthy
  - Spikes over 95% can be healthy if recovering
- **Memory**
  - <40% is healthy
  - Spikes over 95% can be healthy if recovering

Hardware

Application Server

- **Monitoring and Alerting should be enabled**
- **CPU**
  - 60-70% is healthy
  - Spikes over 95% can be healthy if recovering
- **Memory**
  - Depending on how JVM heap memory is configured, this can be difficult to monitor
    - Xmx = Xms vs Xmx > Xms
    - Should look at current java memory usage rather than OS memory usage
  - JVM heap should recover with Garbage Collection

Network

Latency

- Minimize latency between application server and database
  - 300 microsecond ping
    - If this is not achievable, it can affect performance but should not be a “showstopper”
    - 1ms is common
  - Database performance test
    - 9/17/20
    - If this is not achievable, it can affect performance but should not be a “showstopper”
      - 2-3x are common and depends on certain factors

Software

Database

- **Indexing**
  - Custom indexing may be required for optimal performance
- **Maintenance**
  - Index Rebuilds
  - Statistics Updates
- **Asynchronous replication for Disaster Recovery scenarios**

Application Server

- **Configurations**
  - Secure Connection Configuration
  - HTTP2 for better UI interactions
- **Java**
  - Max memory (Xmx) recommended at minimum 3GB per CPU
  - 32GB of Heap is a relative maximum
    - If >32GB is required, go to 48GB or more

IdentityIQ

IdentityIQ Configurations

- **iiq.properties**
  - **Database Connection Properties**
    - maxTotal > 200
    - Idle Connection Configuration
    - Test Connection Properties
    - Unique Keystore (Security)
  - **Task Threads**
    - scheduler.quartzProperties.org.quartz.threadPool.threadCount
    - Default value is 5
    - Value often doubled to 10
  - **Beanshell Parallel Instances**
    - ruleRunnerPoolConfig.maxTotalPerKey
    - Default value is 8
    - Value typcially set to Aggregation/Refresh maxThreads

IdentityIQ

IdentityIQ Configurations

- **Request Definition Threads**
  - Most maxThreads configurations set to 1.5 to 2x CPU count and is dependent on Application and Database Server Health
    - Aggregation
    - Refresh
    - Certification
    - Perform Maintenance
      - Pruning
      - Workitem – set to 3 or 4 due to WorkflowCase size
    - Workflow – set to 3 or 4 due to WorkflowCase size
    - Custom Request queues can be set up
  - Request processing throttle, which will lead to blocking but can help preserve computing resources, can be set
    - Request ServiceDefinition object -> maxThreads
    - Server object -> maxRequestThreads which overrides the Request ServiceDefinition object

IdentityIQ

IdentityIQ Configurations

- **Aggregations**
  - Enable Partitioning when supported at the Connector Layer
    - Balance the partitions as much as possible
  - Disable optimization only if logic (Correlation) changes
- **Refresh**
  - Enable Delta and Consolidate Options
  - Enable Partitioning
- **Enable Role and Policy Cache**
  - Not enabled by default
- **Certification**
  - Targeted
- **Workflow**
  - foregroundProvisioning to false
  - backgroundApprovalCompletion or backgroundApprovalCompletionIfLocked
  - > 20 items can cause delays in request submission

IdentityIQ

IdentityIQ Data Profiles

- **Accounts (Links)**
  - > 250 can be non-performant
  - Could be common with Server or Database Administrators
  - With non-human Identities, this should be avoided
- **Entitlements**
  - > 10,000 can be non-performant
  - Watch for fine grained entitlements that are not “actionable” by the connector
- **Direct Reports**
  - > 250 can be non-performant
  - Valid large counts can be related to a contractor or external organization manager
  - Look for inactive managers and inactive direct reports
- **Workgroups**
  - > 300 can be non-performant
  - Watch for duplicate (or close) memberships

IdentityIQ

IdentityIQ Data

- **Large Counts of Transactional Data**
  - Set Retention Policies (operational vs audit)
    - Task Results
    - Provisioning Transaction
    - Syslog Event
    - Identity Snapshot
    - Identity Request
  - Workflow Case should only have currently executing workflow instances
    - >30 days should be analyzed and cleaned
  - Identity Archive used systematically by Aggregation and Refresh (Lifecycle Events)
    - >1 should troubleshoot and may be related to commitTransaction executed during refresh
    - > a few days should be analyzed and cleaned
- **Large Counts of Warehouse Data**
  - Prune if possible
    - Identity
    - Application

IdentityIQ

IdentityIQ Data

- **Archival Use Cases**
  - 8.4 has a Data Export Task that can push any IdentityIQ data to an external endpoint
    - Audit Event
  - Export and Prune from Active Tables
- **Reporting Use Cases**
  - Data Export for Certification and User Access Reports
  - 8.4 Data Export

Performance Tuning

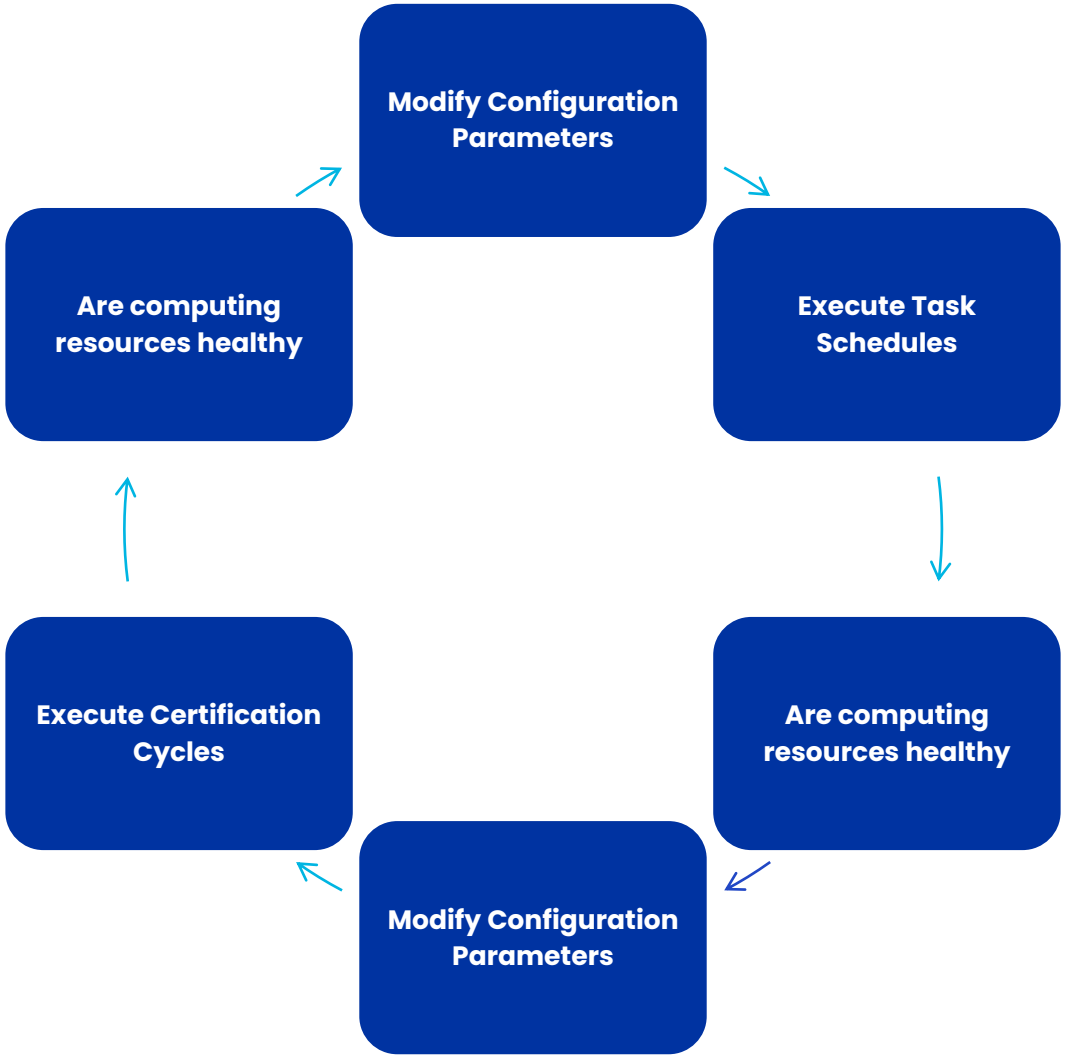

Custom Coding

Projection Searches

- Preferred search
  - Minimal data returned
  - Must be a “property” of the object
    - A database column
- Flush the iterator

Decaching

- Bloated hibernate cache cause degraded performance
- Only if the current logic has fetched the results
- Can lead to lazy initialization errors

Object Locking

- Object Locking is GOOD!! Only bad if held for longer than expected
  - ATOMIC Database Transaction
  - If a parent process does not own a lock, always obtain the lock of an object before updating (saveObject + commitTransaction) to prevent corruption

Troubleshooting

Support Plugin

- https://community.sailpoint.com/t5/Plugin-Framework/SailPoint-IdentityIQ-Support-Data-Collector-Plugin/ta-p/139576
- Allows for Cross Server Data Gathering
  - Logging
  - Thread Dumps
  - If server experiencing CPU or Memory issues, gather logs and thread dumps via OS
- Gather Environment Information
- Gather Database Performance Metrics

Thank You!