Health Check and Common Troubleshooting

Health Check and Common Troubleshooting

How to maximize system performance within Informer 5

Kevin Houk, Implementation Specialist, Entrinsik

Data Governance Tenets

- Data Integrity
- Data Level Security
- Content Level Security
- Functional Level Security
- Transparency
- Traceability
- Data Quality
- Controlled Reuse
- Governed Blended Data

ENTRINSICON
LEVEL UP!

Transparency & Traceability

- System Setup
- Auditing
- Logging
- System Health Check

System Setup

Windows server specs *minimum*

- Windows Server 2016, 2019, 2022 (all updates)
  - No IIS or other applications on server
- Quad-core CPUs
- 16GB RAM (32GB recommended)
- 80GB+ disk space
- Client browser: Anything but Internet Explorer!
- https://informer5.zendesk.com/hc/en-us/articles/11500498926-Hardware-and-Software-Requirements

ENTRINSICON
LEVEL UP!

System Setup

Docker / Linux server specs

- Hardware specs same as Windows Server
- Supported Linux Distributions
  - Ubuntu, Debian, Fedora (must be able to support Docker)
- Docker
  - Docker Community Edition (free! use me!)
  - Docker Enterprise Edition (not free – but if you have it, use it)
- Docker Compose
- https://informer5.zendesk.com/hc/en-us/articles/115004498926-Hardware-and-Software-Requirements#Linux_Requirements

System Setup

Windows Tuning - Elasticsearch

- Careful when tweaking with just minimum hardware specs
- **Elasticsearch**
  - Check Java memory usage in Task Manager
  - 4 GB of memory typical if server hardware meets minimum requirements
  - Can be increased – check with Entrinsik Support
  - Consider increasing ES memory if using a lot of large Datasets, Visuals, Dashboards and have a lot of users

Elasticsearch 8

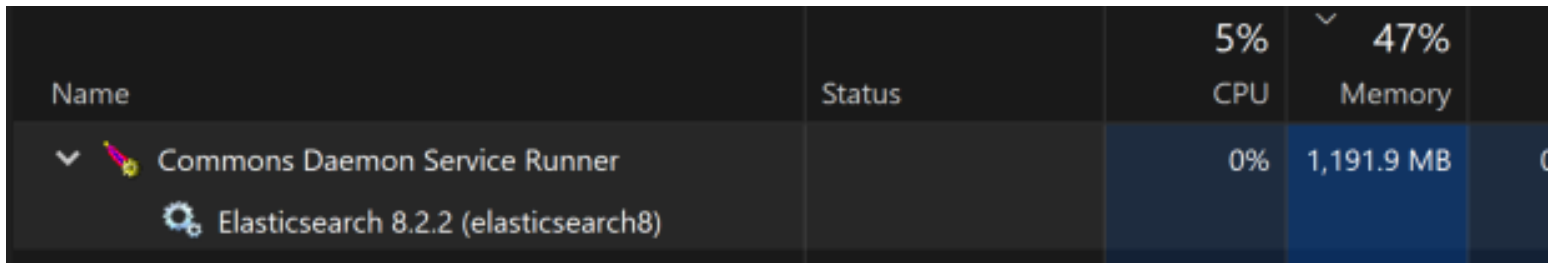

ENTRINSICON
LEVEL UP!

System Setup

Windows Tuning - Elasticsearch (cont.)

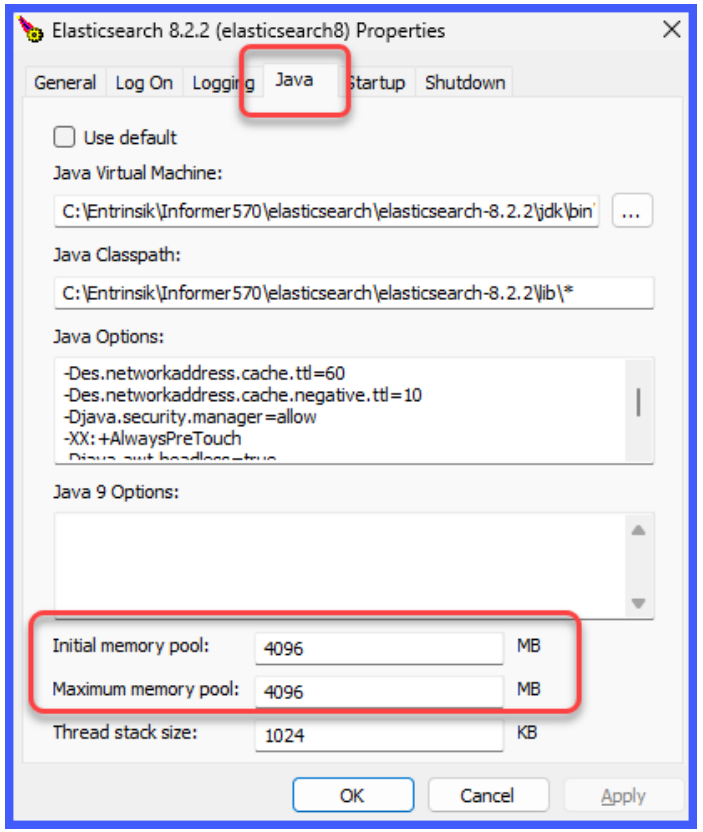

C:\Entrinsik\Informer5\elasticsearch\elasticsearch-8.2.2\bin**elasticsearch-service.bat manager elasticsearch8**

Initial and Maximum Memory Pool ***MUST*** be the same

Leave Thread stack size at 1024 KB

Restart ES Service if changes made

ENTRINSICON
LEVEL UP!

System Setup

Windows Tuning - Load Balancer

- Careful when tweaking with just minimum hardware specs
- The Load Balancer sets up separate processes for the different core components of Informer 5
  - **Queries** - handles all database query executions
  - **API** - the main web server component, handling all UI requests
  - **Scheduler** - handles all Job executions
  - **Reverse Proxy** - acts as a proxy server and is the main entry point for all requests
- Spreading the load across different threads can improve performance of the Informer application when handling larger volumes of requests and/or data
- **C:\Entrinsik\Informer5\service.config.json**

System Setup

Windows Tuning - Load Balancer (cont.)

Default

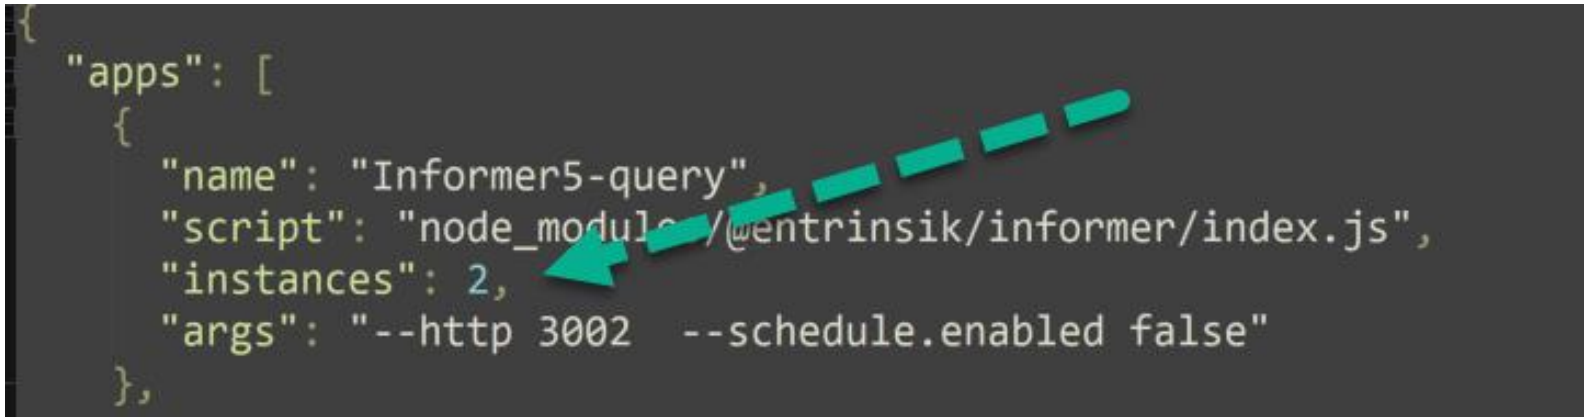

Two instances – each using 1GB memory

ENTRINSICON
LEVEL UP!

System Setup

Windows Tuning - Load Balancer (cont.)

Enhanced

```
{
  "apps": [
    {
      "name": "Informer5-query",
      "script": "node_modules/@entrinsik/informer/index.js",
      "instances": 2,
      "args": "--http 3002  --schedule.enabled false",
      "node_args": "--max-old-space-size=2048",
      "max_memory_restart" : "2500M"
    },
```

node_args - Query thread using 2GB instead of 1GB. Two instances each using 2GB gives 4GB to the Query thread

max_memmory_restart - prevents a thread from locking up if memory exceeds 2GB, this setting will cause the thread to automatically restart if it hits 2.5 GB

System Setup

Windows Tuning - Load Balancer (cont.)

```
{
  "apps": [
  {
    "name": "Informer5-query",
    "script": "node_modules/@entrinsik/informer/index.js",
    "instances": 2,
    "args": "--http 3002 --schedule.enabled false",
    "node_args": "--max-old-space-size=2048",
    "max_memory_restart" : "2500M"
  },
    {
    "name": "Informer5-api",
    "script": "node_modules/@entrinsik/informer/index.js",
    "instances": 2,
    "args": "--http 3002 --schedule.enabled false",
    "node_args": "--max-old-space-size=2048",
    "max_memory_restart" : "2500M"
  },
```

ENTRINSICON
LEVEL UP!

System Setup

Windows Tuning - Load Balancer (cont.)

```
19        {
20          "name": "Informer5-reverse-proxy",
21          "script": "node_modules/@entrinsik/i5-load-balancer/index.js",
22          "instances": 1
23      },
24        {
25        "name": "Informer5-schedules",
26        "script": "node_modules/@entrinsik/informer/index.js",
27        "instances": 2,
28        "args": "--http 3002 --schedule.enabled false",
29        "node_args": "--max-old-space-size=2048",
30        "max_memory_restart" : "2500M"
31    },
32      ]
33    }
```

ENTRINSICON
LEVEL UP!

System Setup

Windows Tuning - Load Balancer (cont.)

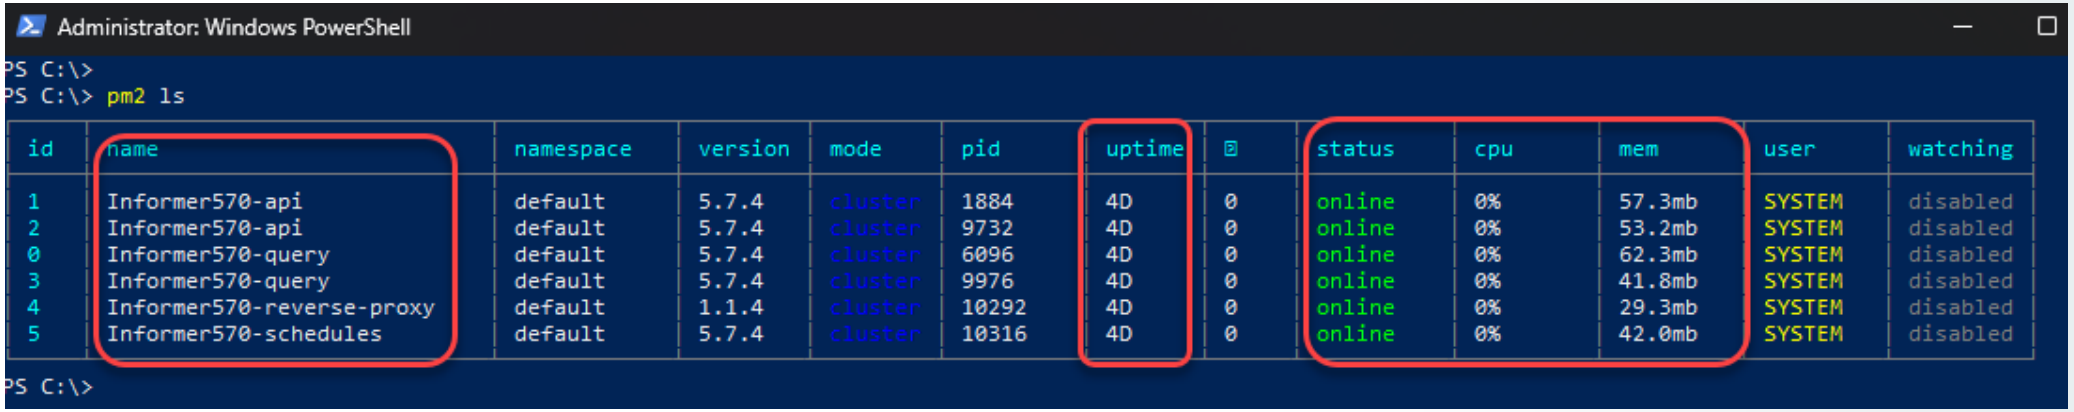

C:\> pm2 ls

- Use this command to see useful data about your worker threads
- Uptime and Status – has any thread crashed and/or restarted
- How much memory and CPU is currently being used?
- Snapshot in time

ENTRINSICON
LEVEL UP!

System Setup

Windows Tuning - Redis

Redis is a temporary data storage area for certain processes in Informer. Imports. Exports. Query Samples, Flow Steps, etc.

C:\Program Files\Redis\redis.windows.conf

Disable Snapshots saving to disk. Saves a lot of disk I/O

```
193
194    # save 900 1
195    # save 300 10
196    # save 60 10000
```

ENTRINSICON
LEVEL UP!

System Setup

Windows Tuning - Redis

Disable Redis Snapshot Persistence:

```
211    stop-writes-on-bgsave-error no
```

C:\Program Files\Redis\redis-cli.exe

info memory
flushall

Transparency and Traceability - Auditing

System Audit

- Which queries are consuming the most time?
- Which users are creating them?
- Can the queries be made more efficient?
- What content isn't being used?
- Are there 10 Jobs queued to run at the same time?

Data Governance

- Clear where data is coming from
- Clear on who is doing what to relative content
- Verify only authorized people able to access what they're supposed to
- Eliminate lack of confidence in data backing important business decisions

Auditing

System Audit

- The new way to monitor system resources to keep Informer running efficiently.
- Who is using Informer? And how?
- How fast are queries running?
- What content isn’t being used?
- See how Users are using Informer, optimize slow-running queries across the system and clean out unused content from Informer.

Auditing

System Audit (Administration > Settings > Audit)

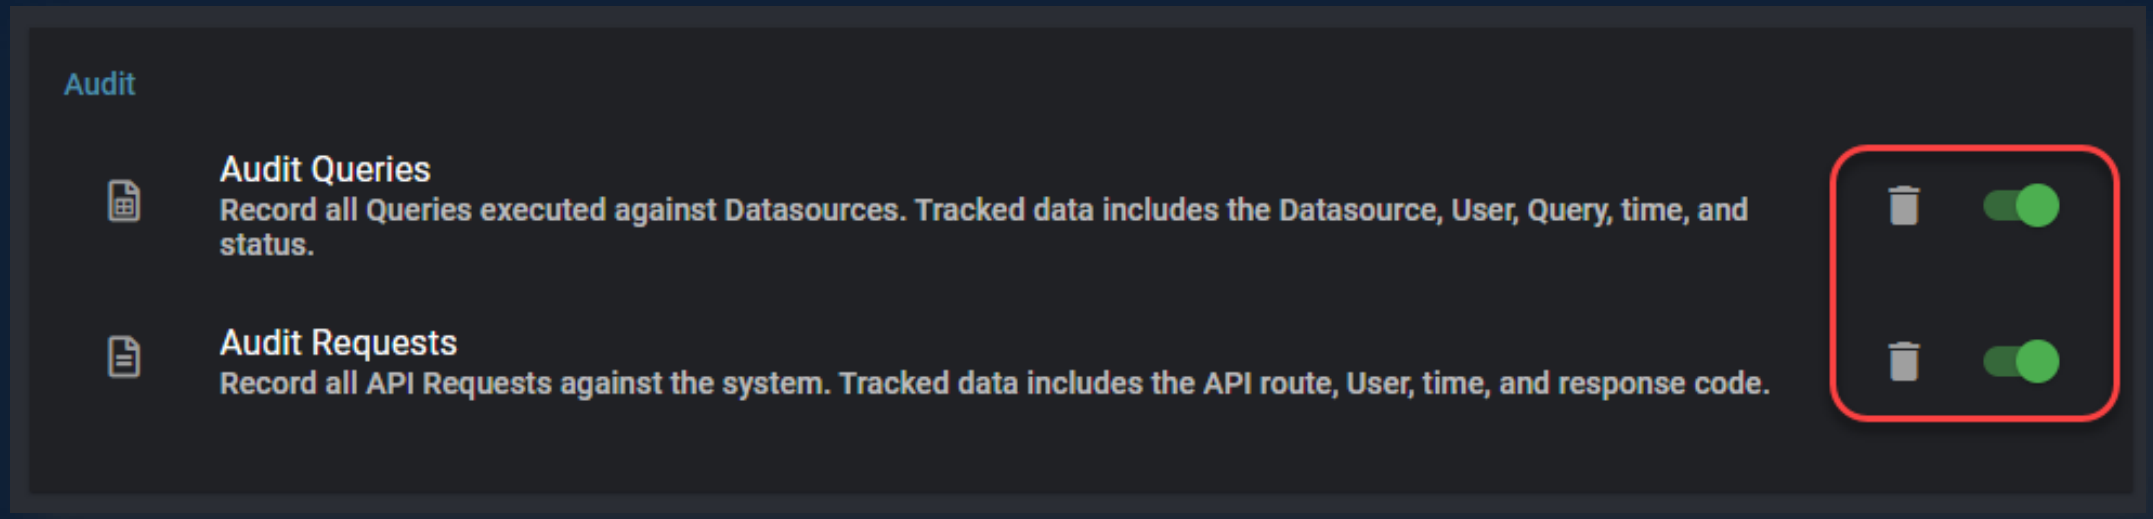

Auditing

Audit Queries

- Records all Queries executed against Datasources
- Tracked data includes the Datasource, User, Query time and status
- Allows Administrators to track what parts of Informer are being used the most often and when and who is using which parts

Auditing

Audit Requests

- Records all API Requests made against Informer.
- Tracked data includes the API route, User, Time and Response Code.
- Useful to Admins who want to keep Informer running at peak performance

Transparency & Traceability - Logging

Request Logs: Contains information about requests made of Informer, such as the time it took for a request to be processed, the type of request and additional information about a specific log message.

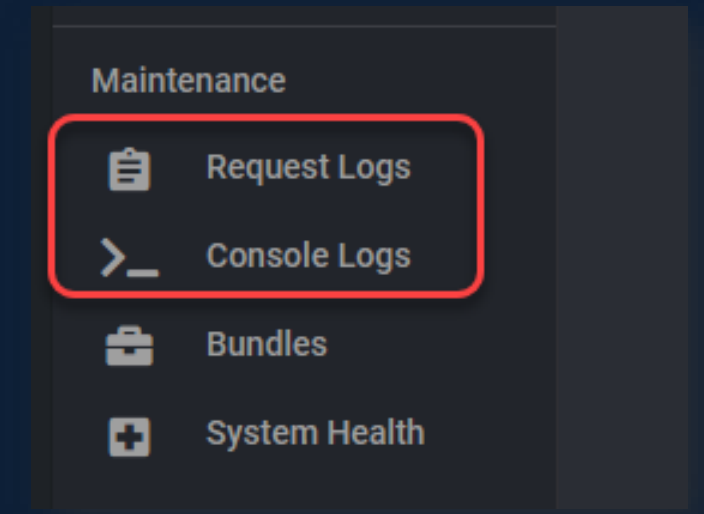

Console Logs: Contains a live feed to API calls that Informer is making as they are made.

Logging

Request Logs and Console Logs

- Unless troubleshooting something specific, do not set Log Level to above `Info`
- Request Logs can be filtered by Time, Tags, Result and other parameters, while Console Logs cannot

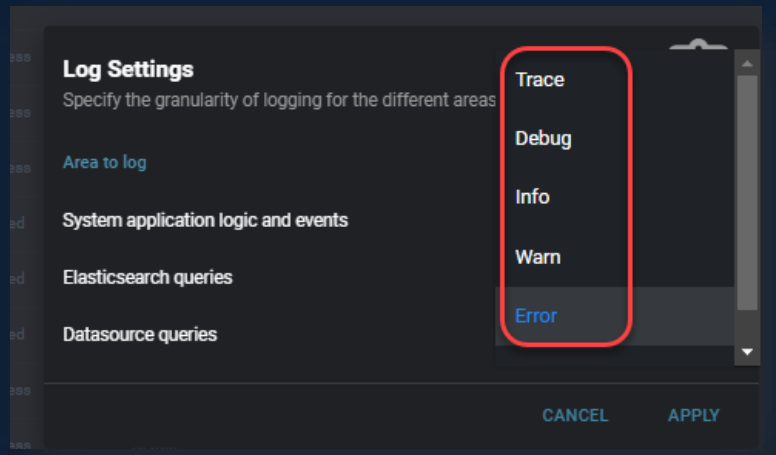

Logging

Log Areas

- System application logic and events
- Elasticsearch queries
- Datasource queries

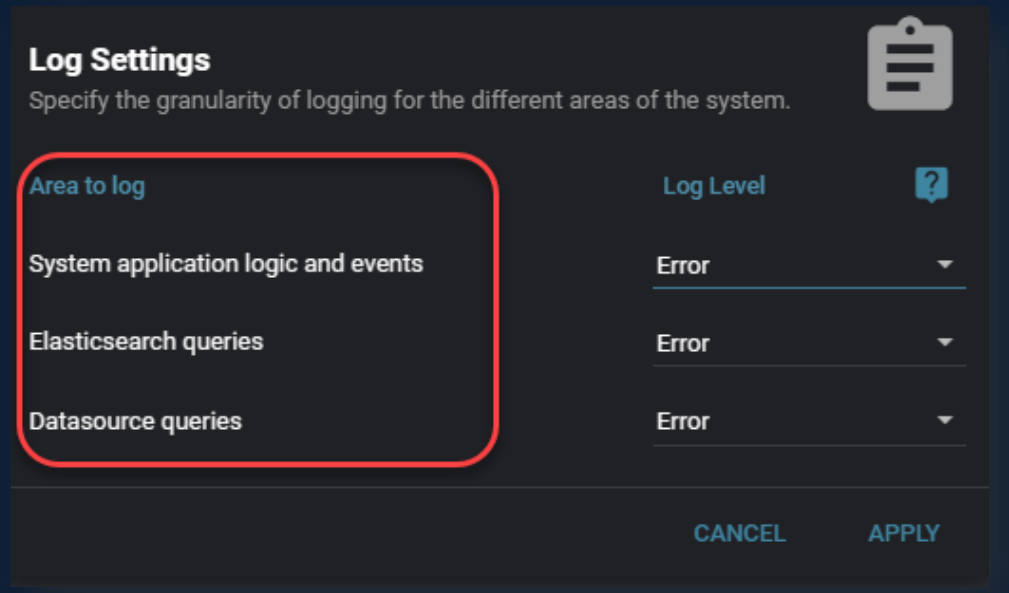

ENTRINSICON
LEVEL UP!

Logging

Log Levels

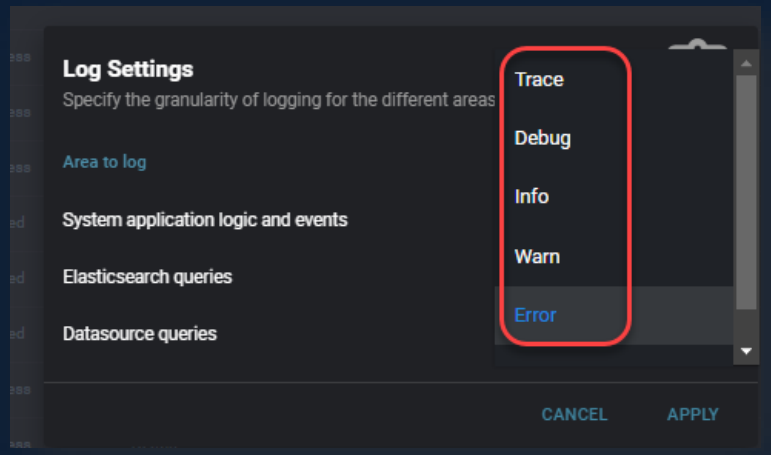

- None – No logs collected for area
- Error – Captures only events that prevents an action from completing
- Warn – Captures logs anytime something unexpected happens that Informer can handle

Logging

Log Levels

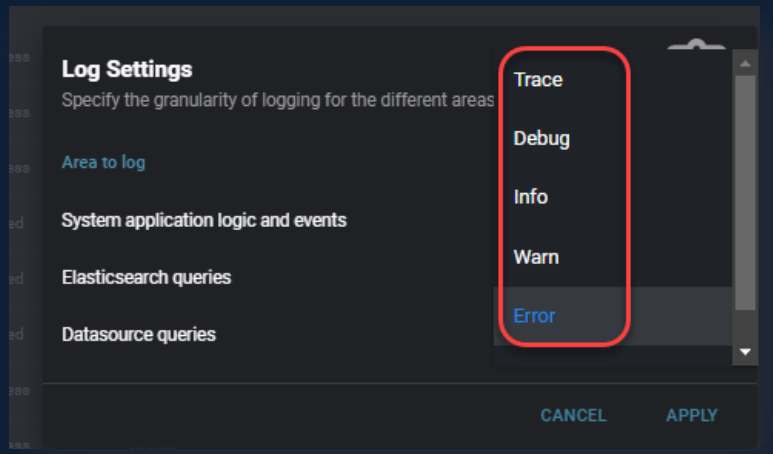

- Info – Lowest level of logs that also include *successful* requests and logs. *Also contains all warnings and errors*.
- Debug – 2nd highest volume of logging. Contains all debugging info.
- Trace – Highest level of logging. Captures all the logs included in Informer giving highest detail of what's happening all of the time.

Logging

Log Filters

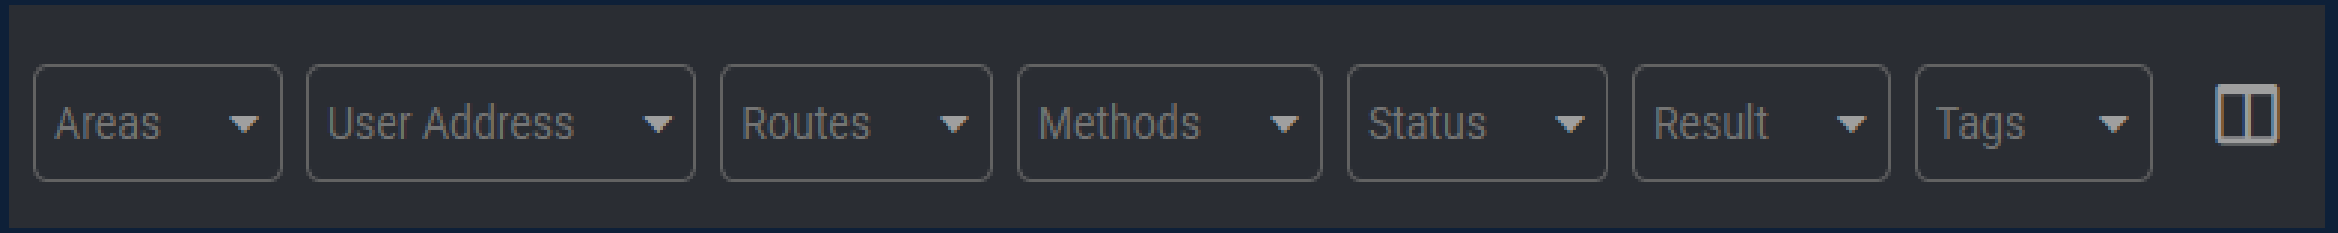

Logging

Browser Logging Settings

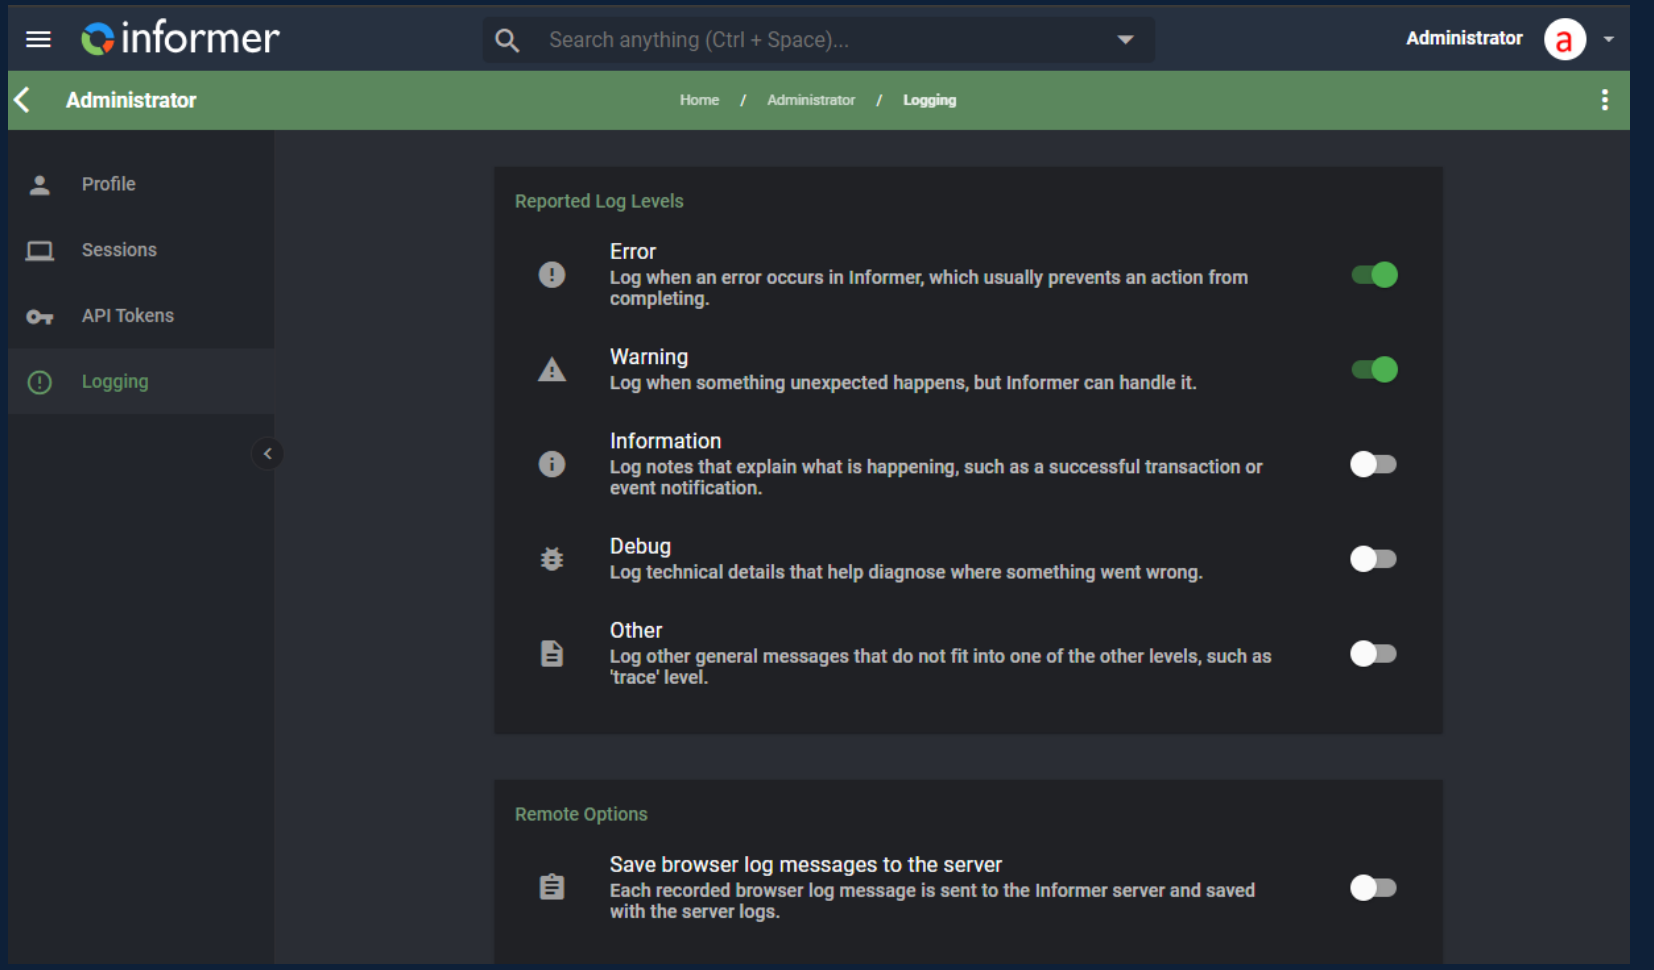

ENTRINSICON
LEVEL UP!

Transparency & Traceability – System Health

- Allows users to see information on Datasets and Elasticsearch indices

- Can help Admins understand what is happening in Elasticsearch from within Informer

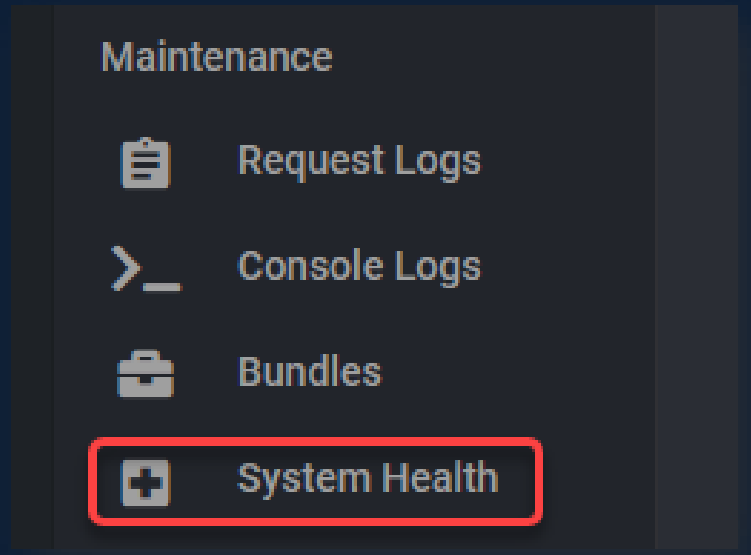

ENTRINSICON
LEVEL UP!

System Health

Unindexed Datasets

A Dataset that doesn’t have any data because it’s not been refreshed, or the data has been cleared. As well, the Dataset could be missing an index if it’s deleted or corrupted at the Elasticsearch level.

Can drilldown to open a window that shows specifics

System Health

Unindexed Datasets

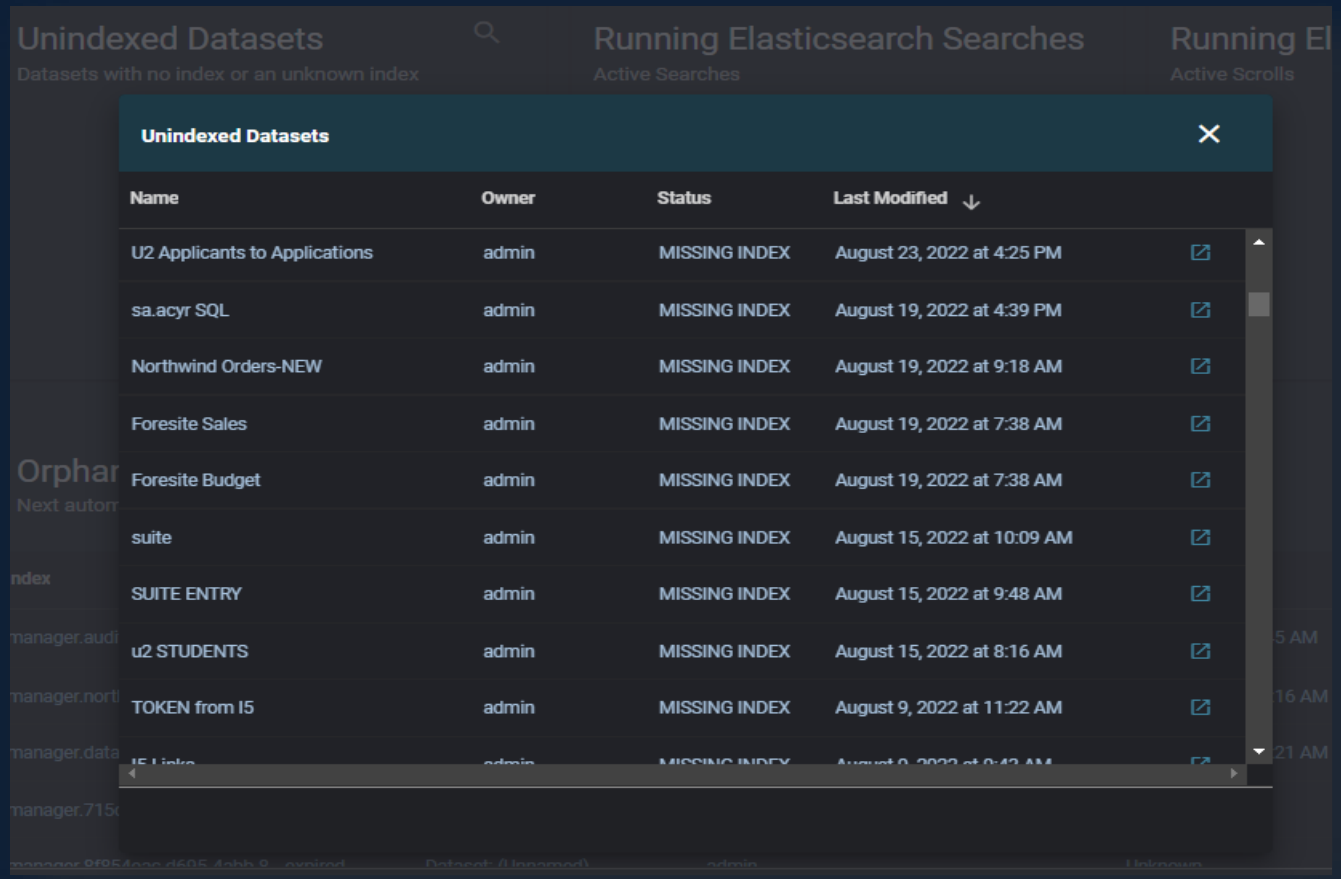

ENTRINSICON
LEVEL UP!

System Health

Running Elasticsearch Searches & Scrolls

- **Searches**: Returns hits that match the Query definition in the query's request
- **Scrolls**: Retrieves a stream of search results matching the query
- *Note:* High numbers can result in performance issues

System Health

Running Elasticsearch Searches & Scrolls

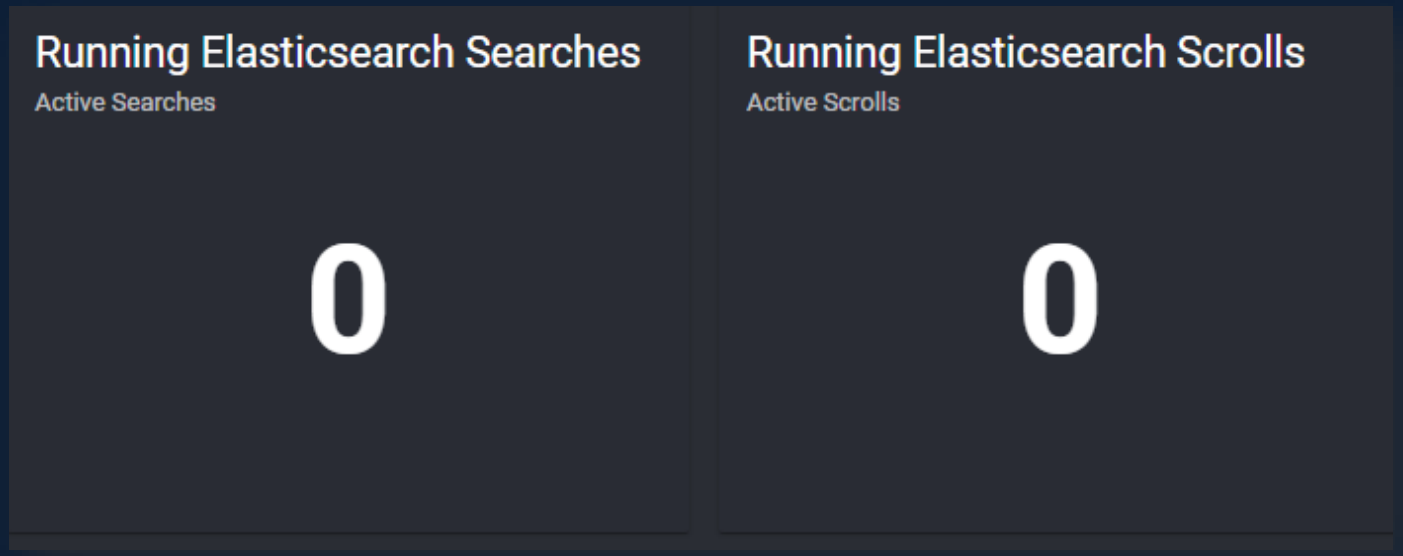

System Health

Orphaned & Expired Indices

- Indices (results) that don’t have an associated Dataset or Report. Their time on the system has expired but are still hanging around.

- System delivered Auto Cleanup job should be running every night

System Health

Orphaned & Expired Indices

ENTRINSICON
LEVEL UP!

System Health

Query Benchmarking

- Helps identify two possible bottlenecks in a query's performance
- Indexing: evaluating the connection speed between Informer and Elasticsearch
- Flow Steps: are there slowdowns because of a particular Flow Step?

System Health

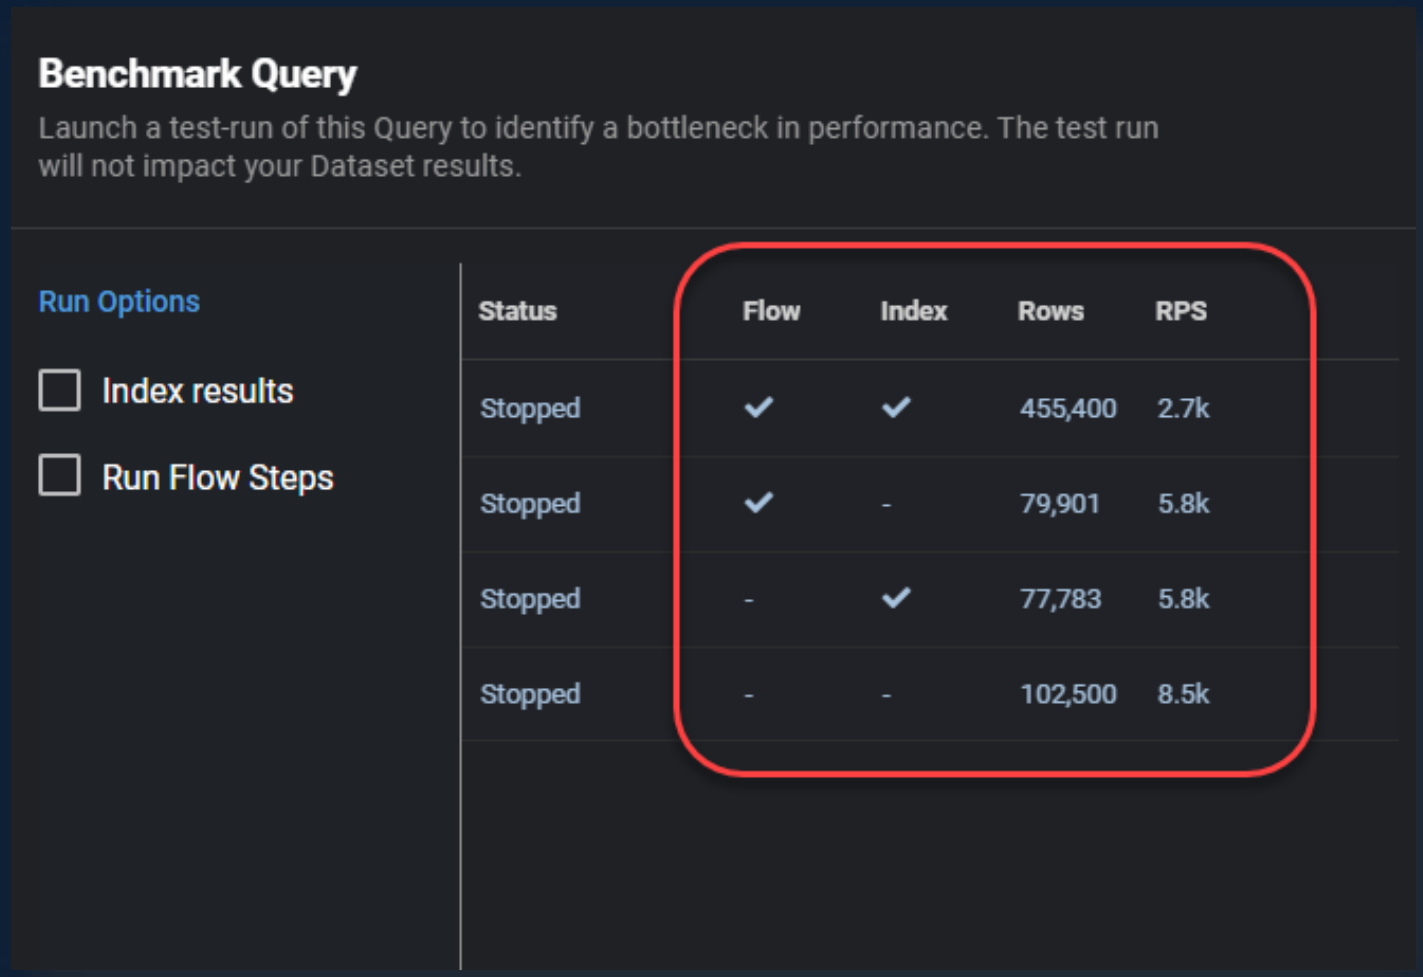

ENTRINSICON
LEVEL UP!

Windows Tuning – Job Setup

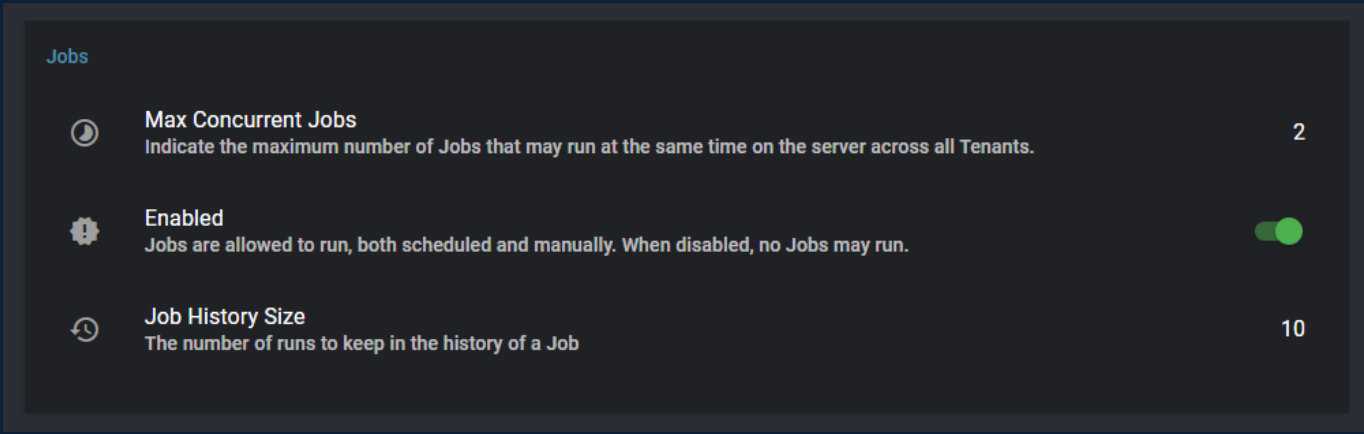

Spread out Jobs – give them room to run and complete

Lower Job History Size. The larger this is, the response time in the Job area will be reduced.

Q & A

ENTRINSICON
LEVEL UP!