

Migration Utils Metrics

The migration-utils offers optional server metrics using [Prometheus](#) for those who have long running migrations.

To enable metrics collection, the `--enable-metrics` command line option must be used when running the migration-utils. With this option on, the migration-utils exposes an endpoint that can be used by Prometheus to collect metrics at `http://localhost:8080/prometheus`.

Example usage

```
java -jar target/migration-utils-6.0.0-SNAPSHOT-driver.jar --enable-metrics --source-type=akubra --target-dir=target/test/ocfl --objects-dir=src/test/resources/legacyFS/objects --datastreams-dir=src/test/resources/legacyFS/datastreams
```

The migration-utils requires that port 8080 on your local machine is available in order to serve the endpoint for Prometheus.

Example Metrics Collection Setup

The following is an example of how you could setup Prometheus and [Grafana](#) using [Docker](#) to collect and monitor metrics. It is a quick and easy setup, and not intended to demonstrate a production ready install. This setup assumes that you already have Docker and [Docker Compose](#) installed on your system.

Linux Setup

Create a Prometheus config file, named `prometheus.yml`, that looks similar to the following:

prometheus.yml

```
scrape_configs:
  - job_name: 'fcrepo'
    scrape_interval: 30s
    metrics_path: '/prometheus'
    static_configs:
      - targets: ['localhost:8080']
```

Next, create a Docker compose file in the same directory, named `docker-compose.yml`, that looks like the following:

docker-compose.yml

```
version: '3.8'
services:
  prometheus:
    image: prom/prometheus
    ports:
      - "9090:9090"
    volumes:
      - ./prometheus.yml:/etc/prometheus/prometheus.yml
    network_mode: "host"

  grafana:
    image: grafana/grafana
    ports:
      - "3000:3000"
    network_mode: "host"
```

Mac Setup

Create a Prometheus config file, named `prometheus.yml`, that looks similar to the following:

prometheus.yml

```
scrape_configs:
  - job_name: 'migration-utils'
    scrape_interval: 30s
    metrics_path: '/prometheus'
    static_configs:
      - targets: ['host.docker.internal:8080']
```

Next, create a Docker compose file in the same directory, named `docker-compose.yml`, that looks like the following:

docker-compose.yml

```
version: '3.8'
services:
  prometheus:
    image: prom/prometheus
    ports:
      - "9090:9090"
    volumes:
      - ./prometheus.yml:/etc/prometheus/prometheus.yml

  grafana:
    image: grafana/grafana
    ports:
      - "3000:3000"
```

Collecting Metrics

From within this directory, execute `docker-compose up` to startup the services. You can confirm the services started by browsing to <http://localhost:9090>, the Prometheus endpoint, and <http://localhost:3000>, the Grafana endpoint.

Note, this will run the process in the foreground. If you prefer to run it in the background, use the `-d` flag.

The final bit of setup is to connect Grafana to Prometheus as follows:

1. Log into Grafana with the default credentials: `admin/admin`
2. Navigate to Configuration -> Datasources
3. Add a Prometheus datasource
4. For the Prometheus URL, if you're on linux specify <http://localhost:9090>, and if you're on Mac use <http://host.docker.internal:9090>. For amazon linux 2 use <http://prometheus:9090>
5. Save the configuration

Now that you have the datasource added, you can create a dashboard to display metrics. Here is an example dashboard with some pre-configured graphs:

Migration Utils Example Dashboard

```
{
  "__inputs": [
    {
      "name": "DS_PROMETHEUS",
      "label": "Prometheus",
      "description": "",
      "type": "datasource",
      "pluginId": "prometheus",
      "pluginName": "Prometheus"
    }
  ],
  "__requires": [
    {
      "type": "grafana",
      "id": "grafana",
      "name": "Grafana",
      "version": "7.5.3"
    },
    {
```

```

    "type": "panel",
    "id": "graph",
    "name": "Graph",
    "version": ""
  },
  {
    "type": "datasource",
    "id": "prometheus",
    "name": "Prometheus",
    "version": "1.0.0"
  },
  {
    "type": "panel",
    "id": "timeseries",
    "name": "Time series",
    "version": ""
  }
],
"annotations": {
  "list": [
    {
      "builtIn": 1,
      "datasource": "-- Grafana --",
      "enable": true,
      "hide": true,
      "iconColor": "rgba(0, 211, 255, 1)",
      "name": "Annotations & Alerts",
      "type": "dashboard"
    }
  ]
},
"editable": true,
"gnetId": null,
"graphTooltip": 0,
"id": null,
"links": [],
"panels": [
  {
    "aliasColors": {},
    "bars": false,
    "dashLength": 10,
    "dashes": false,
    "datasource": "${DS_PROMETHEUS}",
    "fieldConfig": {
      "defaults": {},
      "overrides": []
    },
    "fill": 1,
    "fillGradient": 0,
    "gridPos": {
      "h": 8,
      "w": 12,
      "x": 0,
      "y": 0
    },
    "hiddenSeries": false,
    "id": 6,
    "legend": {
      "avg": false,
      "current": false,
      "max": false,
      "min": false,
      "show": true,
      "total": false,
      "values": false
    },
    "lines": true,
    "linewidth": 1,
    "nullPointMode": "null",
    "options": {
      "alertThreshold": true
    }
  }
]

```

```

    },
    "percentage": false,
    "pluginVersion": "7.5.3",
    "pointradius": 2,
    "points": false,
    "renderer": "flot",
    "seriesOverrides": [],
    "spaceLength": 10,
    "stack": false,
    "steppedLine": false,
    "targets": [
      {
        "expr": "jvm_memory_used_bytes",
        "interval": "",
        "legendFormat": "{{id}}",
        "refId": "A"
      }
    ],
    "thresholds": [],
    "timeFrom": null,
    "timeRegions": [],
    "timeShift": null,
    "title": "JVM Memory Usage",
    "tooltip": {
      "shared": true,
      "sort": 0,
      "value_type": "individual"
    },
    "type": "graph",
    "xaxis": {
      "buckets": null,
      "mode": "time",
      "name": null,
      "show": true,
      "values": []
    },
    "yaxes": [
      {
        "format": "bytes",
        "label": null,
        "logBase": 1,
        "max": null,
        "min": null,
        "show": true
      },
      {
        "format": "short",
        "label": null,
        "logBase": 1,
        "max": null,
        "min": null,
        "show": true
      }
    ],
    "yaxis": {
      "align": false,
      "alignLevel": null
    }
  },
  {
    "aliasColors": {},
    "bars": false,
    "dashLength": 10,
    "dashes": false,
    "datasource": "${DS_PROMETHEUS}",
    "fieldConfig": {
      "defaults": {},
      "overrides": []
    },
    "fill": 1,
    "fillGradient": 0,

```

```

"gridPos": {
  "h": 9,
  "w": 12,
  "x": 12,
  "y": 0
},
"hiddenSeries": false,
"id": 10,
"legend": {
  "avg": false,
  "current": false,
  "max": false,
  "min": false,
  "show": true,
  "total": false,
  "values": false
},
"lines": true,
"linewidth": 1,
"nullPointMode": "null",
"options": {
  "alertThreshold": true
},
"percentage": false,
"pluginVersion": "7.5.3",
"pointRadius": 2,
"points": false,
"renderer": "flot",
"seriesOverrides": [],
"spaceLength": 10,
"stack": false,
"steppedLine": false,
"targets": [
  {
    "expr": "histogram_quantile(0.99, sum(rate(fcrepo_storage_ocfl_object_seconds_bucket[1m])) by (le,
operation))",
    "interval": "",
    "legendFormat": "{{operation}} p99",
    "refId": "A"
  },
  {
    "expr": "rate(fcrepo_storage_ocfl_object_seconds_sum[1m])/rate
(fcrepo_storage_ocfl_object_seconds_count[1m])",
    "interval": "",
    "legendFormat": "{{operation}} avg",
    "refId": "B"
  }
],
"thresholds": [],
"timeFrom": null,
"timeRegions": [],
"timeShift": null,
"title": "OCFL Object Latency",
"tooltip": {
  "shared": true,
  "sort": 0,
  "value_type": "individual"
},
"type": "graph",
"xaxis": {
  "buckets": null,
  "mode": "time",
  "name": null,
  "show": true,
  "values": []
},
"yaxes": [
  {
    "format": "s",
    "label": null,
    "logBase": 1,

```

```

        "max": null,
        "min": null,
        "show": true
    },
    {
        "format": "short",
        "label": null,
        "logBase": 1,
        "max": null,
        "min": null,
        "show": true
    }
],
"yaxis": {
    "align": false,
    "alignLevel": null
}
},
{
    "aliasColors": {},
    "bars": false,
    "dashLength": 10,
    "dashes": false,
    "datasource": null,
    "fieldConfig": {
        "defaults": {},
        "overrides": []
    },
    "fill": 1,
    "fillGradient": 0,
    "gridPos": {
        "h": 8,
        "w": 12,
        "x": 0,
        "y": 8
    },
    "hiddenSeries": false,
    "id": 14,
    "legend": {
        "avg": false,
        "current": false,
        "max": false,
        "min": false,
        "show": true,
        "total": false,
        "values": false
    },
    "lines": true,
    "linewidth": 1,
    "nullPointMode": "null",
    "options": {
        "alertThreshold": true
    },
    "percentage": false,
    "pluginVersion": "7.5.3",
    "pointradius": 2,
    "points": false,
    "renderer": "flot",
    "seriesOverrides": [],
    "spaceLength": 10,
    "stack": false,
    "steppedLine": false,
    "targets": [
        {
            "expr": "rate(jvm_gc_pause_seconds_sum[1m])/rate(jvm_gc_pause_seconds_count[1m])",
            "interval": "",
            "legendFormat": "{{cause}} avg",
            "refId": "A"
        },
        {
            "expr": "histogram_quantile(0.99, sum(rate(jvm_gc_pause_seconds_bucket[1m])) by (le, cause))",

```

```

        "interval": "",
        "legendFormat": "{{cause}} p99",
        "refId": "B"
    }
},
"thresholds": [],
"timeFrom": null,
"timeRegions": [],
"timeShift": null,
"title": "GC Time",
"tooltip": {
    "shared": true,
    "sort": 0,
    "value_type": "individual"
},
"type": "graph",
"xaxis": {
    "buckets": null,
    "mode": "time",
    "name": null,
    "show": true,
    "values": []
},
"yaxes": [
    {
        "format": "dtdurations",
        "label": null,
        "logBase": 1,
        "max": null,
        "min": null,
        "show": true
    },
    {
        "format": "short",
        "label": null,
        "logBase": 1,
        "max": null,
        "min": null,
        "show": true
    }
],
"yaxis": {
    "align": false,
    "alignLevel": null
}
},
{
    "aliasColors": {},
    "bars": false,
    "dashLength": 10,
    "dashes": false,
    "datasource": "${DS_PROMETHEUS}",
    "fieldConfig": {
        "defaults": {},
        "overrides": []
    },
    "fill": 1,
    "fillGradient": 0,
    "gridPos": {
        "h": 9,
        "w": 12,
        "x": 12,
        "y": 9
    },
    "hiddenSeries": false,
    "id": 18,
    "legend": {
        "avg": false,
        "current": false,
        "max": false,
        "min": false,

```

```

    "show": true,
    "total": false,
    "values": false
  },
  "lines": true,
  "linewidth": 1,
  "nullPointMode": "null",
  "options": {
    "alertThreshold": true
  },
  "percentage": false,
  "pluginVersion": "7.5.3",
  "pointradius": 2,
  "points": false,
  "renderer": "flot",
  "seriesOverrides": [],
  "spaceLength": 10,
  "stack": false,
  "steppedLine": false,
  "targets": [
    {
      "exemplar": true,
      "expr": "histogram_quantile(0.99, sum(rate(fcrepo_storage_foxml_object_seconds_bucket[1m])) by (le, operation))",
      "interval": "",
      "legendFormat": "{{operation}} p99",
      "refId": "A"
    },
    {
      "exemplar": true,
      "expr": "rate(fcrepo_storage_foxml_object_seconds_sum[1m])/rate(fcrepo_storage_foxml_object_seconds_count[1m])",
      "interval": "",
      "legendFormat": "{{operation}} avg",
      "refId": "B"
    }
  ],
  "thresholds": [],
  "timeFrom": null,
  "timeRegions": [],
  "timeShift": null,
  "title": "Foxml Object Latency",
  "tooltip": {
    "shared": true,
    "sort": 0,
    "value_type": "individual"
  },
  "type": "graph",
  "xaxis": {
    "buckets": null,
    "mode": "time",
    "name": null,
    "show": true,
    "values": []
  },
  "yaxes": [
    {
      "format": "s",
      "label": null,
      "logBase": 1,
      "max": null,
      "min": null,
      "show": true
    },
    {
      "format": "short",
      "label": null,
      "logBase": 1,
      "max": null,
      "min": null,
      "show": true
    }
  ]

```



```

    }
  ],
  "yaxis": {
    "align": false,
    "alignLevel": null
  }
},
{
  "aliasColors": {},
  "bars": false,
  "dashLength": 10,
  "dashes": false,
  "datasource": null,
  "fieldConfig": {
    "defaults": {},
    "overrides": []
  },
  "fill": 1,
  "fillGradient": 0,
  "gridPos": {
    "h": 8,
    "w": 12,
    "x": 0,
    "y": 16
  },
  "hiddenSeries": false,
  "id": 16,
  "legend": {
    "avg": false,
    "current": false,
    "max": false,
    "min": false,
    "show": true,
    "total": false,
    "values": false
  },
  "lines": true,
  "linewidth": 1,
  "nullPointMode": "null",
  "options": {
    "alertThreshold": true
  },
  "percentage": false,
  "pluginVersion": "7.5.3",
  "pointRadius": 2,
  "points": false,
  "renderer": "flot",
  "seriesOverrides": [],
  "spaceLength": 10,
  "stack": false,
  "steppedLine": false,
  "targets": [
    {
      "expr": "rate(jvm_gc_pause_seconds_count[1m])",
      "interval": "",
      "legendFormat": "{{cause}}",
      "refId": "A"
    }
  ],
  "thresholds": [],
  "timeFrom": null,
  "timeRegions": [],
  "timeShift": null,
  "title": "GC Rate",
  "tooltip": {
    "shared": true,
    "sort": 0,
    "value_type": "individual"
  },
  "type": "graph",
  "xaxis": {

```

```

    "buckets": null,
    "mode": "time",
    "name": null,
    "show": true,
    "values": []
  },
  "yaxes": [
    {
      "format": "reqps",
      "label": null,
      "logBase": 1,
      "max": null,
      "min": null,
      "show": true
    },
    {
      "format": "short",
      "label": null,
      "logBase": 1,
      "max": null,
      "min": null,
      "show": true
    }
  ],
  "yaxis": {
    "align": false,
    "alignLevel": null
  }
},
{
  "datasource": "${DS_PROMETHEUS}",
  "fieldConfig": {
    "defaults": {
      "color": {
        "mode": "palette-classic"
      },
      "custom": {
        "axisLabel": "",
        "axisPlacement": "auto",
        "barAlignment": 0,
        "drawStyle": "line",
        "fillOpacity": 10,
        "gradientMode": "none",
        "hideFrom": {
          "graph": false,
          "legend": false,
          "tooltip": false
        },
        "lineInterpolation": "linear",
        "lineStyle": {
          "fill": "solid"
        },
        "lineWidth": 1,
        "pointSize": 8,
        "scaleDistribution": {
          "log": 10,
          "type": "log"
        },
        "showPoints": "never",
        "spanNulls": false
      },
      "mappings": [],
      "thresholds": {
        "mode": "absolute",
        "steps": [
          {
            "color": "green",
            "value": null
          },
          {
            "color": "red",

```

```

        "value": 80
      }
    ]
  },
  "unit": "Bps"
},
"overrides": []
},
"gridPos": {
  "h": 9,
  "w": 12,
  "x": 12,
  "y": 18
},
"id": 17,
"options": {
  "graph": {},
  "legend": {
    "calcs": [],
    "displayMode": "list",
    "placement": "bottom"
  },
  "tooltipOptions": {
    "mode": "single"
  }
},
"pluginVersion": "7.5.3",
"targets": [
  {
    "exemplar": true,
    "expr": "rate(fcrepo_storage_bytes_total[1m])",
    "format": "time_series",
    "instant": false,
    "interval": "",
    "legendFormat": "{{operation}}",
    "refId": "A"
  }
],
"timeFrom": null,
"timeShift": null,
"title": "Bytes Processed",
"type": "timeseries"
}
],
"refresh": "5s",
"schemaVersion": 27,
"style": "dark",
"tags": [],
"templating": {
  "list": []
},
"time": {
  "from": "now-5m",
  "to": "now"
},
"timepicker": {
  "hidden": false,
  "refresh_intervals": [
    "5s",
    "10s",
    "30s",
    "1m",
    "5m",
    "15m",
    "30m",
    "1h",
    "2h",
    "1d"
  ]
},
"timezone": "",

```

```
"title": "Example Migration Utils Dashboard",  
"uid": "lAH8f4MNs",  
"version": 2  
}
```