

# @push.rocks/smart metrics

A package for easy collection and reporting of system and process metrics.

- [readme.md for @push.rocks/smartmetrics](#)
- [changelog.md for @push.rocks/smartmetrics](#)

# readme.md for @push.rocks/smartmetrics

Powerful system metrics collection for Node.js applications with Prometheus integration



## Issue Reporting and Security

For reporting bugs, issues, or security vulnerabilities, please visit [community.foss.global/](https://community.foss.global/). This is the central community hub for all issue reporting. Developers who sign and comply with our contribution agreement and go through identification can also get a [code.foss.global/](https://code.foss.global/) account to submit Pull Requests directly.

## What is SmartMetrics?

SmartMetrics is a comprehensive metrics collection library that monitors your Node.js application's resource usage in real-time. It tracks CPU usage, memory consumption, and system metrics across your main process and all child processes, providing insights through both JSON and Prometheus formats.

## Key Features

- **Real-time Metrics Collection** - Monitor CPU and memory usage across all processes
- **Automatic Child Process Tracking** - Aggregates metrics from main and child processes via `pidtree` + `pidusage`
- **Docker-Aware** - Detects container memory limits from cgroup automatically
- **Prometheus Integration** - Built-in HTTP endpoint for Prometheus scraping with `prom-client`
- **Flexible Output Formats** - Get metrics as JSON objects or Prometheus text exposition format
- **Automatic Heartbeat Logging** - Optional periodic metrics logging via `@push.rocks/smartlog`

- **Zero Configuration** - Works out of the box with sensible defaults

# Installation

```
pnpm install @push.rocks/smartmetrics
# or
npm install @push.rocks/smartmetrics
```

# Quick Start

```
import { SmartMetrics } from '@push.rocks/smartmetrics';
import { Smartlog } from '@push.rocks/smartlog';

// Create a logger instance
const logger = new Smartlog({
  logContext: null,
  minimumLogLevel: 'info',
});
logger.enableConsole();

// Initialize SmartMetrics
const metrics = new SmartMetrics(logger, 'my-service');

// Get metrics on-demand
const currentMetrics = await metrics.getMetrics();
console.log(`CPU Usage: ${currentMetrics.cpuUsageText}`);
console.log(`Memory: ${currentMetrics.memoryUsageText}`);

// Enable automatic heartbeat logging (every 20 seconds)
metrics.start();

// Enable Prometheus endpoint
metrics.enablePrometheusEndpoint(9090);
// Metrics now available at http://localhost:9090/metrics
```

```
// Clean shutdown
metrics.stop();
```

# Core Concepts

## Process Aggregation

SmartMetrics doesn't just monitor your main process - it automatically discovers and aggregates metrics from all child processes spawned by your application using `pidtree`. This gives you a complete picture of your application's resource footprint, not just the parent process.

## Memory Limit Detection

The library automatically detects available memory whether running on bare metal, in Docker containers, or with Node.js heap restrictions. It picks the most restrictive of:

1. **System total memory** (`os.totalmem()`)
2. **Docker cgroup limit** - supports both cgroup v2 (`/sys/fs/cgroup/memory.max`) and cgroup v1 (`/sys/fs/cgroup/memory/memory.limit_in_bytes`)
3. **V8 heap size limit** (`v8.getHeapStatistics().heap_size_limit`)

This ensures accurate percentage calculations regardless of environment.

## Dual Output Formats

- **JSON Format** (`getMetrics()`) - Ideal for application monitoring, custom dashboards, and programmatic access
- **Prometheus Format** (`getPrometheusFormattedMetrics()`) - Perfect for integration with Prometheus/Grafana monitoring stacks

## API Reference

```
new SmartMetrics(logger, sourceName)
```

Creates a new SmartMetrics instance.

Parameter	Type	Description
logger	Smartlog	A <code>@push.rocks/smartlog</code> logger instance
sourceName	string	Identifier for your service/application

## `async getMetrics():` `Promise<IMetricsSnapshot>`

Retrieves current system metrics as a structured object.

**Returns** `IMetricsSnapshot`:

```
{
  process_cpu_seconds_total: number; // Total CPU time in seconds
  nodejs_active_handles_total: number; // Always 0 (deprecated Node.js API; real values
  tracked by Prometheus default collectors)
  nodejs_active_requests_total: number; // Always 0 (deprecated Node.js API; real values
  tracked by Prometheus default collectors)
  nodejs_heap_size_total_bytes: number; // V8 heap size in bytes
  cpuPercentage: number; // Aggregated CPU usage across all child processes
  cpuUsageText: string; // Human-readable CPU usage (e.g. "12.5 %")
  memoryPercentage: number; // Memory usage as percentage of detected limit
  memoryUsageBytes: number; // Total memory in bytes across all child processes
  memoryUsageText: string; // Human-readable (e.g. "45% | 920 MB / 2 GB")
}
```

**Example:**

```
const metrics = await smartMetrics.getMetrics();
if (metrics.cpuPercentage > 80) {
  console.warn('High CPU usage detected!');
}
```

## `start(): void`

Starts automatic metrics collection and heartbeat logging every 20 seconds via the provided `Smartlog` instance.

```
smartMetrics.start();
// Logs: "sending heartbeat for my-service with metrics" every 20 seconds
```

## stop(): void

Stops automatic metrics collection, closes heartbeat loop, and shuts down any open Prometheus endpoints.

## async getPrometheusFormattedMetrics():

## Promise<string>

Returns all metrics in Prometheus text exposition format, including default Node.js collectors and custom SmartMetrics gauges.

```
const promMetrics = await smartMetrics.getPrometheusFormattedMetrics();
// Returns:
// # HELP smartmetrics_cpu_percentage Current CPU usage percentage
// # TYPE smartmetrics_cpu_percentage gauge
// smartmetrics_cpu_percentage 15.2
// # HELP smartmetrics_memory_percentage Current memory usage percentage
// # TYPE smartmetrics_memory_percentage gauge
// smartmetrics_memory_percentage 45.3
// # HELP smartmetrics_memory_usage_bytes Current memory usage in bytes
// # TYPE smartmetrics_memory_usage_bytes gauge
// smartmetrics_memory_usage_bytes 965214208
// ... plus all default Node.js metrics from prom-client
```

## enablePrometheusEndpoint(port?: number):

## void

Starts an HTTP server that exposes metrics at `/metrics` for Prometheus scraping.

Parameter	Type	Default	Description
<code>port</code>	<code>number</code>	<code>9090</code>	Port to listen on

```
smartMetrics.enablePrometheusEndpoint(3000);
// GET http://localhost:3000/metrics → Prometheus text format
```

```
// GET http://localhost:3000/anything-else → 404
```

## disablePrometheusEndpoint(): void

Gracefully shuts down the Prometheus HTTP server.

## formatBytes(bytes, decimals?): string

Utility method to convert byte values to human-readable strings.

```
smartMetrics.formatBytes(1073741824); // "1 GB"
smartMetrics.formatBytes(1536, 1);    // "1.5 KB"
```

# Use Cases

## ☐ Application Performance Monitoring

```
const metricsBefore = await smartMetrics.getMetrics();
await performHeavyOperation();
const metricsAfter = await smartMetrics.getMetrics();

console.log(`Operation consumed ${
  metricsAfter.process_cpu_seconds_total - metricsBefore.process_cpu_seconds_total
} CPU seconds`);
```

## ☐ Resource Limit Enforcement

```
async function checkResources() {
  const metrics = await smartMetrics.getMetrics();

  if (metrics.memoryPercentage > 90) {
    throw new Error('Memory usage too high, refusing new operations');
  }

  if (metrics.cpuPercentage > 95) {
```

```
    await delay(1000); // Back off when CPU is stressed
  }
}
```

## ☐☐ Prometheus + Grafana Stack

```
smartMetrics.enablePrometheusEndpoint(9090);

// prometheus.yml:
// scrape_configs:
//   - job_name: 'my-app'
//     scrape_interval: 15s
//     static_configs:
//       - targets: ['localhost:9090']
```

## ☐☐ Container Resource Monitoring

```
// Automatically detects Docker/cgroup memory limits
const metrics = await smartMetrics.getMetrics();
console.log(metrics.memoryUsageText);
// Output: "45% | 920 MB / 2 GB" (container limit detected)
```

## ☐☐ Health Check Endpoint

```
import express from 'express';

const app = express();

app.get('/health', async (req, res) => {
  const metrics = await smartMetrics.getMetrics();

  res.json({
    status: metrics.memoryPercentage < 90 ? 'healthy' : 'degraded',
    cpu: metrics.cpuUsageText,
    memory: metrics.memoryUsageText,
  });
});
```

```
});
```

## ☐☐ Graceful Restart on High Memory (PM2)

```
setInterval(async () => {  
  const metrics = await smartMetrics.getMetrics();  
  
  if (metrics.memoryPercentage > 95) {  
    console.error('Memory limit reached, requesting restart');  
    process.exit(0); // PM2 will restart the process  
  }  
}, 10000);
```

## License and Legal Information

This repository contains open-source code licensed under the MIT License. A copy of the license can be found in the [LICENSE](#) file.

**Please note:** The MIT License does not grant permission to use the trade names, trademarks, service marks, or product names of the project, except as required for reasonable and customary use in describing the origin of the work and reproducing the content of the NOTICE file.

## Trademarks

This project is owned and maintained by Task Venture Capital GmbH. The names and logos associated with Task Venture Capital GmbH and any related products or services are trademarks of Task Venture Capital GmbH or third parties, and are not included within the scope of the MIT license granted herein.

Use of these trademarks must comply with Task Venture Capital GmbH's Trademark Guidelines or the guidelines of the respective third-party owners, and any usage must be approved in writing. Third-party trademarks used herein are the property of their respective owners and used only in a descriptive manner, e.g. for an implementation of an API or similar.

## Company Information

Task Venture Capital GmbH Registered at District Court Bremen HRB 35230 HB, Germany

For any legal inquiries or further information, please contact us via email at [hello@task.vc](mailto:hello@task.vc).

By using this repository, you acknowledge that you have read this section, agree to comply with its terms, and understand that the licensing of the code does not imply endorsement by Task Venture Capital GmbH of any derivative works.

# changelog.md for @push.rocks/smartmetrics

## 2026-03-23 - 3.0.3 - fix(prometheus)

clean up default Prometheus metric collectors on stop

- Return cleanup handlers from default metric registration for event loop and GC observers.
- Dispose registered Prometheus default metric collectors when stopping Smartmetrics to prevent lingering observers.

## 2026-03-02 - 3.0.2 - fix(pidusage)

prune history entries for PIDs not present in the requested set to avoid stale data and memory growth

- Deletes entries from the history map when a PID is not included in the current pids array
- Prevents accumulation of stale PID histories and potential memory growth
- Change implemented in ts/smartmetrics.pidusage.ts alongside the metrics result construction

## 2026-02-19 - 3.0.1 - fix(smartmetrics)

no code changes detected; no version bump or release required

- git diff contained no modifications
- current package.json version is 3.0.0

- no dependency or file changes to warrant a release

## 2026-02-19 - 3.0.0 - BREAKING CHANGE(smartmetrics)

add system-wide metrics collection, Prometheus gauges, and normalized CPU reporting

- Add new sysusage plugin (ts/smartmetrics.sysusage.ts) that reads /proc/stat and /proc/meminfo (with os fallback) and returns system CPU, memory and load averages.
- Expose system-wide Prometheus gauges: smartmetrics\_system\_cpu\_percent, smartmetrics\_system\_memory\_used\_percent, smartmetrics\_system\_memory\_used\_bytes, smartmetrics\_system\_load\_avg\_1, smartmetrics\_system\_load\_avg\_5, smartmetrics\_system\_load\_avg\_15.
- Extend IMetricsSnapshot with system fields: systemCpuPercent, systemMemTotalBytes, systemMemAvailableBytes, systemMemUsedBytes, systemMemUsedPercent, systemLoadAvg1, systemLoadAvg5, systemLoadAvg15 (this is a breaking TypeScript API change).
- Normalize per-process CPU in pidusage by adding cpuCoreCount and cpuNormalizedPercent and use cpuNormalizedPercent when aggregating CPU across the process tree.
- Export the new sysusage plugin from ts/smartmetrics.plugins.ts and wire system metrics into metric collection and Prometheus gauge updates.

## 2026-02-19 - 2.0.11 - fix(deps)

bump dependencies, update build script, expand README and npm metadata

- Bumped runtime deps: @push.rocks/smartdelay ^3.0.5, @push.rocks/smartlog ^3.1.11
- Updated devDependencies: @git.zone/tsbuild, tsbundle, tsrun, tstest and @types/node versions
- Changed build script: "(tsbuild --web)" → "(tsbuild tsfolders)"
- Updated npmextra.json: renamed keys (gitzone → @git.zone/cli, tsdoc → @git.zone/tsdoc), added release registries and accessLevel, and added @ship.zone/szci entry
- Extensive README improvements: installation notes (pnpm), clearer API docs, examples, added Issue Reporting & Security section and utility docs (formatBytes)

## 2025-06-09 - 2.0.9 - fix(readme)

Update documentation with detailed usage instructions, API references and integration examples.

- Overhauled README to provide a clearer explanation of SmartMetrics features and API.
- Added a quick start guide, detailed examples, and code snippets for various integrations (Express, PM2, custom dashboards).
- Reorganized documentation sections to better highlight core concepts including process aggregation and memory limit detection.
- Updated installation instructions and usage examples to reflect the latest functionality.

## 2025-06-09 - 2.0.8 - fix(smartmetrics)

Refactor metrics calculation and update Prometheus integration documentation

- Removed dependency on `registry.getMetricsAsJSON` by directly calculating Node.js metrics
- Updated `getMetrics` to compute CPU time, heap size, and set default values for deprecated metrics
- Enhanced documentation with comprehensive Prometheus integration examples
- Improved logging on shutdown of the Prometheus endpoint

## 2025-06-09 - 2.0.7 - Prometheus Metrics Integration

feat: Implement Prometheus metrics exposure in SmartMetrics

- Added Prometheus gauges for CPU and memory metrics.
- Implemented HTTP server to expose metrics at the `/metrics` endpoint.
- Created methods to enable and disable the Prometheus endpoint.
- Updated `getMetrics()` to set gauge values.
- Added tests for Prometheus metrics functionality.
- Updated documentation plan for Prometheus integration.

# 2023-07-02 to 2023-08-08 - 2.0.0 to 2.0.6 - Maintenance Updates

Over this period several releases were published with iterative fixes and minor organizational changes.

- Applied multiple core fixes and routine maintenance updates.
- Switched to new organization scheme (recorded on 2023-07-10).
- Performed several version bumps and configuration updates.

# 2021-08-12 to 2022-07-27 - 1.0.1 to 1.0.17 - Maintenance and Breaking Changes

During this interval a series of minor fixes were combined with a significant breaking change.

- BREAKING CHANGE: Switched to ESM in 1.0.17 (2022-07-27).
- Numerous maintenance updates and core fixes were applied.