

@push.rocks/taskbuffer

A flexible task management library supporting TypeScript, allowing for task buffering, scheduling, and execution with dependency management.

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

readme.md for @push.rocks/taskbuffer

“ Modern TypeScript task orchestration and service lifecycle management with constraint-based concurrency, smart buffering, scheduling, health checks, and real-time event streaming

npm version TypeScript License: MIT

Issue Reporting and Security

For reporting bugs, issues, or security vulnerabilities, please visit 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/ account to submit Pull Requests directly.

☐ Features

- ☐ **Type-Safe Task Management** — Full TypeScript support with generics and type inference
- ☐ **Constraint-Based Concurrency** — Per-key mutual exclusion, group concurrency limits, cooldown enforcement, sliding-window rate limiting, and result sharing via `TaskConstraintGroup`
- ☐ **Real-Time Progress Tracking** — Step-based progress with percentage weights
- ☐ **Smart Buffering** — Intelligent request debouncing and batching
- ☐ **Cron Scheduling** — Schedule tasks with cron expressions
- ☐ **Task Chains & Parallel Execution** — Sequential and parallel task orchestration
- ☐ **Labels** — Attach arbitrary `Record<string, string>` metadata (userId, tenantId, etc.) for multi-tenant filtering
- ☐ **Push-Based Events** — rxjs `Subject<ITaskEvent>` on every Task and TaskManager for real-time state change notifications

- **❏ Error Handling** — Configurable error propagation with `catchErrors`, error tracking, and clear error state
- **❏ Service Lifecycle Management** — `Service` and `ServiceManager` for long-running components (databases, servers, queues) with health checks, auto-restart, dependency ordering, and instance access
- **❏ Web Component Dashboard** — Built-in Lit-based dashboard for real-time task visualization
- **❏ Distributed Coordination** — Abstract coordinator for multi-instance task deduplication

❏ Installation

```
pnpm add @push.rocks/taskbuffer
# or
npm install @push.rocks/taskbuffer
```

❏ Quick Start

Basic Task

```
import { Task } from '@push.rocks/taskbuffer';

const greetTask = new Task({
  name: 'Greet',
  taskFunction: async (name) => {
    return `Hello, ${name}!`;
  },
});

const result = await greetTask.trigger('World');
console.log(result); // "Hello, World!"
```

Task with Typed Data ❏

Every task can carry a typed data bag — perfect for constraint matching, routing, and metadata:

```
const task = new Task<undefined, [], { domain: string; priority: number }>({
  name: 'update-dns',
  data: { domain: 'example.com', priority: 1 },
  taskFunction: async () => {
    // task.data is fully typed here
    console.log(`Updating DNS for ${task.data.domain}`);
  },
});

task.data.domain; // string – fully typed
task.data.priority; // number – fully typed
```

Task with Steps & Progress □□

```
const deployTask = new Task({
  name: 'Deploy',
  steps: [
    { name: 'build', description: 'Building app', percentage: 30 },
    { name: 'test', description: 'Running tests', percentage: 20 },
    { name: 'deploy', description: 'Deploying to server', percentage: 40 },
    { name: 'verify', description: 'Verifying deployment', percentage: 10 },
  ] as const,
  taskFunction: async () => {
    deployTask.notifyStep('build');
    await buildApp();

    deployTask.notifyStep('test');
    await runTests();

    deployTask.notifyStep('deploy');
    await deployToServer();

    deployTask.notifyStep('verify');
    await verifyDeployment();

    return 'Deployment successful!';
  },
});
```

```
await deployTask.trigger();
console.log(deployTask.getProgress()); // 100
console.log(deployTask.getStepsMetadata()); // Step details with status
```

“ **Note:** `notifyStep()` is fully type-safe — TypeScript only accepts step names you declared in the `steps` array when you use `as const`.”

☐ Task Constraints — Concurrency, Mutual Exclusion & Cooldowns

`TaskConstraintGroup` is the unified mechanism for controlling how tasks run relative to each other. It replaces older patterns like task runners, blocking tasks, and execution delays with a single, composable, key-based constraint system.

Per-Key Mutual Exclusion

Ensure only one task runs at a time for a given key (e.g. per domain, per tenant, per resource):

```
import { Task, TaskManager, TaskConstraintGroup } from '@push.rocks/taskbuffer';

const manager = new TaskManager();

// Only one DNS update per domain at a time
const domainMutex = new TaskConstraintGroup<{ domain: string }>({
  name: 'domain-mutex',
  maxConcurrent: 1,
  constraintKeyForExecution: (task, input?) => task.data.domain,
});

manager.addConstraintGroup(domainMutex);

const task1 = new Task<undefined, [], { domain: string }>({
```

```

name: 'update-a.com',
data: { domain: 'a.com' },
taskFunction: async () => { /* update DNS for a.com */ },
});

const task2 = new Task<undefined, [], { domain: string }>({
  name: 'update-a.com-2',
  data: { domain: 'a.com' },
  taskFunction: async () => { /* another update for a.com */ },
});

manager.addTask(task1);
manager.addTask(task2);

// task2 waits until task1 finishes (same domain key)
await Promise.all([
  manager.triggerTask(task1),
  manager.triggerTask(task2),
]);

```

Group Concurrency Limits

Cap how many tasks can run concurrently across a group:

```

// Max 3 DNS updaters running globally at once
const dnsLimit = new TaskConstraintGroup<{ group: string }>({
  name: 'dns-concurrency',
  maxConcurrent: 3,
  constraintKeyForExecution: (task) =>
    task.data.group === 'dns' ? 'dns' : null, // null = skip constraint
});

manager.addConstraintGroup(dnsLimit);

```

Cooldowns (Rate Limiting)

Enforce a minimum time gap between consecutive executions for the same key:

```
// No more than one API call per domain every 11 seconds
const rateLimiter = new TaskConstraintGroup<{ domain: string }>({
  name: 'api-rate-limit',
  maxConcurrent: 1,
  cooldownMs: 11000,
  constraintKeyForExecution: (task) => task.data.domain,
});

manager.addConstraintGroup(rateLimiter);
```

Global Concurrency Cap

Limit total concurrent tasks system-wide:

```
const globalCap = new TaskConstraintGroup({
  name: 'global-cap',
  maxConcurrent: 10,
  constraintKeyForExecution: () => 'all', // same key = shared limit
});

manager.addConstraintGroup(globalCap);
```

Composing Multiple Constraints

Multiple constraint groups stack — a task only runs when **all** applicable constraints allow it:

```
manager.addConstraintGroup(globalCap); // max 10 globally
manager.addConstraintGroup(domainMutex); // max 1 per domain
manager.addConstraintGroup(rateLimiter); // 11s cooldown per domain

// A task must satisfy ALL three constraints before it starts
await manager.triggerTask(dnsTask);
```

Selective Constraints

Return `null` from `constraintKeyForExecution` to exempt a task from a constraint group:

```

const constraint = new TaskConstraintGroup<{ priority: string }>({
  name: 'low-priority-limit',
  maxConcurrent: 2,
  constraintKeyForExecution: (task) =>
    task.data.priority === 'low' ? 'low-priority' : null, // high priority tasks skip this
  constraint
});

```

Input-Aware Constraints □□

The `constraintKeyForExecution` function receives both the **task** and the **runtime input** passed to `trigger(input)`. This means the same task triggered with different inputs can be constrained independently:

```

const extractTLD = (domain: string) => {
  const parts = domain.split('.');
  return parts.slice(-2).join('.');
};

// Same TLD → serialized. Different TLDs → parallel.
const tldMutex = new TaskConstraintGroup({
  name: 'tld-mutex',
  maxConcurrent: 1,
  constraintKeyForExecution: (task, input?: string) => {
    if (!input) return null;
    return extractTLD(input); // "example.com", "other.org", etc.
  },
});

manager.addConstraintGroup(tldMutex);

// These two serialize (same TLD "example.com")
const p1 = manager.triggerTaskConstrained(getCert, 'app.example.com');
const p2 = manager.triggerTaskConstrained(getCert, 'api.example.com');

// This runs in parallel (different TLD "other.org")
const p3 = manager.triggerTaskConstrained(getCert, 'my.other.org');

```

You can also combine `task.data` and `input` for composite keys:

```
const providerDomain = new TaskConstraintGroup<{ provider: string }>({
  name: 'provider-domain',
  maxConcurrent: 1,
  constraintKeyForExecution: (task, input?: string) => {
    return `${task.data.provider}:${input || 'default'}`;
  },
});
```

Pre-Execution Check with `shouldExecute` □

The `shouldExecute` callback runs right before a queued task executes. If it returns `false`, the task is skipped and its promise resolves with `undefined`. This is perfect for scenarios where a prior execution's outcome makes subsequent queued tasks unnecessary:

```
const certCache = new Map<string, string>();

const certConstraint = new TaskConstraintGroup({
  name: 'cert-mutex',
  maxConcurrent: 1,
  constraintKeyForExecution: (task, input?: string) => {
    if (!input) return null;
    return extractTLD(input);
  },
  shouldExecute: (task, input?: string) => {
    if (!input) return true;
    // Skip if a wildcard cert already covers this TLD
    return certCache.get(extractTLD(input)) !== 'wildcard';
  },
});

const getCert = new Task({
  name: 'get-certificate',
  taskFunction: async (domain: string) => {
    const cert = await acme.getCert(domain);
    if (cert.isWildcard) certCache.set(extractTLD(domain), 'wildcard');
    return cert;
  },
});
```

```

manager.addConstraintGroup(certConstraint);
manager.addTask(getCert);

const r1 = manager.triggerTaskConstrained(getCert, 'app.example.com'); // runs, gets wildcard
const r2 = manager.triggerTaskConstrained(getCert, 'api.example.com'); // queued → skipped!
const r3 = manager.triggerTaskConstrained(getCert, 'my.other.org'); // parallel (different
TLD)

const [cert1, cert2, cert3] = await Promise.all([r1, r2, r3]);
// cert2 === undefined (skipped because wildcard already covers example.com)

```

`shouldExecute` semantics:

- Runs right before execution (after slot acquisition, before `trigger()`)
- Also checked on immediate (non-queued) triggers
- Returns `false` → skip execution, deferred resolves with `undefined`
- Can be async (return `Promise<boolean>`)
- Has closure access to external state modified by prior executions
- If multiple constraint groups have `shouldExecute`, **all** must return `true`

Sliding Window Rate Limiting

Enforce "N completions per time window" with burst capability. Unlike `cooldownMs` (which forces even spacing between executions), `rateLimit` allows bursts up to the cap, then blocks until the window slides:

```

// Let's Encrypt style: 300 new orders per 3 hours
const acmeRateLimit = new TaskConstraintGroup({
  name: 'acme-rate',
  constraintKeyForExecution: () => 'acme-account',
  rateLimit: {
    maxPerWindow: 300,
    windowMs: 3 * 60 * 60 * 1000, // 3 hours
  },
});

manager.addConstraintGroup(acmeRateLimit);

// All 300 can burst immediately. The 301st waits until the oldest
// completion falls out of the 3-hour window.
for (const domain of domains) {

```

```
manager.triggerTaskConstrained(certTask, { domain });
}
```

Compose multiple rate limits for layered protection:

```
// Per-domain weekly cap AND global order rate
const perDomainWeekly = new TaskConstraintGroup({
  name: 'per-domain-weekly',
  constraintKeyForExecution: (task, input) => input.registeredDomain,
  rateLimit: { maxPerWindow: 50, windowMs: 7 * 24 * 60 * 60 * 1000 },
});

const globalOrderRate = new TaskConstraintGroup({
  name: 'global-order-rate',
  constraintKeyForExecution: () => 'global',
  rateLimit: { maxPerWindow: 300, windowMs: 3 * 60 * 60 * 1000 },
});

manager.addConstraintGroup(perDomainWeekly);
manager.addConstraintGroup(globalOrderRate);
```

Combine with `maxConcurrent` and `cooldownMs` for fine-grained control:

```
const throttled = new TaskConstraintGroup({
  name: 'acme-throttle',
  constraintKeyForExecution: () => 'acme',
  maxConcurrent: 5,           // max 5 concurrent requests
  cooldownMs: 1000,         // 1s gap after each completion
  rateLimit: {
    maxPerWindow: 300,
    windowMs: 3 * 60 * 60 * 1000,
  },
});
```

Result Sharing — Deduplication for Concurrent Requests

When multiple callers request the same resource concurrently, `resultSharingMode: 'share-latest'` ensures only one execution occurs. All queued waiters receive the same result:

```

const certMutex = new TaskConstraintGroup({
  name: 'cert-per-tld',
  constraintKeyForExecution: (task, input) => extractTld(input.domain),
  maxConcurrent: 1,
  resultSharingMode: 'share-latest',
});

manager.addConstraintGroup(certMutex);

const certTask = new Task({
  name: 'obtain-cert',
  taskFunction: async (input) => {
    return await acmeClient.obtainWildcard(input.domain);
  },
});
manager.addTask(certTask);

// Three requests for *.example.com arrive simultaneously
const [cert1, cert2, cert3] = await Promise.all([
  manager.triggerTaskConstrained(certTask, { domain: 'api.example.com' }),
  manager.triggerTaskConstrained(certTask, { domain: 'www.example.com' }),
  manager.triggerTaskConstrained(certTask, { domain: 'mail.example.com' }),
]);

// Only ONE ACME request was made.
// cert1 === cert2 === cert3 – all callers got the same cert object.

```

Result sharing semantics:

- `shouldExecute` is NOT called for shared results (the task's purpose was already fulfilled)
- Error results are NOT shared — queued tasks execute independently after a failure
- `lastResults` persists until `reset()` — for time-bounded sharing, use `shouldExecute` to control staleness
- Composable with rate limiting: rate-limited waiters get shared results without waiting for the window

How It Works

When you trigger a task through `TaskManager` (via `triggerTask`, `triggerTaskByName`, `addExecuteRemoveTask`, or `cron`), the manager:

1. Evaluates all registered constraint groups against the task and input
2. If no constraints apply (all matchers return `null`) → checks `shouldExecute` → runs or skips
3. If all applicable constraints have capacity → acquires slots → checks `shouldExecute` → runs or skips
4. If any constraint blocks → enqueues the task; when a running task completes, the queue is drained
5. Cooldown/rate-limit-blocked tasks auto-retry after the shortest remaining delay expires
6. Queued tasks check for shared results first (if any group has `resultSharingMode: 'share-latest'`)
7. Queued tasks re-check `shouldExecute` when their turn comes — stale work is automatically pruned

☐ Core Concepts

Task Buffering — Intelligent Request Management

Prevent overwhelming your system with rapid-fire requests:

```
const apiTask = new Task({
  name: 'APIRequest',
  buffered: true,
  bufferMax: 5, // Maximum 5 concurrent executions
  taskFunction: async (endpoint) => {
    return await fetch(endpoint).then((r) => r.json());
  },
});

// Rapid fire 100 calls – only bufferMax execute concurrently
for (let i = 0; i < 100; i++) {
  apiTask.trigger(`/api/data/${i}`);
}
```

Buffer Behavior:

- First `bufferMax` calls execute immediately
- Additional calls are queued
- When buffer is full, new calls overwrite the last queued item
- Perfect for real-time data streams where only recent data matters

Task Chains — Sequential Workflows □□

Build complex workflows with automatic data flow between tasks:

```
import { Taskchain } from '@push.rocks/taskbuffer';

const fetchTask = new Task({
  name: 'Fetch',
  taskFunction: async () => {
    const res = await fetch('/api/data');
    return res.json();
  },
});

const transformTask = new Task({
  name: 'Transform',
  taskFunction: async (data) => {
    return data.map((item) => ({ ...item, transformed: true }));
  },
});

const saveTask = new Task({
  name: 'Save',
  taskFunction: async (transformedData) => {
    await database.save(transformedData);
    return transformedData.length;
  },
});

const pipeline = new Taskchain({
  name: 'DataPipeline',
  taskArray: [fetchTask, transformTask, saveTask],
});

const savedCount = await pipeline.trigger();
console.log(`Saved ${savedCount} items`);
```

Taskchain also supports dynamic mutation:

```
pipeline.addTask(newTask); // Append to chain
pipeline.removeTask(oldTask); // Remove by reference (returns boolean)
pipeline.shiftTask(); // Remove & return first task
```

Error context is rich — a chain failure includes the chain name, failing task name, task index, and preserves the original error as `.cause`.

Parallel Execution — Concurrent Processing ⚡

Execute multiple tasks simultaneously:

```
import { Taskparallel } from '@push.rocks/taskbuffer';

const parallel = new Taskparallel({
  taskArray: [emailTask, smsTask, pushNotificationTask, webhookTask],
});

await parallel.trigger(notificationData);
```

Debounced Tasks — Smart Trigger Coalescing 📦

Coalesce rapid triggers into a single execution after a quiet period:

```
import { TaskDebounced } from '@push.rocks/taskbuffer';

const searchTask = new TaskDebounced({
  name: 'Search',
  debounceTimeInMillis: 300,
  taskFunction: async (query) => {
    return await searchAPI(query);
  },
});

// Rapid calls – only the last triggers after 300ms of quiet
```

```
searchTask.trigger('h');
searchTask.trigger('he');
searchTask.trigger('hel');
searchTask.trigger('hello'); // ← this one fires
```

TaskOnce — Single-Execution Guard

Ensure a task only runs once, regardless of how many times it's triggered:

```
import { TaskOnce } from '@push.rocks/taskbuffer';

const initTask = new TaskOnce({
  name: 'Init',
  taskFunction: async () => {
    await setupDatabase();
    console.log('Initialized!');
  },
});

await initTask.trigger(); // Runs
await initTask.trigger(); // No-op
await initTask.trigger(); // No-op
console.log(initTask.hasTriggered); // true
```

☐ Labels — Multi-Tenant Task Filtering

Attach arbitrary key-value labels to any task for filtering, grouping, or multi-tenant isolation:

```
const task = new Task({
  name: 'ProcessOrder',
  labels: { userId: 'u-42', tenantId: 'acme-corp', priority: 'high' },
  taskFunction: async (order) => {
    /* ... */
  },
});
```

```
// Manipulate labels at runtime
task.setLabel('region', 'eu-west');
task.getLabel('userId'); // 'u-42'
task.hasLabel('tenantId', 'acme-corp'); // true
task.removeLabel('priority'); // true

// Labels are included in metadata snapshots
const meta = task.getMetadata();
console.log(meta.labels); // { userId: 'u-42', tenantId: 'acme-corp', region: 'eu-west' }
```

Filtering Tasks by Label in TaskManager

```
const manager = new TaskManager();
manager.addTask(orderTask1); // labels: { tenantId: 'acme' }
manager.addTask(orderTask2); // labels: { tenantId: 'globex' }
manager.addTask(orderTask3); // labels: { tenantId: 'acme' }

const acmeTasks = manager.getTasksByLabel('tenantId', 'acme');
// → [orderTask1, orderTask3]

const acmeMetadata = manager.getTasksMetadataByLabel('tenantId', 'acme');
// → [ITaskMetadata, ITaskMetadata]
```

☐ Push-Based Events — Real-Time Task Lifecycle

Every `Task` exposes an rxjs `Subject<ITaskEvent>` that emits events as the task progresses through its lifecycle:

```
import type { ITaskEvent } from '@push.rocks/taskbuffer';

const task = new Task({
  name: 'DataSync',
  steps: [
```

```

    { name: 'fetch', description: 'Fetching data', percentage: 50 },
    { name: 'save', description: 'Saving data', percentage: 50 },
  ] as const,
  taskFunction: async () => {
    task.notifyStep('fetch');
    const data = await fetchData();
    task.notifyStep('save');
    await saveData(data);
  },
});

// Subscribe to individual task events
task.eventSubject.subscribe((event: ITaskEvent) => {
  console.log(`[${event.type}] ${event.task.name} @ ${new
Date(event.timestamp).toISOString()}`);
  if (event.type === 'step') console.log(` Step: ${event.stepName}`);
  if (event.type === 'failed') console.log(` Error: ${event.error}`);
});

await task.trigger();
// [started] DataSync @ 2025-01-26T...
// [step] DataSync @ 2025-01-26T...
// Step: fetch
// [step] DataSync @ 2025-01-26T...
// Step: save
// [completed] DataSync @ 2025-01-26T...

```

Event Types

Type	When	Extra Fields
'started'	Task begins execution	—
'step'	<code>notifyStep()</code> is called	<code>stepName</code>
'completed'	Task finishes successfully	—
'failed'	Task throws an error	<code>error</code> (message string)

Every event includes a full `ITaskMetadata` snapshot (including labels) at the time of emission.

Aggregated Events on TaskManager

`TaskManager` automatically aggregates events from all added tasks into a single `taskSubject`:

```
const manager = new TaskManager();
manager.addTask(syncTask);
manager.addTask(reportTask);
manager.addTask(cleanupTask);

// Single subscription for ALL task events
manager.taskSubject.subscribe((event) => {
  sendToMonitoringDashboard(event);
});

// Events stop flowing for a task after removal
manager.removeTask(syncTask);
```

`manager.stop()` automatically cleans up all event subscriptions.

☐ Error Handling

By default, `trigger()` **rejects** when the task function throws — errors propagate naturally:

```
const task = new Task({
  name: 'RiskyOp',
  taskFunction: async () => {
    throw new Error('something broke');
  },
});

try {
  await task.trigger();
} catch (err) {
  console.error(err.message); // "something broke"
}
```

Swallowing Errors with `catchErrors`

Set `catchErrors: true` to swallow errors and return `undefined` instead of rejecting:

```
const task = new Task({
  name: 'BestEffort',
  catchErrors: true,
  taskFunction: async () => {
    throw new Error('non-critical');
  },
});

const result = await task.trigger(); // undefined (no throw)
```

Error State Tracking

Regardless of `catchErrors`, the task tracks errors:

```
console.log(task.lastError); // Error object (or undefined)
console.log(task.errorCount); // Number of failures across all runs
console.log(task.getMetadata().status); // 'failed'

task.clearError(); // Resets lastError to undefined (errorCount stays)
```

On a subsequent successful run, `lastError` is automatically cleared.

☐ TaskManager — Centralized Orchestration

```
const manager = new TaskManager();

// Add tasks
manager.addTask(dataProcessor);
manager.addTask(deployTask);

// Schedule with cron expressions
manager.addAndScheduleTask(backupTask, '0 2 * * *'); // Daily at 2 AM
manager.addAndScheduleTask(healthCheck, '*/* 5 * * * *'); // Every 5 minutes
```

```
// Register constraint groups
manager.addConstraintGroup(globalCap);
manager.addConstraintGroup(perDomainMutex);

// Query metadata
const meta = manager.getTaskMetadata('Deploy');
console.log(meta);
// {
//   name: 'Deploy',
//   status: 'completed',
//   steps: [...],
//   currentProgress: 100,
//   runCount: 3,
//   labels: { env: 'production' },
//   lastError: undefined,
//   errorCount: 0,
//   ...
// }

// All tasks at once
const allMeta = manager.getAllTasksMetadata();

// Scheduled task info
const scheduled = manager.getScheduledTasks();
const nextRuns = manager.getNextScheduledRuns(5);

// Trigger by name (routes through constraints)
await manager.triggerTaskByName('Deploy');

// One-shot: add, execute, collect report, remove
const report = await manager.addExecuteRemoveTask(temporaryTask);
console.log(report);
// {
//   taskName: 'TempTask',
//   startTime: 1706284800000,
//   endTime: 1706284801523,
//   duration: 1523,
//   steps: [...],
```

```
//  stepsCompleted: ['step1', 'step2'],
//  progress: 100,
//  result: any
// }

// Lifecycle
await manager.start(); // Starts cron scheduling + distributed coordinator
await manager.stop(); // Stops scheduling, cleans up event subscriptions
```

Remove Tasks

```
manager.removeTask(task); // Removes from map and unsubscribes event forwarding
manager.descheduleTaskByName('Deploy'); // Remove cron schedule only
```

Remove Constraint Groups

```
manager.removeConstraintGroup('domain-mutex'); // By name
```

☐☐ Service Lifecycle Management

For long-running components like database connections, HTTP servers, and message queues, taskbuffer provides `Service` and `ServiceManager` — a complete lifecycle management system with health checks, dependency ordering, retry, auto-restart, and typed instance access.

Basic Service — Builder Pattern

```
import { Service, ServiceManager } from '@push.rocks/taskbuffer';

const dbService = new Service<DatabasePool>('Database')
  .critical()
  .withStart(async () => {
    const pool = new DatabasePool({ host: 'localhost', port: 5432 });
    await pool.connect();
    return pool; // stored as service.instance
  })
```

```

.withStop(async (pool) => {
  await pool.disconnect(); // receives the instance from start
})
.withHealthCheck(async (pool) => {
  return await pool.ping(); // receives the instance too
});

await dbService.start();
dbService.instance!.query('SELECT 1'); // typed access to the pool
await dbService.stop();

```

The `start()` return value is stored as `service.instance` and automatically passed to `stop()` and `healthCheck()` functions — no need for external closures or shared variables.

Service with Dependencies & Health Checks

```

const cacheService = new Service('Redis')
  .optional()
  .withStart(async () => new RedisClient())
  .withStop(async (client) => client.quit())
  .withHealthCheck(async (client) => client.isReady, {
    intervalMs: 10000, // check every 10s
    timeoutMs: 3000, // 3s timeout per check
    failuresBeforeDegraded: 3, // 3 consecutive failures → 'degraded'
    failuresBeforeFailed: 5, // 5 consecutive failures → 'failed'
    autoRestart: true, // auto-restart when failed
    maxAutoRestarts: 5, // give up after 5 restart attempts
    autoRestartDelayMs: 2000, // start with 2s delay
    autoRestartBackoffFactor: 2, // double delay each attempt
  });

const apiService = new Service('API')
  .critical()
  .dependsOn('Database', 'Redis')
  .withStart(async () => {
    const server = createServer();
    await server.listen(3000);
  });

```

```
    return server;
  })
  .withStop(async (server) => server.close())
  .withStartupTimeout(10000); // fail if start takes > 10s
```

ServiceManager — Orchestration

`ServiceManager` handles dependency-ordered startup, failure isolation, and aggregated health reporting:

```
const manager = new ServiceManager({
  name: 'MyApp',
  startupTimeoutMs: 60000, // global startup timeout
  shutdownTimeoutMs: 15000, // per-service shutdown timeout
  defaultRetry: { maxRetries: 3, baseDelayMs: 1000, backoffFactor: 2 },
});

manager.addService(dbService);
manager.addService(cacheService);
manager.addService(apiService);

await manager.start();
// □ Starts Database first, then Redis (parallel with DB since independent),
//   then API (after both deps are running)
// □ If Database (critical) fails → rollback, stop everything, throw
// △□ If Redis (optional) fails → log warning, continue, health = 'degraded'

// Health aggregation
const health = manager.getHealth();
// { overall: 'healthy', services: [...], startedAt: 1706284800000, uptime: 42000 }

// Cascade restart – stops dependents first, restarts target, then restarts dependents
await manager.restartService('Database');

// Graceful reverse-order shutdown
await manager.stop();
```

Subclass Pattern

For complex services, extend `Service` and override the lifecycle hooks:

```
class PostgresService extends Service<Pool> {
  constructor(private config: PoolConfig) {
    super('Postgres');
    this.critical();
  }

  protected async serviceStart(): Promise<Pool> {
    const pool = new Pool(this.config);
    await pool.connect();
    return pool;
  }

  protected async serviceStop(): Promise<void> {
    await this.instance?.end();
  }

  protected async serviceHealthCheck(): Promise<boolean> {
    const result = await this.instance?.query('SELECT 1');
    return result?.rows.length === 1;
  }
}
```

Waiting for Service Readiness

Programmatically wait for a service to reach a specific state:

```
// Wait for the service to be running (with timeout)
await dbService.waitForRunning(10000);

// Wait for any state
await service.waitForState(['running', 'degraded'], 5000);

// Wait for shutdown
await service.waitForStopped();
```

Service Labels

Tag services with metadata for filtering and grouping:

```
const service = new Service('Redis')
  .withLabels({ type: 'cache', env: 'production', region: 'eu-west' })
  .withStart(async () => new RedisClient())
  .withStop(async (client) => client.quit());

// Query by label in ServiceManager
const caches = manager.getServicesByLabel('type', 'cache');
const prodStatuses = manager.getServicesStatusByLabel('env', 'production');
```

Service Events

Every `Service` emits events via an rxjs `Subject<IServiceEvent>`:

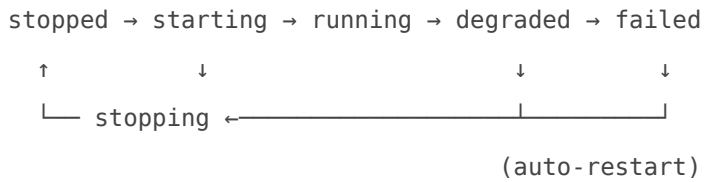
```
service.eventSubject.subscribe((event) => {
  console.log(`[${event.type}] ${event.serviceName} → ${event.state}`);
});

// [started] Database → running
// [healthCheck] Database → running
// [degraded] Database → degraded
// [autoRestarting] Database → failed
// [started] Database → running
// [recovered] Database → running
// [stopped] Database → stopped
```

Event Type	When
'started'	Service started successfully
'stopped'	Service stopped
'failed'	Service start failed or health check threshold exceeded
'degraded'	Health check failures exceeded <code>failuresBeforeDegraded</code>
'recovered'	Health check succeeded while in degraded state
'retrying'	ServiceManager retrying a failed start attempt
'healthCheck'	Health check completed (success or failure)
'autoRestarting'	Auto-restart scheduled after health check failure

`ServiceManager.serviceSubject` aggregates events from all registered services.

Service State Machine



Web Component Dashboard

Visualize your tasks in real-time with the included Lit-based web component:

```
<script type="module">
  import { TaskManager } from '@push.rocks/taskbuffer';
  import '@push.rocks/taskbuffer/dist_ts_web/taskbuffer-dashboard.js';

  const manager = new TaskManager();
  // ... add and schedule tasks ...

  const dashboard = document.querySelector('taskbuffer-dashboard');
  dashboard.taskManager = manager;
  dashboard.refreshInterval = 500; // Poll every 500ms
</script>

<taskbuffer-dashboard></taskbuffer-dashboard>
```

The dashboard provides:

- Real-time progress bars with step indicators
- Task execution history and metadata
- Scheduled task information with next-run times
- Light/dark theme support

Distributed Coordination

For multi-instance deployments, extend `AbstractDistributedCoordinator` to prevent duplicate task execution:

```

import { TaskManager, distributedCoordination } from '@push.rocks/taskbuffer';

class RedisCoordinator extends distributedCoordination.AbstractDistributedCoordinator {
  async fireDistributedTaskRequest(request) {
    // Implement leader election / distributed lock via Redis
    return { shouldTrigger: true, considered: true, rank: 1, reason: 'elected', ...request };
  }
  async updateDistributedTaskRequest(request) {
    /* update status */
  }
  async start() {
    /* connect */
  }
  async stop() {
    /* disconnect */
  }
}

const manager = new TaskManager({
  distributedCoordinator: new RedisCoordinator(),
});

```

When a distributed coordinator is configured, scheduled tasks consult it before executing — only the elected instance runs the task.

☐☐ Advanced Patterns

Pre-Task & After-Task Hooks

Run setup/teardown tasks automatically:

```

const mainTask = new Task({
  name: 'MainWork',
  preTask: new Task({
    name: 'Setup',
    taskFunction: async () => {
      console.log('Setting up...');
    }
  })
});

```

```

    },
  },
  afterTask: new Task({
    name: 'Cleanup',
    taskFunction: async () => {
      console.log('Cleaning up...');
    },
  }),
  taskFunction: async () => {
    console.log('Doing work...');
    return 'done';
  },
});

await mainTask.trigger();
// Setting up... → Doing work... → Cleaning up...

```

One-Time Setup Functions

Run an expensive initialization exactly once, before the first execution:

```

const task = new Task({
  name: 'DBQuery',
  taskSetup: async () => {
    const pool = await createConnectionPool();
    return pool; // This becomes `setupValue`
  },
  taskFunction: async (input, pool) => {
    return await pool.query(input);
  },
});

await task.trigger('SELECT * FROM users'); // Setup runs here
await task.trigger('SELECT * FROM orders'); // Setup skipped, pool reused

```

Database Migration Pipeline

```

const migration = new Taskchain({
  name: 'DatabaseMigration',
  taskArray: [backupTask, validateSchemaTask, runMigrationsTask, verifyIntegrityTask],
});

try {
  await migration.trigger();
  console.log('Migration successful!');
} catch (error) {
  // error includes chain name, failing task name, index, and original cause
  console.error(error.message);
  await rollbackTask.trigger();
}

```

Multi-Tenant SaaS Monitoring

Combine labels + events + constraints for a real-time multi-tenant system:

```

const manager = new TaskManager();

// Per-tenant concurrency limit
const tenantLimit = new TaskConstraintGroup<{ tenantId: string }>({
  name: 'tenant-concurrency',
  maxConcurrent: 2,
  constraintKeyForExecution: (task, input?) => task.data.tenantId,
});
manager.addConstraintGroup(tenantLimit);

// Create tenant-scoped tasks
function createTenantTask(tenantId: string, taskName: string, fn: () => Promise<any>) {
  const task = new Task<undefined, [], { tenantId: string }>({
    name: `${tenantId}:${taskName}`,
    data: { tenantId },
    labels: { tenantId },
    taskFunction: fn,
  });
  manager.addTask(task);
  return task;
}

```

```

}

createTenantTask('acme', 'sync', async () => syncData('acme'));
createTenantTask('globex', 'sync', async () => syncData('globex'));

// Stream events to tenant-specific WebSocket channels
manager.taskSubject.subscribe((event) => {
  const tenantId = event.task.labels?.tenantId;
  if (tenantId) {
    wss.broadcast(tenantId, JSON.stringify(event));
  }
});

// Query tasks for a specific tenant
const acmeTasks = manager.getTasksMetadataByLabel('tenantId', 'acme');

```

API Reference

Classes

Class	Description
<code>Task<T, TSteps, TData></code>	Core task unit with typed data, optional step tracking, labels, and event streaming
<code>TaskManager</code>	Centralized orchestrator with constraint groups, scheduling, label queries, and aggregated events
<code>TaskConstraintGroup<TData></code>	Concurrency, mutual exclusion, and cooldown constraints with key-based grouping
<code>Taskchain</code>	Sequential task executor with data flow between tasks
<code>Taskparallel</code>	Concurrent task executor via <code>Promise.all()</code>
<code>TaskOnce</code>	Single-execution guard
<code>TaskDebounce</code>	Debounce task using rxjs
<code>TaskStep</code>	Step tracking unit (internal, exposed via metadata)
<code>Service<T></code>	Long-running component with start/stop lifecycle, health checks, auto-restart, and typed instance access
<code>ServiceManager</code>	Service orchestrator with dependency ordering, failure isolation, retry, and health aggregation

Task Constructor Options

Option	Type	Default	Description
<code>taskFunction</code>	<code>ITaskFunction<T></code>	<i>required</i>	The async function to execute
<code>name</code>	<code>string</code>	—	Task identifier (required for TaskManager)
<code>data</code>	<code>TData</code>	<code>{}</code>	Typed data bag for constraint matching and routing
<code>steps</code>	<code>ReadonlyArray<{name, description, percentage}></code>	—	Step definitions for progress tracking
<code>buffered</code>	<code>boolean</code>	—	Enable request buffering
<code>bufferMax</code>	<code>number</code>	—	Max buffered calls
<code>preTask</code>	<code>Task () => Task</code>	—	Task to run before
<code>afterTask</code>	<code>Task () => Task</code>	—	Task to run after
<code>taskSetup</code>	<code>() => Promise<T></code>	—	One-time setup function
<code>catchErrors</code>	<code>boolean</code>	<code>false</code>	Swallow errors instead of rejecting
<code>labels</code>	<code>Record<string, string></code>	<code>{}</code>	Initial labels

Task Methods

Method	Returns	Description
<code>trigger(input?)</code>	<code>Promise<any></code>	Execute the task
<code>notifyStep(name)</code>	<code>void</code>	Advance to named step (type-safe)
<code>getProgress()</code>	<code>number</code>	Current progress 0-100
<code>getStepsMetadata()</code>	<code>ITaskStep[]</code>	Step details with status
<code>getMetadata()</code>	<code>ITaskMetadata</code>	Full task metadata snapshot
<code>setLabel(key, value)</code>	<code>void</code>	Set a label
<code>getLabel(key)</code>	<code>string undefined</code>	Get a label value
<code>removeLabel(key)</code>	<code>boolean</code>	Remove a label
<code>hasLabel(key, value?)</code>	<code>boolean</code>	Check label existence / value
<code>clearError()</code>	<code>void</code>	Reset <code>lastError</code> to undefined

Task Properties

Property	Type	Description
<code>name</code>	<code>string</code>	Task identifier
<code>data</code>	<code>TData</code>	Typed data bag
<code>running</code>	<code>boolean</code>	Whether the task is currently executing
<code>idle</code>	<code>boolean</code>	Inverse of <code>running</code>
<code>labels</code>	<code>Record<string, string></code>	Attached labels
<code>eventSubject</code>	<code>Subject<ITaskEvent></code>	rxjs Subject emitting lifecycle events
<code>lastError</code>	<code>Error undefined</code>	Last error encountered
<code>errorCount</code>	<code>number</code>	Total error count across all runs
<code>runCount</code>	<code>number</code>	Total execution count
<code>lastRun</code>	<code>Date undefined</code>	Timestamp of last execution

TaskConstraintGroup Constructor Options

Option	Type	Default	Description
<code>name</code>	<code>string</code>	<i>required</i>	Constraint group identifier
<code>constraintKeyForExecution</code>	<code>(task, input?) => string null</code>	<i>required</i>	Returns key for grouping, or <code>null</code> to skip. Receives both the task and runtime input.
<code>maxConcurrent</code>	<code>number</code>	<code>Infinity</code>	Max concurrent tasks per key
<code>cooldownMs</code>	<code>number</code>	<code>0</code>	Minimum ms between completions per key
<code>shouldExecute</code>	<code>(task, input?) => boolean Promise<boolean></code>	—	Pre-execution check. Return <code>false</code> to skip; deferred resolves <code>undefined</code> .
<code>rateLimit</code>	<code>IRateLimitConfig</code>	—	Sliding window: <code>{ maxPerWindow, windowMs }</code> . Counts running + completed tasks.
<code>resultSharingMode</code>	<code>TResultSharingMode</code>	<code>'none'</code>	<code>'none'</code> or <code>'share-latest'</code> . Queued tasks get first task's result without executing.

TaskConstraintGroup Methods

Method	Returns	Description
<code>getConstraintKey(task, input?)</code>	<code>string null</code>	Get the constraint key for a task + input
<code>checkShouldExecute(task, input?)</code>	<code>Promise<boolean></code>	Run the <code>shouldExecute</code> callback (defaults to <code>true</code>)
<code>canRun(key)</code>	<code>boolean</code>	Check if a slot is available (considers concurrency, cooldown, and rate limit)
<code>acquireSlot(key)</code>	<code>void</code>	Claim a running slot
<code>releaseSlot(key)</code>	<code>void</code>	Release a slot and record completion time + rate-limit timestamp
<code>getCooldownRemaining(key)</code>	<code>number</code>	Milliseconds until cooldown expires
<code>getRateLimitDelay(key)</code>	<code>number</code>	Milliseconds until a rate-limit slot opens
<code>getNextAvailableDelay(key)</code>	<code>number</code>	Max of cooldown + rate-limit delay — unified "when can I run"
<code>getRunningCount(key)</code>	<code>number</code>	Current running count for key
<code>recordResult(key, result)</code>	<code>void</code>	Store result for sharing (no-op if mode is <code>'none'</code>)
<code>getLastResult(key)</code>	<code>{result, timestamp} undefined</code>	Get last shared result for key
<code>hasResultSharing()</code>	<code>boolean</code>	Whether result sharing is enabled
<code>reset()</code>	<code>void</code>	Clear all state (running counts, cooldowns, rate-limit timestamps, shared results)

TaskManager Methods

Method	Returns	Description
<code>addTask(task)</code>	<code>void</code>	Register a task (wires event forwarding)
<code>removeTask(task)</code>	<code>void</code>	Remove task and unsubscribe events
<code>getTaskByName(name)</code>	<code>Task undefined</code>	Look up by name
<code>triggerTaskByName(name)</code>	<code>Promise<any></code>	Trigger by name (routes through constraints)
<code>triggerTask(task)</code>	<code>Promise<any></code>	Trigger directly (routes through constraints)

Method	Returns	Description
<code>triggerTaskConstrained(task, input?)</code>	<code>Promise<any></code>	Core constraint evaluation entry point
<code>addConstraintGroup(group)</code>	<code>void</code>	Register a constraint group
<code>removeConstraintGroup(name)</code>	<code>void</code>	Remove a constraint group by name
<code>addAndScheduleTask(task, cron)</code>	<code>void</code>	Register + schedule
<code>scheduleTaskByName(name, cron)</code>	<code>void</code>	Schedule existing task
<code>descheduleTaskByName(name)</code>	<code>void</code>	Remove schedule
<code>getTaskMetadata(name)</code>	<code>ITaskMetadata null</code>	Single task metadata
<code>getAllTasksMetadata()</code>	<code>ITaskMetadata[]</code>	All tasks metadata
<code>getScheduledTasks()</code>	<code>IScheduledTaskInfo[]</code>	Scheduled task info
<code>getNextScheduledRuns(limit?)</code>	<code>Array<{...}></code>	Upcoming scheduled runs
<code>getTasksByLabel(key, value)</code>	<code>Task[]</code>	Filter tasks by label
<code>getTasksMetadataByLabel(key, value)</code>	<code>ITaskMetadata[]</code>	Filter metadata by label
<code>addExecuteRemoveTask(task, opts?)</code>	<code>Promise<ITaskExecutionReport></code>	One-shot execution with report
<code>start()</code>	<code>Promise<void></code>	Start cron + coordinator
<code>stop()</code>	<code>Promise<void></code>	Stop cron + clean up subscriptions

TaskManager Properties

Property	Type	Description
<code>taskSubject</code>	<code>Subject<ITaskEvent></code>	Aggregated events from all added tasks
<code>taskMap</code>	<code>ObjectMap<Task></code>	Internal task registry
<code>constraintGroups</code>	<code>TaskConstraintGroup[]</code>	Registered constraint groups

Service Builder Methods

Method	Returns	Description
<code>critical()</code>	<code>this</code>	Mark as critical (startup failure aborts ServiceManager)
<code>optional()</code>	<code>this</code>	Mark as optional (startup failure is tolerated)
<code>dependsOn(...names)</code>	<code>this</code>	Declare dependencies by service name

Method	Returns	Description
<code>withStart(fn)</code>	<code>this</code>	Set start function: <code>() => Promise<T></code>
<code>withStop(fn)</code>	<code>this</code>	Set stop function: <code>(instance: T) => Promise<void></code>
<code>withHealthCheck(fn, config?)</code>	<code>this</code>	Set health check: <code>(instance: T) => Promise<boolean></code>
<code>withRetry(config)</code>	<code>this</code>	Set retry config: <code>{ maxRetries, baseDelayMs, maxDelayMs, backoffFactor }</code>
<code>withStartupTimeout(ms)</code>	<code>this</code>	Per-service startup timeout
<code>withLabels(labels)</code>	<code>this</code>	Attach key-value labels

Service Methods

Method	Returns	Description
<code>start()</code>	<code>Promise<T></code>	Start the service (no-op if already running)
<code>stop()</code>	<code>Promise<void></code>	Stop the service (no-op if already stopped)
<code>checkHealth()</code>	<code>Promise<boolean undefined></code>	Run health check manually
<code>waitForState(target, timeoutMs?)</code>	<code>Promise<void></code>	Wait for service to reach a state
<code>waitForRunning(timeoutMs?)</code>	<code>Promise<void></code>	Wait for <code>'running'</code> state
<code>waitForStopped(timeoutMs?)</code>	<code>Promise<void></code>	Wait for <code>'stopped'</code> state
<code>getStatus()</code>	<code>IServiceStatus</code>	Full status snapshot
<code>setLabel(key, value)</code>	<code>void</code>	Set a label
<code>getLabel(key)</code>	<code>string undefined</code>	Get a label value
<code>removeLabel(key)</code>	<code>boolean</code>	Remove a label
<code>hasLabel(key, value?)</code>	<code>boolean</code>	Check label existence / value

Service Properties

Property	Type	Description
<code>name</code>	<code>string</code>	Service identifier
<code>state</code>	<code>TServiceState</code>	Current state (<code>stopped</code> , <code>starting</code> , <code>running</code> , <code>degraded</code> , <code>failed</code> , <code>stopping</code>)

Property	Type	Description
<code>instance</code>	<code>T undefined</code>	The value returned from <code>start()</code>
<code>criticality</code>	<code>TServiceCriticality</code>	<code>'critical'</code> or <code>'optional'</code>
<code>dependencies</code>	<code>string[]</code>	Dependency names
<code>labels</code>	<code>Record<string, string></code>	Attached labels
<code>eventSubject</code>	<code>Subject<IServiceEvent></code>	rxjs Subject emitting lifecycle events
<code>errorCount</code>	<code>number</code>	Total error count
<code>retryCount</code>	<code>number</code>	Retry attempts during last startup

ServiceManager Methods

Method	Returns	Description
<code>addService(service)</code>	<code>void</code>	Register a service
<code>addServiceFromOptions(options)</code>	<code>Service<T></code>	Create and register from options
<code>removeService(name)</code>	<code>void</code>	Remove service (throws if others depend on it)
<code>start()</code>	<code>Promise<void></code>	Start all services in dependency order
<code>stop()</code>	<code>Promise<void></code>	Stop all services in reverse order
<code>restartService(name)</code>	<code>Promise<void></code>	Cascade restart with dependents
<code>getService(name)</code>	<code>Service undefined</code>	Look up by name
<code>getServiceStatus(name)</code>	<code>IServiceStatus undefined</code>	Single service status
<code>getAllStatuses()</code>	<code>IServiceStatus[]</code>	All service statuses
<code>getHealth()</code>	<code>IServiceManagerHealth</code>	Aggregated health report
<code>getServicesByLabel(key, value)</code>	<code>Service[]</code>	Filter services by label
<code>getServicesStatusByLabel(key, value)</code>	<code>IServiceStatus[]</code>	Filter statuses by label

Exported Types

```
import type {
  // Task types
  ITaskMetadata,
  ITaskExecutionReport,
  ITaskExecution,
```

```
IScheduledTaskInfo,  
ITaskEvent,  
TTaskEventType,  
ITaskStep,  
ITaskFunction,  
ITaskConstraintGroupOptions,  
IRateLimitConfig,  
TResultSharingMode,  
StepNames,  
// Service types  
IServiceOptions,  
IServiceStatus,  
IServiceEvent,  
IServiceManagerOptions,  
IServiceManagerHealth,  
IRetryConfig,  
IHealthCheckConfig,  
TServiceState,  
TServiceCriticality,  
TServiceEventType,  
TOverallHealth,  
} from '@push.rocks/taskbuffer';
```

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.

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/taskbuffer

2026-03-23 - 8.0.2 - fix(servicemanager)

cancel startup timeout once service initialization completes

- Replaces the startup timeout race delay with a cancellable Timeout instance
- Prevents the global startup timeout from lingering after startup finishes or fails

2026-03-23 - 8.0.1 - fix(servicemanager)

cancel shutdown timeouts after services stop

- Replace the shutdown race delay with a cancellable Timeout in ServiceManager.
- Prevent timeout handlers from lingering after a service stops successfully during shutdown.

2026-03-20 - 8.0.0 - BREAKING CHANGE(service)

expand service lifecycle management with instance-aware hooks, startup timeouts, labels, readiness waits, and auto-restart support

- Change service stop and health check callbacks to receive the started instance and expose it via `service.instance`
- Add per-service and global startup timeout handling plus `waitForState`, `waitForRunning`, and `waitForStopped` readiness helpers
- Support service labels, label-based manager queries, and auto-restart lifecycle events with configurable backoff

2026-02-15 - 6.1.2 - fix(deps)

bump `@push.rocks/smarttime` to `^4.2.3`

- Updated `@push.rocks/smarttime` from `^4.1.1` to `^4.2.3`
- Non-breaking dependency version bump; increment patch version

2026-02-15 - 6.1.1 - fix(tests)

improve buffered task tests: add chain, concurrency and queue behavior tests

- Replace `tools.delayFor` with `@push.rocks/smartdelay` for more deterministic timing in tests
- Add tests for `afterTask` chaining, `bufferMax` concurrency, `queued-run` limits, and re-trigger behavior
- Rename tasks to descriptive names and fix `afterTask` chaining order to avoid circular references
- Change test runner invocation to export default `tap.start()` instead of calling `tap.start()` directly

2026-02-15 - 6.1.0 - feat(taskbuffer)

add sliding-window rate limiting and result-sharing to `TaskConstraintGroup` and integrate with `TaskManager`

- Added `IRateLimitConfig` and `TResultSharingMode` types and exported them from the public index
- `TaskConstraintGroup`: added `rateLimit` and `resultSharingMode` options, internal completion timestamp tracking, and last-result storage

- TaskConstraintGroup: new helpers - pruneCompletionTimestamps, getRateLimitDelay, getNextAvailableDelay, recordResult, getLastResult, hasResultSharing
- TaskConstraintGroup: rate-limit logic enforces maxPerWindow (counts running + completions) and composes with cooldown/maxConcurrent
- TaskManager: records successful task results to constraint groups and resolves queued entries immediately when a shared result exists
- TaskManager: queue drain now considers unified next-available delay (cooldown + rate limit) when scheduling retries
- Documentation updated: README and hints with usage examples for sliding-window rate limiting and result sharing
- Comprehensive tests added for rate limiting, concurrency interaction, and result-sharing behavior

2026-02-15 - 6.0.1 - fix(taskbuffer)

no changes to commit

- Git diff shows no changes
- package.json current version is 6.0.0; no version bump required

2026-02-15 - 6.0.0 - BREAKING CHANGE(constraints)

make TaskConstraintGroup constraint matcher input-aware and add shouldExecute pre-execution hook

- Rename ITaskConstraintGroupOptions.constraintKeyForTask -> constraintKeyForExecution(task, input?) and update TaskConstraintGroup.getConstraintKey signature
- Add optional shouldExecute(task, input?) hook; TaskManager checks shouldExecute before immediate runs, after acquiring slots, and when draining the constraint queue (queued tasks are skipped when shouldExecute returns false)
- Export ITaskExecution type and store constraintKeys on queued entries (IConstrainedTaskEntry.constraintKeys)
- Documentation and tests updated to demonstrate input-aware constraint keys and shouldExecute pruning

2026-02-15 - 5.0.1 - fix(tests)

add and tighten constraint-related tests covering return values, error propagation, concurrency, cooldown timing, and constraint removal

- Tightened cooldown timing assertion from $\geq 100\text{ms}$ to $\geq 250\text{ms}$ to reflect 300ms cooldown with 50ms tolerance.
- Added tests for queued task return values, error propagation when `catchErrors` is false, and error swallowing behavior when `catchErrors` is true.
- Added concurrency and cooldown interaction tests to ensure `maxConcurrent` is respected and batch timing is correct.
- Added test verifying removing a constraint group unblocks queued tasks and drain behavior completes correctly.

2026-02-15 - 5.0.0 - BREAKING CHANGE(taskbuffer)

Introduce constraint-based concurrency with `TaskConstraintGroup` and `TaskManager` integration; remove legacy `TaskRunner` and several `Task` APIs (breaking); add typed `Task.data` and update exports and tests.

- Add `TaskConstraintGroup` class with per-key `maxConcurrent`, `cooldownMs`, and helper methods (`canRun`, `acquireSlot`, `releaseSlot`, `getCooldownRemaining`, `getRunningCount`, `reset`).
- `Task` generic signature extended to `Task<T, TSteps, TData>` and a new typed data property (`data`) with default `{}`.
- `TaskManager` now supports `addConstraintGroup/removeConstraintGroup`, `triggerTaskConstrained`, queues blocked tasks, drains queue with cooldown timers, and routes `triggerTask/triggerTaskByName` through the constraint system.
- Removed `TaskRunner`, plus `Task` APIs: `blockingTasks`, `execDelay`, `finished` promise and associated behavior have been removed (breaking changes).
- Exports and interfaces updated: `TaskConstraintGroup` and `ITaskConstraintGroupOptions` added; `TaskRunner` removed from public API.
- Updated README and added comprehensive tests for constraint behavior; adjusted other tests to remove `TaskRunner` usage and reflect new APIs.

2026-02-15 - 4.2.1 - fix(deps)

bump @push.rocks/smartlog and @types/node; update dependency list version and license link in docs

- package.json: update @push.rocks/smartlog from ^3.1.10 to ^3.1.11
- package.json: update @types/node from ^25.1.0 to ^25.2.3
- readme.hints.md: update 'Dependencies (as of v4.1.1)' to 'Dependencies (as of v4.2.0)' and reflect bumped dependency versions
- readme.md: change license link text to '[LICENSE](#)'

2026-01-29 - 4.2.0 - feat(ts_web)

support TC39 'accessor' decorators for web components; bump dependencies and devDependencies; rename browser tests to .chromium.ts; move LICENSE to license.md and update readme

- Convert web component class fields to use the TC39 'accessor' keyword in ts_web/taskbuffer-dashboard.ts to be compatible with @design.estate/dees-element v2.1.6
- Bump @design.estate/dees-element to ^2.1.6 and update devDependencies (@git.zone/tsbuild, @git.zone/tsbundle, @git.zone/tsrun, @git.zone/tstest, @types/node) to newer versions
- Replace test/test.10.webcomponent.browser.ts with test/test.10.webcomponent.chromium.ts and update testing guidance in readme.hints.md to prefer .chromium.ts
- Move LICENSE file content to license.md and update readme.md to reference the new license file
- Small test cleanups: remove obsolete tslint:disable comments

2026-01-26 - 4.1.1 - fix(ts_web)

fix web dashboard typings and update generated commit info

- Updated generated commit info file ts_web/00_commitinfo_data.ts to version 4.1.0
- Large changes applied to web/TS build files (net +529 additions, -399 deletions) — likely fixes and typing/refactor improvements in ts_web/dashboard
- package.json remains at 4.1.0; recommend a patch bump to 4.1.1 for these fixes

2026-01-26 - 4.1.0 - feat(task)

add task labels and push-based task events

- Introduce Task labels: Task accepts labels in constructor and exposes setLabel/getLabel/removeLabel/hasLabel; labels are included (shallow copy) in getMetadata().
- Add push-based events: Task.eventSubject (rxjs Subject<ITaskEvent>) emits 'started','step','completed','failed' with timestamp; 'step' includes stepName and 'failed' includes error string.
- Task now emits events during lifecycle: emits 'started' at run start, 'step' on notifyStep, and 'completed' or 'failed' when finished or errored. getMetadata() now includes labels.
- TaskManager aggregates task events into taskSubject, subscribes on addTask and unsubscribes on removeTask/stop; includes helper methods getTasksByLabel and getTasksMetadataByLabel.
- Public API updated: exported ITaskEvent and TTaskEventType in ts/index.ts and interfaces updated (labels in metadata, new event types).
- Tests and docs: added test/test.12.labels-and-events.ts and updated readme.hints.md to document labels and push-based events.

2026-01-25 - 4.0.0 - BREAKING CHANGE(taskbuffer)

Change default Task error handling: trigger() now rejects when taskFunction throws; add catchErrors option (default false) to preserve previous swallow behavior; track errors (lastError, errorCount) and expose them in metadata; improve error propagation and logging across runners, chains, parallels and debounced tasks; add tests and documentation for new behavior.

- Introduce catchErrors option on Task (default: false) — previously errors were swallowed by default
- Tasks now set lastError and increment errorCount when failures occur; clearError() added to reset error state
- getMetadata() now reports status 'failed' and includes lastError and errorCount
- Task.run flow updated to reset error state at start, log errors, and either swallow or rethrow based on catchErrors
- BufferRunner, TaskRunner, Taskchain, Taskparallel, TaskDebounce and TaskManager updated to handle errors, avoid hanging promises, and use logger instead of console
- Added comprehensive tests (test/test.11.errorhandling.ts) and readme hints documenting the new error-handling behavior (v3.6.0+)
- npmextra.json updated for @git.zone/cli and release registries

2025-12-04 - 3.5.0 - feat(core)

Add debounced tasks and step-based progress tracking; upgrade deps and improve dashboard and scheduling

- Add TaskDebounce class to coalesce rapid triggers into a single execution (debounce behavior).
- Introduce step tracking and progress reporting on Task via TaskStep, getProgress(), getStepsMetadata(), getMetadata(), resetSteps(), and completeAllSteps().
- Enhance buffered execution flow: BufferRunner and CycleCounter improvements to better coordinate buffered runs and cycle promises.
- Standardize concurrent runner naming (Taskparallel) and update related exports/usages (ts/index.ts, readme examples).
- Enhance TaskManager scheduling/metadata: getScheduledTasks now returns schedule and nextRun, addExecuteRemoveTask collects execution report metadata and cleans up after execution, distributed coordination hooks retained.
- Add/upgrade web dashboard UI, demos and refresh logic to surface task metadata, scheduled tasks and progress.
- Bump runtime and dev dependencies (multiple @push.rocks packages and @git.zone tooling).
- Update tests: reduce iteration threshold and tighten schedule interval in test/test.4.taskmanager.ts.
- Remove several .serena memory files (project overview, style guides and suggested commands) as cleanup.

2025-09-07 - 3.4.0 -

feat(taskbuffer-dashboard)

Add TaskBuffer dashboard web component, demo and browser tests; add HTML entry and update dependencies

- Introduce a new web component taskbuffer-dashboard for real-time visualization of tasks and schedules (ts_web/taskbuffer-dashboard.ts).
- Add a demo wrapper and interactive UI for the dashboard (ts_web/elements/taskbuffer-dashboard.demo.ts).
- Provide web exports and typings for web usage (ts_web/index.ts) and include an HTML entry (html/index.html).
- Add browser-oriented tests to validate metadata structures for the web component (test/test.10.webcomponent.browser.ts).
- Bump package version to 3.3.0 in package.json as part of this change.
- Update/add dependencies and devDependencies (@design.estate/dees-element added; smartlog, @git.zone/tsbuild and @git.zone/tstest bumped).

2025-09-06 - 3.2.0 - feat(core)

Add step-based progress tracking, task metadata and enhanced TaskManager scheduling/metadata APIs

- Introduce TaskStep class for named, weighted steps with timing and status (pending|active|completed).
- Add step-tracking to Task: notifyStep, getProgress, getStepsMetadata, getMetadata, resetSteps and internal step lifecycle handling.
- Task now records runCount and lastRun; Task.run flow resets/cleans steps and aggregates progress.
- TaskManager enhancements: schedule/deschedule improvements, performDistributedConsultation, and new metadata-focused APIs: getTaskMetadata, getAllTasksMetadata, getScheduledTasks, getNextScheduledRuns, addExecuteRemoveTask (exec + collect report).
- Exports updated: TaskStep and related types exported from index, plus Task metadata interfaces.
- Comprehensive README updates documenting step-based progress tracking, metadata, TaskManager and examples.
- New/updated tests added for step behavior and metadata (test/test.9.steps.ts) and other TS additions.
- Minor build/script change: build script updated to use 'tsbuild tsfolders'.

2025-08-26 - 3.1.10 - fix(task)

Implement core Task execution flow, buffering and lifecycle; update README with generics and buffer docs

- Implement Task.runTask including preTask/afterTask chaining, touched-task cycle prevention and error handling.
- Add Task helpers: extractTask, isTask, isTaskTouched and emptyTaskFunction (resolved promise).
- Introduce task lifecycle coordination: finished promise, resolveFinished, and blockingTasks to await dependent tasks.
- Support taskSetup/setupValue, execDelay handling, and wait-for-blocking-tasks before execution.
- Wire up trigger() to choose buffered vs unbuffered execution (triggerBuffered / triggerUnBuffered) and integrate BufferRunner.
- Improve logging and safer promise handling (caught errors are logged).
- Update README with extended TypeScript generics examples and expanded buffer behavior and strategies documentation.

2025-08-26 - 3.1.9 - fix(tests)

Update CI workflows, fix tests and refresh README/package metadata

- CI: switch Docker image to code.foss.global/host.today/ht-docker-node:npmci and adjust NPMCI_COMPUTED_REPOURL; replace npmci installer package name from @shipzone/npmci to @ship.zone/npmci in Gitea workflows
- Tests: update test imports to use @git.zone/tstest/tapbundle and apply small formatting fixes to test files
- Package metadata: update bugs URL and homepage to code.foss.global, add a pnpm.overrides placeholder in package.json
- .gitignore: add AI/tooling directories (.claude, .serena) and reorganize custom section
- Code style/TS fixes: minor formatting changes across ts sources (trailing commas, line breaks, consistent object/argument commas) and small API surface formatting fixes
- Documentation: whitespace/formatting cleanups in README and add changelog entry for 3.1.8

2025-08-26 - 3.1.8 - fix(tests)

Update test runner and imports, refresh README and package metadata, add project tooling/config files

- Replaced test imports from '@push.rocks/tapbundle' to '@git.zone/tstest/tapbundle' across test files
- Updated test script in package.json to run tstest with --verbose --logfile --timeout 120
- Bumped devDependency @git.zone/tstest to ^2.3.5 and adjusted package.json fields (typings, packageManager)
- Expanded and rewrote README with detailed examples, API reference, and usage guidance
- Refactored TaskManager tests (removed duplicate both-file and added consolidated test/test.4.taskmanager.ts)
- Added development/project tooling and metadata files (.claude settings, .serena project/memories) to aid local development and CI

2024-05-29 - 3.1.7 - maintenance/config

Updated package metadata and build configuration.

- Updated package description.
- Multiple TypeScript configuration updates (tsconfig).
- Updated npmextra.json githost entries (changes across 2024-03-30, 2024-04-01, 2024-04-14).

2023-08-04 - 3.0.15 - feat(Task)

Tasks can now be blocked by other tasks.

- Introduced task blocking support in the Task implementation.
- Release contains related maintenance and patch fixes.

2023-01-07 to 2023-10-20 - 3.0.4..3.1.6 - maintenance

Series of patch releases focused on core fixes and stability.

- Numerous core fixes and small adjustments across many patch versions.
- General maintenance: bug fixes, internal updates and stability improvements.

2022-03-25 - 2.1.17 - BREAKING(core)

Switched module format to ESM (breaking).

- BREAKING CHANGE: project now uses ESM module format.
- Release includes the version bump and migration to ESM.

2019-11-28 - 2.0.16 - feat(taskrunner)

Introduce a working task runner.

- Added/activated a working taskrunner implementation.
- Improvements to task execution and orchestration.

2019-09-05 to 2022-11-14 - 2.0.3..2.1.16 - maintenance

Ongoing maintenance and incremental fixes between 2.0.x and 2.1.x series.

- Multiple fixes labeled as core maintenance updates.
- CI, packaging and small doc/test fixes rolled out across these releases.

2018-08-04 - 2.0.0 - major

Major release and scope change with CI/test updates.

- Released 2.0.0 with updated docs.
- BREAKING CHANGE: package scope switched to @pushrocks (scope migration).
- CI and testing updates (moved to new ttest), package.json adjustments.

2017-07-12 - 1.0.21 - enhancements

Feature additions around task utilities and manager.

- Introduced TaskOnce.
- Implemented TaskManager (added TaskManager class and improvements across 1.0.10-1.0.16).
- Implemented execDelay for tasks.
- Documentation and test improvements.

2016-08-03 - 1.0.6 - types

Type and promise improvements.

- Now returns correct Promise types.
- Dependency and typings updates.

2016-08-01 - 1.0.0 - stable

First stable 1.0.0 release.

- Exported public interfaces.
- Base API stabilized for 1.x line.

2016-05-15 to 2016-05-06 - 0.1.0..0.0.5 - initial features

Initial implementation of core task primitives and utilities.

- Added Taskparallel class to execute tasks in parallel.
- Introduced basic Task class and working taskchain.
- Added logging and initial task buffering behavior.
- Improvements to README, typings and packaging.
- Early CI and build setup (Travis/GitLab CI).