

Archetype ECS lib Documentation

Integrated content management

PirateJL

Copyright © 2026 PirateJL

Table of contents

1. Archetype ECS Lib	4
1.1 Install	4
1.2 Quick start	4
1.3 Notes & limitations	4
1.4 License	5
2. Changelog	6
2.1 v1.0.0	6
2.2 v0.6.1	7
2.3 v0.6.0	8
2.4 v0.5.0	9
2.5 v0.4.2	10
2.6 v0.4.1	10
3. Explanation	11
3.1 Understanding ECS Debugging & Profiling	11
3.2 Architecture overview	13
3.3 ECS and the game loop	14
3.4 Integrating an ECS with Three.js	17
3.5 What people mean by a “full ECS”	20
3.6 Why archetype ECS?	22
3.7 Why deferred commands exist in an archetype ECS	25
3.8 Why use Events in ECS?	28
3.9 World Snapshots (Save and Load)	30
4. How To Guides	32
4.1 How to Debug Your ECS Application	32
4.2 How to add InputState + AssetCache as Resources and use them in systems	35
4.3 How to add/remove components at runtime	39
4.4 How to despawn entities safely	40
4.5 How to filter queries with <code>with</code> and <code>without</code>	41
4.6 How to have multiple Worlds (globe vs ground simulation)	43
4.7 How to integrate ECS into a game loop	44
4.8 How to run logic conditionally	45
4.9 How to save and load a World	46
4.10 How to split logic into multiple system phases	48
4.11 How to use ECS alongside Three.js	49
4.12 How to use Events to decouple systems across phases	50

5. Reference	52
5.1 Archetypes	52
5.2 Commands	54
5.3 Components	57
5.4 Debugging & Profiling Reference	59
5.5 Entity	63
5.6 Reference: Events API	65
5.7 Non goals	67
5.8 Query — Reference	68
5.9 Resources (Singletons / World Globals)	73
5.10 Save and Load (World Snapshot) Reference	77
5.11 Schedule	80
5.12 Systems	84
5.13 World	86
6. Tutorials	93
6.1 Debugging Your First ECS Application	93
6.2 Tutorial 1 - Your first ECS World	97
6.3 Tutorial 2 - Components & archetypes	99
6.4 Tutorial 3 - Deferred structural changes	102
6.5 Tutorial 4 - Writing systems	106
6.6 Tutorial 5 - ECS + Three.js (render-sync + safe spawn/despawn)	109
6.7 Tutorial 6 - Save and load	114

coverage 97.18%

1. Archetype ECS Lib

A tiny **archetype-based ECS** (Entity Component System) for TypeScript.

This documentation is split into 4 parts :

- **Explanation** the general operation of the library
- Find information in the **Reference**
- Target a specific goal using the **How-To Guides**
- Learn through the **Tutorials**: step-by-step guidance
- See what changed in the **Changelog**

1.1 Install

[NPM package available here](#)

```
1 npm i archetype-ecs-lib
```

1.2 Quick start

```

1 import { World, Schedule } from "archetype-ecs-lib";
2
3 class Position { constructor(public x = 0, public y = 0) {} }
4 class Velocity { constructor(public x = 0, public y = 0) {} }
5
6 const world = new World();
7
8 // Spawn immediately
9 const e = world.spawn();
10 world.add(e, Position, new Position(0, 0));
11 world.add(e, Velocity, new Velocity(1, 0));
12
13 // A simple system
14 world.addSystem(w => {
15   for (const { e, c1: pos, c2: vel } of w.query(Position, Velocity)) {
16     pos.x += vel.x * dt;
17     pos.y += vel.y * dt;
18   }
19   // Defer structural changes safely
20   if (pos.x > 10) w.cmd().despawn(e);
21 }
22 });
23
24 world.update(1 / 60);

```

Note: `SystemFn` is typed as `(world: WorldApi, dt) => void`.

Checkout the [tutorials](#) for more!

1.3 Notes & limitations

- This is intentionally minimal: **no parallelism**, no borrow-checking, no automatic conflict detection.
- Query results use `c1/c2/...` fields for stability and speed; you can wrap this in helpers if you prefer tuple returns.
- `TypeId` assignment is process-local and based on constructor identity (`WeakMap`).

1.4 License

This code is distributed under the terms and conditions of the [MIT license](#).

🕒 March 23, 2026

🕒 January 3, 2026

2. Changelog

All notable changes to this project will be documented in this file.

2.1 v1.0.0

2.1.1 API Stabilization & Ergonomics

Goal: harden the public API surface before freezing `WorldApi`.

Breaking Changes

- `spawnMany` **renamed to** `spawnWith` — the old name was misleading (`spawnMany` implied spawning multiple entities, like `despawnMany`). `spawnWith` clearly conveys "spawn one entity with these components". Update all call sites: `world.spawnMany(...)` → `world.spawnWith(...)`, `cmd().spawnMany(...)` → `cmd().spawnWith(...)`.

New Features

- `cmd().addMany()` / `cmd().removeMany()` **as batched commands** — previously these methods decomposed into N individual `add` / `remove` commands, causing N archetype migrations on flush. They now enqueue a single `addMany` / `removeMany` command, matching the single-migration behaviour of the direct `world.addMany()` / `world.removeMany()` calls.
- `world.destroy()` **guard on all public methods** — after `world.destroy()` is called, every public method now throws "Cannot use a destroyed World." immediately. Previously only `spawn()` had this check; other methods produced confusing stale-entity errors or silently no-oped.

Audit Fixes

- **@internal leaks removed** — `_scheduleSystems`, `_getSystemCount`, and `_warnAboutLifecycleConflict` are no longer part of the public `WorldApi` interface.
- **flatted dependency bumped to ^3.4.2** — resolves [GHSA-25h7-pfq9-p65f](#) (unbounded recursion DoS in `parse()` revive phase).
- `npm run lint` passes clean with no warnings.

Performance & Archetype Graph

- **Archetype graph edge cache** — add/remove transitions between archetypes are now cached (`addEdges` / `removeEdges` maps on each `Archetype`). Repeated single-component `add` / `remove` on the same type no longer re-compute `signatureKey` on hot paths.
- **Query result caching** — `_matchingArchetypes()` caches results per query signature. New archetypes are scanned incrementally; unchanged queries pay no rescan cost.
- **Benchmark suite** — baseline numbers for spawn throughput, query iteration, and component add/remove are tracked in `test/bench.bench.ts`.
- Per-entity allocations in `query()` reduced via V8 object-shape stability (per-arity `switch` branches).
- Documentation added for when to prefer `queryTables` vs `query` vs `queryEach`.

Query Filters

- **without exclusion filter** — exclude entities that have a given component: `world.query(Position, { without: [Dead] })`.

- with **presence filter** — require a component without reading its value (marker components): `world.query(Position, { with: [Active] })`.
- Filter support across all three query styles (`query`, `queryTables`, `queryEach`).
- Typed overloads and `QueryRow` / `QueryTable` types updated accordingly.

WorldApi Freeze

No breaking changes to `WorldApi` will be made after this release.

Documentation

- **Versioned docs** — documentation is now versioned using `mike`. Each release tag (`v1.0.0`, `v1.1.0`, ...) publishes a dedicated version at its own URL; the `latest` alias always points to the most recent stable release.
- **Preview channel** — every push to `master` deploys a live `preview` version of the docs, reflecting unreleased changes. The preview build is skipped automatically when `docs/` and `mkdocs.yml` are identical to the latest release tag.
- **PDF per version** — each version (including `preview`) generates its own downloadable PDF.

2.2 v0.6.1

2.2.1 Security

Resolved two vulnerabilities via `npm audit fix` — both are transitive dependencies from the Jest / TypeScript ESLint toolchain.

Package	Severity	Advisory
<code>ajv < 6.14.0</code>	Moderate	GHSA-2g4f-4pwh-qvx6 — ReDoS when using <code>\$data</code>
<code>minimatch ≤ 3.1.3 / 9.0.0-9.0.6</code>	High	GHSA-3ppc-4f35-3m26 , GHSA-7r86-cg39-jmmj , GHSA-23c5-xmqv-rm74 — Multiple ReDoS

All tests and builds pass with no regressions after the update.

2.2.2 Bug Fixes

`World.update()` — finally block robustness (`src/ecs/World.ts`)

If `flush()` threw inside the `finally` block (e.g. from a double-despawn), all subsequent statements were skipped, causing three cascading failures:

- `swapEvents()` was never called — events emitted that frame were delivered one frame late
- `_profAddPhase` and `_profEndFrame` were never called — `_frameCounter` incremented but history arrays were not pushed, breaking `stats()`
- If a system also threw, the system error was masked by the flush error

Fix: Wrapped `flush()` in its own `try/catch` so `swapEvents()`, `_profAddPhase`, and `_profEndFrame` always execute. A new test covers the flush-throws path.

`Schedule.run()` — profiling cleanup on system throw (`src/ecs/Schedule.ts`)

`_profBeginFrame(dt)` was called at the start of each run but there was no `try/finally` around the phase loop. If a system threw, `_profEndFrame()` and `updateOverlay()` were never called, causing `stats().frame` to permanently diverge from `statsHistory().size` and the overlay to stop rendering.

Fix: Added a `try/finally` around the phase loop, mirroring the pattern already used in `World.update()`.

Query generator — `_iterateDepth` leak on abandoned iterator (`src/ecs/world.ts`)

`world.query()` and `world.queryTables()` increment `_iterateDepth` on first `next()` and decrement it in a `finally` block. The `finally` only runs when the generator is exhausted or explicitly closed via `.return()` / `.throw()`. Callers who advanced the generator manually via `.next()` and dropped the reference without cleanup permanently leaked `_iterateDepth` until GC — during that window, every structural change threw `"Cannot do structural change while iterating"`, completely breaking the world.

Note: `for...of` (including with `break`) is safe and remains the recommended pattern — it always calls `.return()` on early exit.

Fix: Added a `FinalizationRegistry` to detect abandoned generators and immediately decrement `_iterateDepth`.

2.2.3 Upgrading

Drop-in patch release with no API changes.

- Run `npm audit fix` (or update to `0.6.1`) to pick up the security patches.
- No migration steps required.
- All existing `for...of` query patterns continue to work without changes.

2.3 v0.6.0

2.3.1 Schedule Upgrades, Debugging Tools, Zero-alloc Queries, World Snapshot/Restore

This release brings four major improvements focused on **performance, determinism, and developer experience**.

Schedule Upgrades

- Keep `Schedule.run(phases)` while adding:
- **Before/after constraints:** `schedule.add(world, "sim", fn).after("input")`
- **Phase order registration:** `schedule.setOrder([...])`
- **Optional flush control:** automatic flush-after-phase points or manual flush mode

Debugging Tools (profiling + overlay)

- **Stats overlay** configurable via `new World({ statsOverlayOptions: ... })`
- Supports position/size, thresholds (`targetFrameMs`, `slowFrameMs`), history size (`maxSamples`), and custom parent element
- **World stats APIs**
- `world.stats()` → live counts + timing (entities, archetypes, systems, resources, events, pending commands; `frameMs`, `phaseMs`, `systemMs`, `dt`, `frame`)
- `world.statsHistory()` → rolling arrays (`dt`, `frameMs`, `phaseMs`, `systemMs`)
- **Profiling controls**
- `world.setProfilingEnabled(enabled)` toggles timing collection
- `world.setProfilingHistorySize(frames)` changes history capacity
- **Overlay controls**
- `world.setDebugging(enabled)` toggles overlay
- `world.destroyOverlay()` removes overlay and listeners
- `world.updateOverlay(stats, history)` manually redraws
- Overlay UI: draggable title bar, console logging toggle, collapse/expand, frame-time history graph and threshold legend

New Query APIs (zero/low alloc, additive)

- `world.queryTables(...ctors)` — yields per-archetype SoA tables, avoids per-entity object allocation
- `world.queryEach(...ctors, fn)` — iterates entities with a callback (no yielded objects)

World Snapshot / Restore (Save & Load)

Deterministic save/load via pure-data snapshots:

- `world.snapshot(): WorldSnapshot`
- `world.restore(snapshot: WorldSnapshot): void`

Persisted (opt-in): entities, components, resources. **Not persisted:** systems, events, schedules/phases, commands/queues (rebuilt on load).

2.4 v0.5.0

2.4.1 Resources, Events, Better Typing, Bundle Helpers, Docs Site

This release focuses on **ergonomics and strong typing** while staying fast, simple, and backwards-compatible.

New: Resources (Singletons) + Events API

- **Resources:** one instance per type (Bevy-style singletons) — `setResource`, `getResource`, `requireResource`, `hasResource`, `removeResource`, `initResource`
- **Events:** emit during systems, drain/consume after a phase — `emit`, `events`, `drainEvents`, `clearEvents`, `swapEvents`

Improved: Strong Typing for `query()` Results

`query()` now infers component types so systems don't need any .

```
1 for (const { e, c1: pos, c2: vel } of world.query(Position, Velocity)) {
2   // pos: Position, vel: Velocity - fully typed
3 }
```

Includes overloads up to 6 components.

Improved: SystemFn Typing

- Introduces the exported `WorldApi` interface.
- `SystemFn` now targets `WorldApi` so systems type-check cleanly without casts.

New: Bundle Helpers

- `world.spawnWith([Position, new Position()], [Velocity, new Velocity()])` — spawn one entity with multiple components
- `world.addMany(e, [Ctor, value], ...)` — add multiple components in one call
- `world.removeMany(e, CtorA, CtorB, ...)` — remove multiple components in one call

New: Documentation Site

Published via MkDocs and deployed to GitHub Pages: <https://piratejl.github.io/archetype-ecs-lib/>

2.5 v0.4.2

2.5.1 Iteration Safety, Command Flushing, Diagnostics

Patch focused on **iteration safety, more predictable command flushing, better diagnostics**, and small internal performance wins. No API changes.

Fixed

- **Robust iteration lock** — iteration state is now tracked with a depth counter instead of a boolean, preventing a nested query from accidentally unlocking structural changes while an outer iteration is still running.
- `flush()` **forbidden during iteration** — calling `flush()` during iteration now throws the same way as `add / remove / despawn`.

Changed

- `flush()` **drains until stable** — `flush()` now repeatedly drains and applies commands until the queue is empty, supporting the `cmd().spawn(init)` pattern where `init` may enqueue further commands.

Improved

- **Better error context** — errors now include schedule phase name, entity id/gen, and component constructor name where applicable.

Performance

- `query()` reduces repeated allocations per call.
- `Commands.drain()` swaps buffers instead of copying.

Migration

No API changes. If your code called `flush()` inside a system while iterating it will now throw — move those structural updates to phase boundaries.

2.6 v0.4.1

Initial public release on [npm](#).

- Create and despawn entities
- Register and query components
- Basic system loop via `world.addSystem()` / `world.update(dt)`
- Archetypes (tables) store entities in SoA layout (one column per component type)
- Queries iterate matching archetypes efficiently
- Commands defer structural changes (`spawn / despawn / add / remove`) safely
- Minimal `schedule` runs systems by phase and flushes commands between phases

🕒 April 1, 2026

🕒 March 23, 2026

3. Explanation

3.1 Understanding ECS Debugging & Profiling

3.1.1 Why Debugging ECS is Different

Entity Component System architectures present unique debugging challenges compared to traditional object-oriented code. Instead of inspecting individual objects with encapsulated state, you're dealing with:

- **Entities:** Just numeric IDs with no inherent meaning
- **Components:** Data scattered across archetype tables
- **Systems:** Functions that operate on query results, not objects
- **Deferred operations:** Commands that execute later, not when called

This separation of data and logic makes traditional breakpoint debugging less intuitive. You can't simply "step into" an entity to see what it's doing.

3.1.2 The Stats Overlay Philosophy

The built-in stats overlay provides **runtime observability** rather than step-through debugging. It answers questions like:

- "How many entities exist right now?"
- "Which systems are slow?"
- "Are my archetypes fragmenting?"
- "Is my frame budget being exceeded?"

This approach aligns with how games and simulations are typically debugged—by observing the running system rather than pausing it.

3.1.3 What the Metrics Mean

Entity & Archetype Counts

The relationship between **alive entities**, **archetypes**, and **rows** tells you about your data layout:

Metric	Healthy Sign	Warning Sign
Archetypes	Stable, low count	Growing unboundedly
Rows \approx Entities	Close match	Large gap indicates fragmentation
Entities	Predictable growth	Unexpected spikes

Archetype explosion occurs when entities frequently add/remove components, creating many unique component combinations. Each unique combination requires its own archetype table.

Frame Timing

The overlay tracks two time measurements:

- **dt:** Delta time passed to `update()` (what your game logic sees)
- **frame:** Actual wall-clock time spent in the update (what the CPU experiences)

When `frame` consistently exceeds `dt`, you're falling behind real-time.

Phase & System Timing

With the `Schedule` API, systems are grouped into phases. The overlay shows:

- Per-phase total time
- Per-system time (when profiling is enabled)

This helps identify which phase or system is the bottleneck.

3.1.4 Profiling History

The overlay maintains a rolling history (default: 120 frames) displayed as a bar graph:

- **Blue bars:** Frames within budget (\leq slow threshold)
- **Red bars:** Slow frames ($>$ slow threshold)
- **White line:** Target frame time

Patterns in this graph reveal:


- **Periodic spikes:** Often garbage collection or physics sync
- **Gradual increase:** Memory leak or unbounded growth
- **Consistent red:** Fundamental performance problem

3.1.5 Debug Logging

The `/` toggle enables console output of phase timings each frame. This is useful for:

- Capturing timing data for analysis
- Correlating visual stutters with logged spikes
- Automated performance regression testing

 February 3, 2026

 February 3, 2026

3.2 Architecture overview

3.2.1 Architecture Overview

```

1 World – Central owner of all ECS state
2 |
3 |— EntityManager       ID + generation counter, free-list allocator
4 |
5 |— Archetypes          One table per unique component set, SoA columns, add/remove edge cache
6 |   |— Query Cache     Incremental archetype matching, invalidated on new archetype
7 |
8 |— TypeRegistry        Maps component constructor – TypeId, used to build signatures
9 |
10 |— Resources           Global singletons, keyed by constructor
11 |
12 |— Event Channels     Double-buffered per type, write this phase – read next phase
13 |
14 |— Commands           Deferred queue (spawn · despawn · add · remove), flushed after each phase
15 |
16 |— Snapshot Store     Opt-in serialize / restore, codec registered per type
17 |
18 |— StatsOverlay       Per-frame profiling, optional DOM HUD
19 |
20 |— world.update(dt)   Single-phase loop, auto flush + swap events
21 |   |— SystemFn[]
22 |
23 |— Schedule            Multi-phase, topological ordering, .after() / .before() constraints,
24 |   |— Phase – SystemFn[]
25 |

```

Data flow per frame



Two update modes (mutually exclusive)

	world.update(dt)	schedule.run(world, dt)
Phases	Single	Multiple, ordered
Flush	Once at end of frame	After each phase
Event swap	Once at end of frame	After each phase
Use case	Simple loop / prototyping	Production, multi-system games

🕒 March 23, 2026

🕒 March 16, 2026

3.3 ECS and the game loop

ECS is best understood as the **way you organize game state and game logic**, not as the thing that *does everything*. In a typical game, the loop still has input, rendering, audio, physics, networking, etc. ECS provides a **consistent place** for runtime data (components) and behavior (systems), plus a **schedule** that defines when that behavior runs. This library already models this explicitly with `World.update(dt)` and with a phase-based `Schedule` that flushes between phases.

3.3.1 Frame phases

A “frame” is rarely just “update then draw”. Most games are structured in phases, even if informally. A common conceptual breakdown:

1. **Input**: read devices/events, translate into game intent
2. **Simulation**: movement, AI, gameplay rules, timers
3. **Physics** (optional separate step): integrate, solve collisions, constraints
4. **Post-sim**: resolve gameplay outcomes, spawn/despawn, apply state transitions
5. **Render prep**: build renderable data, sort, cull
6. **Render**: submit to GPU / engine renderer
7. **End-of-frame**: cleanup, present frame, etc.

The `Schedule` is designed exactly for this idea: you define phases (strings) and run them in order, with `flush()` after each phase.

3.3.2 Where ECS fits

ECS typically fits in the **simulation and render-prep** parts of the loop:

- **World** holds the mutable runtime state (entities + components)
- **Systems** implement the game logic by querying components and mutating them
- **Commands** allow safe structural changes during those systems (`cmd()` → `flush()`)
- **Schedule** provides deterministic ordering and safe mutation boundaries between phases

A useful mental model:

- Rendering engines want a *renderable snapshot* (meshes, transforms, materials, draw lists).
- Input systems produce *intent/state* (move left, fire, target position).
- Physics engines operate on *physical representations* (bodies, colliders).

ECS sits in the middle coordinating these, not replacing them.

A concrete mapping using this primitives

- **Input phase**: read input → write `InputState` component / resource → enqueue spawns/despawns if needed
- `flush()`
- **Sim phase**: run movement/AI/gameplay using queries → update `Position`, `Velocity`, etc.
- `flush()`
- **Render phase**: build lightweight render data (`RenderTransform`, `Visible`, etc.) → hand off to renderer

This is why “flush points” exist in an ECS schedule: they define when the world structure is allowed to change and when the next phase sees those changes.

3.3.3 Why ECS does not replace rendering, input, or physics engines

Rendering

A renderer is a specialized pipeline:

- GPU resources, shaders, batching, sorting, culling
- frame graph / render passes
- platform-specific backends

ECS is not a GPU pipeline. What ECS does well is:

- storing render-related data as components (`Transform`, `Renderable`, `MaterialRef`, etc.)
- running systems that prepare and synchronize data for the renderer

So ECS often produces a **render list** or updates engine scene objects, but the renderer still does the rendering.

Input

Input is inherently eventful and platform-driven:

- OS/window events
- device state polling
- mapping raw events to game actions

ECS can *store* input state (`InputAxis`, `ActionPressed`, etc.) and *process* it in systems, but it doesn't replace the platform input layer. In practice:

- platform collects input
- ECS system transforms it into gameplay-friendly state

Physics

Physics engines are optimized solvers:

- broadphase / narrowphase collision detection
- integrators and constraint solvers
- continuous collision, joints, sleeping, etc.

ECS can represent physics **data** (mass, collider type, desired forces) and drive the physics engine, but the solver itself is a dedicated subsystem.

A common integration pattern:

- ECS → write forces/desired velocity into physics engine
- Physics step happens
- Physics results → write back transforms/velocities into ECS

3.3.4 The key idea: ECS is the *coordination model*

ECS shines when you treat it as:

- **a data model** for game state (components)
- **a behavior model** for game logic (systems)
- **an execution model** for ordering (schedule + phases + flush points)

But rendering/input/physics are specialized domains with their own constraints and pipelines. ECS coordinates them by being the “truth” for game state and by running the logic that translates between subsystems.

🕒 January 4, 2026

🕒 January 4, 2026

3.4 Integrating an ECS with Three.js

Three.js is a **rendering engine** (scene graph + GPU submission). This ECS is a **simulation architecture** (data in components, behavior in systems, ordered by a schedule, with safe structural changes via deferred commands + flush points). Integrating them well means **letting each do what it's good at**, and defining clean “hand-off” boundaries.

3.4.1 The mental model: ECS drives state, Three.js draws it

A practical split that scales:

- **ECS World** = authoritative game/sim state (position, velocity, health, selection, etc.)
- **Three.js Scene** = visual representation (Object3D transforms, meshes, materials, lights)

So the goal is not “put Three.js inside ECS”, but:

Systems write simulation state → a render-sync step pushes that state into Three.js objects.

3.4.2 Where ECS fits in the Three.js render loop

Three.js typically runs:

1. update (your code)
2. `renderer.render(scene, camera)`

With ECS, your “update” becomes scheduled phases, e.g.:

- `input` (read DOM/input, write components/resources)
- `sim` (gameplay, movement, AI)
- `render` (sync ECS → Three.js, then render)

The `Schedule` already supports this exact idea and flushes commands between phases to make entity/component creation/removal deterministic.

3.4.3 Why flush points matter for Three.js integration

Spawning/despawning and add/remove are **structural changes** in this ECS and are expected to be deferred while iterating queries/systems.

That maps perfectly to Three.js object lifecycle:

- **During sim:** decide “this entity should appear/disappear” → enqueue ECS commands
- **At flush boundary:** ECS structure becomes stable
- **Render-sync phase:** create/remove corresponding `Object3D` safely, because you're no longer mid-iteration on archetype tables

This is the same reason this ECS has `cmd()` / `flush()` and why `Schedule` flushes between phases.

3.4.4 A clean integration pattern: “Renderable bridge” components

Common approach:

- A `Transform` component (position/rotation/scale) is owned by ECS.
- A `Renderable` component carries a reference/handle to what Three.js should draw (mesh id, model key, material key...).
- A render-sync system queries (`Transform`, `Renderable`) and applies changes to the corresponding `Object3D`.

Key idea: **ECS components store “what it is” and “where it is”**, while the actual `Mesh/Object3D` lives in Three.js.

This keeps:

- ECS portable (not tied to Three.js types everywhere)
 - Three.js free to manage GPU resources
-

3.4.5 One-way vs two-way sync (pick a source of truth)

Integration gets messy when both ECS and Three.js “own” transforms.

A scalable default:

- **ECS is the source of truth** for gameplay transforms.
- Three.js `Object3D` is just the projection of that state.

Only do **two-way sync** when you truly need it (editor gizmos, drag interactions). Even then, treat it as a controlled input step:

- read `Object3D` change in `input` or `tools` phase
 - write back to ECS components
 - let sim proceed from ECS again
-

3.4.6 Why ECS does not replace Three.js (and shouldn't try)

Even with a “full ECS” architecture, Three.js still owns:

- scene graph concerns (parenting, cameras, lights)
- GPU resource lifetimes (buffers, textures, materials)
- draw submission, sorting, batching, culling strategies

ECS complements that by making **simulation state and logic** scalable: archetype tables + queries + systems + scheduling.

3.4.7 Scaling tips (when entity counts grow)

When you have many similar visuals:

- prefer **InstancedMesh** in Three.js
- let ECS systems produce instance transforms (dense arrays) from queries
- upload those transforms once per frame

This aligns with why archetype ECS exists: tight iteration over dense component columns.

🕒 January 4, 2026

🕒 January 4, 2026

3.5 What people mean by a “full ECS”

“ECS” can mean **just a storage model** (entities + components in some container), or it can mean an **entire game/app architecture** where *most runtime state and behavior* flows through an ECS **world + schedule + systems**.

A “full ECS” is typically an architecture where:

- **Entities** are only IDs/handles (no behavior).
- **Components** are only data.
- **Systems** are where behavior lives (pure-ish functions operating on data).
- A **World** is the single source of truth for runtime state.
- A **Scheduler** (or “app loop”) defines *when* systems run and in what order.
- Structural changes are controlled (often via a **command buffer**) so iteration stays safe and fast.

This library already contains several “full ECS” building blocks: **archetype tables (SoA)**, **queries**, **deferred commands**, and a **phase-based schedule**.

What makes it “full” is less “do you use archetypes?” and more “does the ECS define the whole program’s execution model?”

3.5.1 ECS as architecture, not just storage

Storage-only ECS (not “full”)

This is common in small libs or quick implementations:

- Entities: IDs
- Components: data bags
- “Systems”: often just loops in user code
- Little/no scheduling model
- No consistent lifecycle for input → simulation → rendering
- Structural changes are ad-hoc

You *can* build a game with this, but the ECS isn’t the **organizing principle**—it’s a container.

Architecture ECS (“full ECS”)

Here, ECS is the **spine of the app**:

- There’s a **main schedule** (often phases like `input → sim → render`).
- Systems are registered, ordered, and executed consistently each tick.
- Cross-cutting state is handled intentionally (resources/singletons, events, time, config).
- Structural changes are made safe/deterministic (command buffers, flush points).
- You get a uniform pattern for new features: “add data + add system”.

The `Schedule` explicitly models *phase ordering + flush barriers*, which is a key “architecture ECS” ingredient.

3.5.2 Difference between a library ECS and an engine ECS

Library ECS

Goal: provide **core ECS mechanics**.

Typical traits:

- Focus on **storage + query performance** (archetypes/SoA)
- Minimal assumptions about the rest of the program
- Simple scheduling (or none), often single-threaded
- You (the user) integrate input, rendering, physics, assets, scenes, etc.

Engine ECS (Bevy / Unity DOTS / etc.)

Goal: ECS is the **entire runtime framework**.

Engine ECS usually includes (beyond a library):

- A full **app lifecycle** (startup, update, fixed update, shutdown)
- Integrated **input, rendering, audio, physics, animation, UI**
- Asset pipeline + hot reload + serialization
- Advanced scheduling: dependency graphs, system sets, run criteria, fixed timesteps
- Often **parallel execution** + conflict detection
- Tooling/editor integration

So: **library ECS = the “ECS core”**. **engine ECS = ECS core + everything around it**, with ECS as the central organizing model.

🕒 January 4, 2026

🕒 January 4, 2026

3.6 Why archetype ECS?

An **archetype ECS** organizes entities into **tables** where every entity in a table shares the same component set, stored in **SoA** form (one column per component). This library explicitly follows this model: “Archetypes (tables) store entities in a SoA layout... Queries iterate matching archetypes efficiently... Commands defer structural changes...”

The “why” is mostly about making the *common case* (systems that iterate lots of entities with the same components) extremely fast and predictable.

3.6.1 Cache locality

Most game/sim systems look like:

- “for all entities with `Position` and `Velocity`, update position”
- “for all entities with `Transform` and `Renderable`, build render data”

With archetypes, those entities live together in a table, and each component is a dense column:

- `Position[]` contiguous
- `Velocity[]` contiguous

So the CPU reads memory sequentially, which is what caches and prefetchers love. That’s the practical meaning of **cache locality**: fewer cache misses, more work per nanosecond.

In the library, this is literally the storage promise: SoA archetype tables + queries over matching archetypes.

3.6.2 Branch elimination (and “no-join” iteration)

In many ECS designs, the core loop must constantly ask:

- “does this entity have `Velocity`?”
- “if yes, fetch it; if not, skip”

That creates branches and scattered memory access.

With archetypes, the *membership check is moved up*:

1. pick archetypes that already contain all required components
2. iterate their rows

Inside the inner loop, there’s no per-entity “has component?” branching—every row is guaranteed to match. The API reflects that by querying required component types and yielding direct component references (`c1`, `c2`, ...).

This is what people mean by **branch elimination** in archetype ECS: fewer conditional checks in the hot loop, more straight-line code.

3.6.3 Predictable iteration

Archetype iteration tends to be predictable because:

- You iterate dense arrays (rows/columns), not sparse IDs.
- Results are shaped consistently (`e`, `c1`, `c2`, ... in argument order).
- Structural changes are controlled: this library emphasizes deferring structural changes via `cmd()` and applying them at `flush()` points.
- `schedule` adds explicit “phase barriers” by flushing between phases, making the world structure stable during each phase’s iteration.

That predictability is less about “deterministic order of entities” and more about **deterministic rules for when the world can change shape**.

3.6.4 Comparison with sparse-set ECS

A **sparse-set ECS** typically stores each component type separately (often as a dense array + sparse index by entity id). It’s excellent for:

- fast lookup for a single component type (`Position` alone)
- cheap per-component iteration
- simple storage and often cheaper structural changes for *single* components

But when a system needs **multiple components** (`Position + Velocity + Mass + Forces`), sparse-set often needs some form of **join**:

- iterate one component pool, check membership in the others
- or intersect sets / hop through indirections

That can introduce:

- more branching (`if has(...)`)
- more random memory access (chasing indices across pools)

Archetypes flip that trade-off:

- multi-component iteration is the “happy path” (no join inside the hot loop)
- but structural changes can be more expensive because adding/removing a component may move an entity between tables.

Rule of thumb

- If your game spends most time in **systems that read/write several components per entity**, archetypes tend to shine.
- If your workload is lots of **single-component iteration** and **high churn** (constant add/remove), sparse-set can be simpler and sometimes cheaper.

3.6.5 The real trade-off (why it’s not “always archetypes”)

Archetype ECS wins by making the hot loops fast, but it pays for it with:

- **structural churn cost** (moving entities between tables on add/remove)
- **many archetypes** if you have lots of component combinations
- a stronger need for **command buffering + flush boundaries** to keep iteration safe.

That’s why a “full ECS” architecture often includes commands + scheduling: it’s the natural partner to archetype storage.

🕒 January 14, 2026

🕒 January 4, 2026

3.7 Why deferred commands exist in an archetype ECS

In an archetype ECS, **deferred commands** (a command buffer) are not a “nice-to-have”. They exist because **the fastest storage model makes certain mutations unsafe during iteration**. The library API expresses this directly with `world.cmd()`, `world.flush()`, and `Schedule.run(...)/flush barriers`.

3.7.1 Archetypes are tables, and queries walk those tables

An archetype ECS stores entities in **tables**:

- one archetype = one *component set*
- one row = one entity
- one column per component type (SoA)

A query like `world.query(Position, Velocity)` does not “scan entities”. It first selects archetypes that contain the required component columns, then iterates **dense rows** in those tables.

This density is where the performance comes from.

3.7.2 The core problem: structural changes move entities between tables

A **structural change** is anything that changes the component *set* of an entity:

- `spawn()`
- `despawn(e)`
- `add(e, Ctor, value)`
- `remove(e, Ctor)`

In an archetype ECS, `add/remove` usually means:

1. remove the entity’s row from its current archetype table
2. insert a row into another archetype table
3. update internal bookkeeping (where the entity lives now)

That is fundamentally different from `set(e, Ctor, value)`, which just updates a value *inside the same row/column*.

So: **structural change = table move**.

3.7.3 Why it’s unsafe to do structural changes during a query

When you iterate a query, you are conceptually doing:

- “for each matching archetype table”
- “for each row index in that table”
- “read columns at that row”

If you structurally change any entity during this loop, you can break the iteration invariants:

1) Swap-remove can invalidate the current row

Many archetype implementations remove rows with **swap-remove** ($O(1)$): the last row is swapped into the removed row index.

If you remove entity A at row `i`, entity B may be swapped into row `i`.

- If your loop then increments `i`, entity B might be **skipped**.
- Or processed twice depending on iteration strategy.

2) Moving entities changes which archetypes match

Adding/removing a component can move an entity into or out of the set of archetypes that the query is iterating.

If you mutate membership while iterating:

- you can end up iterating an archetype that didn’t exist in the matching set at the start
- or miss entities that moved into a matching archetype

3) Internal indices can become stale mid-loop

The library `World` tracks where an entity lives (which archetype + row). A structural change updates those indices. If you mutate while holding references from the iteration, you can end up with:

- stale row pointers
- stale bookkeeping
- inconsistent state if multiple mutations occur

Even if you “think it works”, it’s fragile and will eventually bite.

3.7.4 Deferred commands are the solution: separate “read/iterate” from “mutate structure”

A command buffer enforces a clean two-step model:

1. **During iteration:** read data, compute decisions, mutate *component values* (safe)
2. **At a safe boundary:** apply structural changes in a batch (safe)

That’s exactly what the library documents:

- `world.cmd()` enqueues structural operations
- `world.flush()` applies queued commands
- `world.update(dt)` runs systems, then flushes at frame end
- `Schedule.run(...)` flushes **between phases**, providing deterministic barriers

This is why deferred commands exist: they preserve **iteration correctness** without giving up **table-based performance**.

3.7.5 Why flushing in phases is architecturally important

The library `Schedule` explicitly flushes after each phase.

This is not just “nice ordering”. It creates **deterministic points** where the world’s structure is allowed to change.

Example mental model:

- **Input phase:** decide spawns/despawns based on input → enqueue commands
- **Flush:** apply those spawns so they exist for simulation
- **Simulation phase:** move things, detect collisions → enqueue structural changes
- **Flush:** apply spawns/despawns/removals before render
- **Render phase:** build render data from a stable world snapshot

That separation reduces “action at a distance” bugs and makes debugging easier:

- “why does entity exist in sim but not render?” → check which phase flushed it.

3.7.6 What you gain by deferring

Correctness

- No skipped entities
- No double-processing due to swap-remove effects
- Stable iteration semantics

Determinism

- Structural changes occur at explicit boundaries
- Easier to reason about ordering

Performance

- Keeps archetype iteration tight and cache-friendly
- Batching structural operations reduces churn

3.7.7 What to do inside a system

Inside a system (or a query loop), follow this rule:

- mutate component values directly (e.g. `pos.x += ...`)
- enqueue structural changes via `cmd()`
- don't call structural `World` ops directly mid-iteration

3.7.8 Summary: the “why” in one sentence

Deferred commands exist because **archetype queries iterate dense tables**, and **structural changes move rows between tables**, which can invalidate iteration—so we **queue structural changes** and apply them at **safe flush boundaries** (`flush()` / schedule phases).

🕒 January 4, 2026

🕒 January 4, 2026

3.8 Why use Events in ECS?

3.8.1 Events solve a different problem than Components and Resources

ECS has three kinds of data:

- **Components:** persistent, per-entity state (Position, Velocity, Health)
- **Resources:** persistent, global state (Input, Time, Config, caches)
- **Events:** transient messages (Hit happened, Click happened, Play sound)

Trying to represent “something happened” as a component usually causes awkward designs:

- adding/removing “Event components” becomes structural churn
- you need cleanup systems to remove them
- multiple systems race to observe/remove them

Events avoid that by being explicitly transient.

3.8.2 Events reduce coupling between systems

Without events:

- `combatSystem` might call `audioSystem` directly
- or it might mutate a shared global array

With events:

- producers don’t know consumers exist
- consumers don’t know who produced the messages

This keeps systems reusable and easy to rearrange in `Schedule`.

3.8.3 Why double-buffering?

A common bug in event systems is “events appear while I’m iterating”.

Double-buffering prevents that:

- consumers read a stable snapshot (`read buffer`)
- producers write to a different buffer (`write buffer`)
- swap happens at deterministic boundaries

No surprises. No iterator invalidation. No mid-phase visibility.

3.8.4 Why phase-scoped delivery?

This ECS already has a concept of **phase boundaries**:

- structural changes are deferred via Commands
- `flush()` applies them between phases

Events align with the same boundary:

- `swapEvents()` delivers events between phases

This makes it easy to design pipelines:


- `input` produces actions → `beforeUpdate` consumes
 - `update` produces gameplay events → `afterUpdate` consumes
 - `render` produces UI/VFX events → `afterRender` consumes
 - `audio` consumes sound events
-

3.8.5 Trade-offs (and the forwarding pattern)

With phase-scoped delivery, an event is visible in the **next phase only**. To deliver an event across multiple phases (e.g., from `update` to `audio`), you forward it by draining and re-emitting.

This is deliberate:

- it keeps pipelines explicit
- prevents “stale” events lingering through unrelated phases
- makes delivery deterministic and easy to debug

 January 14, 2026

 January 13, 2026

3.9 World Snapshots (Save and Load)

`World.snapshot()` and `World.restore()` let you persist ECS **data state** and rebuild the world later.

The key idea is: save only data that belongs to simulation state, and reconstruct runtime behavior from code.

3.9.1 What is persisted

Snapshot payloads include:

- alive entities
 - registered components (serialized through codecs)
 - registered resources (serialized through codecs)
 - allocator state (`nextId`, free-list, generations) so entity id reuse stays deterministic after load
-

3.9.2 What is not persisted

Snapshot payloads intentionally exclude runtime-only behavior:

- systems
- schedules and phase graph
- pending commands
- event buffers/channels

After `restore(...)`, these are clean runtime state and should be set up by normal app boot code.

3.9.3 Why codecs are explicit and opt-in

Snapshot registration is explicit (`registerComponentSnapshot`, `registerResourceSnapshot`) for three reasons:

1. **Stable identity** Runtime `TypeId` values are process-local. Snapshots need stable string keys like `"comp.position"`.
 2. **Schema control** Not every field should be saved. A codec chooses exactly which data is persisted.
 3. **Long-term compatibility** Save files can outlive one run. Explicit keys and payload shapes make migration/versioning manageable.
-

3.9.4 Determinism goals

Snapshot export is deterministic for the same world state:

- entities are emitted in entity-id order
- component entries are emitted in snapshot-key order
- resource entries are emitted in snapshot-key order

This makes snapshots useful for debugging, test fixtures, and replay checkpoints.

3.9.5 Long-term saves vs quick saves

The same snapshot format supports both:

- quick save/load in memory
- persisted saves (`JSON.stringify(snapshot)` to disk/DB/localStorage)

For long-term saves:

- keep codec keys stable
 - if schema changes, migrate old payloads before `restore()`
-

3.9.6 Mental model

Think of world snapshots as **state serialization**, not **engine serialization**:

- engine runtime is rebuilt
- simulation data is restored

That separation keeps the ECS data-oriented and engine-agnostic.

🕒 February 8, 2026

🕒 February 8, 2026

4. How To Guides

4.1 How to Debug Your ECS Application

4.1.1 How to Enable the Stats Overlay

The stats overlay is automatically created when you instantiate a `World` in a browser environment. However, to enable and show it, use the function `setDebugging` to true.

```
1 import { World } from "archetype-ecs-lib";
2
3 const world = new World();
4 world.setDebugging(true);
5 // Overlay appears automatically in top-left corner
```

Customizing Overlay Position and Appearance

```
1 const world = new World({
2   statsOverlayOptions: {
3     left: 10,           // X position in pixels
4     top: 10,           // Y position in pixels
5     width: 400,        // Canvas width
6     height: 100,       // Canvas height
7     targetFrameMs: 16.67, // Target frame time (60 FPS)
8     slowFrameMs: 20,   // Threshold for "slow" frames
9     maxSamples: 240    // History length in frames
10  }
11 });
```

Attaching to a Custom Container

```
1 const debugPanel = document.getElementById("debug-panel");
2
3 const world = new World({
4   statsOverlayOptions: {
5     parent: debugPanel
6   }
7 });
```

4.1.2 How to Identify Slow Systems

Step 1: Enable Profiling (On by Default)

```
1 world.setProfilingEnabled(true);
```

Step 2: Check the Overlay

Look at the "Phases" line in the stats overlay. Systems are shown with their execution time:

```
1 Phases: update=5.23ms render=2.10ms physics=1.50ms
```

Step 3: Enable Console Logging for Details

Click the button to toggle console debug logging. Each frame will log phase timings:

```
1 Phases: update=5.23ms render=2.10ms physics=1.50ms
```

Step 4: Analyze the History Graph

Red bars indicate frames exceeding the slow threshold. Look for patterns:

- Isolated spikes → Likely GC or async operations
- Periodic spikes → Check fixed-interval systems
- Sustained red → System needs optimization

4.1.3 How to Debug Entity Lifecycle Issues

Finding "Lost" Entities

If entity counts don't match expectations:

```

1 // Check stats for current counts
2 const stats = world.stats();
3 console.log(`Alive: ${stats.aliveEntities}, Rows: ${stats.rows}`);
4
5 // Iterate to find entities with specific components
6 for (const { e, c1 } of world.query(Position)) {
7   console.log(`Entity ${e.id} has Position:`, c1);
8 }

```

Detecting Stale Entity References

```

1 const entity = world.spawn();
2 world.despawn(entity);
3
4 // Later...
5 if (!world.isAlive(entity)) {
6   console.warn("Entity was despawned:", entity);
7 }

```

Tracking Pending Commands

The overlay shows "Pending commands: true/false". If commands aren't being applied:

```

1 // Commands are deferred
2 world.cmd().spawn();
3 console.log(world.stats().pendingCommands); // true
4
5 // After flush
6 world.flush();
7 console.log(world.stats().pendingCommands); // false

```

4.1.4 How to Debug Archetype Fragmentation

Symptoms

- High archetype count relative to entity count
- Performance degradation over time
- Memory growth

Diagnosis

```

1 const stats = world.stats();
2 const ratio = stats.archetypes / stats.aliveEntities;
3
4 if (ratio > 0.1) { // More than 1 archetype per 10 entities
5   console.warn("Possible archetype fragmentation");
6 }

```

Common Causes

1. Frequent component add/remove cycles

```

1 // Bad: Creates new archetypes constantly
2 world.add(e, TempMarker, {});
3 world.remove(e, TempMarker);

```

2. Unique component combinations per entity

```

1 // Bad: Each entity gets unique archetype
2 world.add(e, UniqueId, { id: generateUuid() });
3 world.add(e, CreatedAt, { time: Date.now() });

```

Solutions


- Use resources for singleton data
- Use component values instead of presence/absence for flags
- Batch component changes with `addMany()` / `removeMany()`


4.1.5 How to Disable the Overlay in Production

```

1 // Option 1: Conditional creation
2 const world = new World();
3 world.setDebugging(process.env.NODE_ENV === "development");
4
5 // Option 2: Destroy after creation
6 const world = new World();
7 world.destroyOverlay();

```

 February 3, 2026

 February 3, 2026

4.2 How to add InputState + AssetCache as Resources and use them in systems

4.2.1 Goal

Store **Input state** and an **Asset cache** as world **Resources**, then access them inside systems using `requireResource()`.

Example InputStateRes

```

1  export class InputStateRes
2  {
3      public keysDown = new Set<string>();
4      public keysPressed = new Set<string>(); // pressed this frame
5      public keysReleased = new Set<string>(); // released this frame
6
7      public mouseX = 0;
8      public mouseY = 0;
9      public mouseButtonsDown = new Set<number>();
10     public mousePressed = new Set<number>(); // pressed this frame
11     public mouseReleased = new Set<number>(); // released this frame
12     public wheelDeltaY = 0;
13
14     beginFrame(): void
15     {
16         this.keysPressed.clear();
17         this.keysReleased.clear();
18         this.mousePressed.clear();
19         this.mouseReleased.clear();
20         this.wheelDeltaY = 0;
21     }
22
23     keyDown(code: string): void
24     {
25         if (!this.keysDown.has(code)) this.keysPressed.add(code);
26         this.keysDown.add(code);
27     }
28
29     keyUp(code: string): void
30     {
31         if (this.keysDown.has(code)) this.keysReleased.add(code);
32         this.keysDown.delete(code);
33     }
34
35     mouseMove(x: number, y: number): void {
36         this.mouseX = x;
37         this.mouseY = y;
38     }
39
40     mouseDown(btn: number): void
41     {
42         if (!this.mouseButtonsDown.has(btn)) this.mousePressed.add(btn);
43         this.mouseButtonsDown.add(btn);
44     }
45
46     mouseUp(btn: number): void
47     {
48         if (this.mouseButtonsDown.has(btn)) this.mouseReleased.add(btn);
49         this.mouseButtonsDown.delete(btn);
50     }
51
52     wheel(deltaY: number): void
53     {
54         this.wheelDeltaY += deltaY;
55     }
56 }

```

Example AssetCacheRes

```

1  export class AssetCacheRes
2  {
3    private images = new Map<string, HTMLImageElement>();
4    private pending = new Map<string, Promise<HTMLImageElement>>();
5
6    /** Loads once, dedupes concurrent calls, returns the same instance thereafter. */
7    public getImage(url: string): Promise<HTMLImageElement>
8    {
9      const ready = this.images.get(url);
10     if (ready) return Promise.resolve(ready);
11
12     const p = this.pending.get(url);
13     if (p) return p;
14
15     const promise = new Promise<HTMLImageElement>((resolve, reject) => {
16       const img = new Image();
17       img.onload = () => {
18         this.images.set(url, img);
19         this.pending.delete(url);
20         resolve(img);
21       };
22       img.onerror = (e) => {
23         this.pending.delete(url);
24         reject(e);
25       };
26       img.src = url;
27     });
28
29     this.pending.set(url, promise);
30     return promise;
31   }
32
33   /** Returns the image if already loaded; otherwise undefined. */
34   public peekImage(url: string): HTMLImageElement | undefined {
35     return this.images.get(url);
36   }
37 }

```

4.2.2 1) Register the resources at startup

```

1  world.initResource(InputStateRes, () => new InputStateRes());
2  world.initResource(AssetCacheRes, () => new AssetCacheRes());

```

That's the only “required” setup. Everything else assumes these exist.

4.2.3 2) Wire DOM events into InputStateRes

Attach listeners at once:

```

1  export function attachInput(world: WorldApi): void
2  {
3    const input = world.requireResource(InputStateRes);
4
5    window.addEventListener("keydown", e => input.keyDown(e.code));
6    window.addEventListener("keyup", e => input.keyUp(e.code));
7    window.addEventListener("mousemove", e => input.mouseMove(e.clientX, e.clientY));
8    window.addEventListener("mousedown", e => input.mouseDown(e.button));
9    window.addEventListener("mouseup", e => input.mouseUp(e.button));
10   window.addEventListener("wheel", e => input.wheel(e.deltaY, { passive: true }));
11 }

```

Call it after `initResource(...)`.

4.2.4 3) Reset “pressed/released” flags once per frame

Add a phase/system that runs before gameplay update:

```
1 export function beginFrameSystem(w: WorldApi, _dt: number): void
2 {
3   w.requireResource(InputStateRes).beginFrame();
4 }
```

4.2.5 4) Read input from systems

Example “move player” system:

```
1 export function playerMoveSystem(w: WorldApi, dt: number): void
2 {
3   const input = w.requireResource(InputStateRes);
4
5   let dx = 0, dy = 0;
6   if (input.keysDown.has("KeyW")) dy -= 1;
7   if (input.keysDown.has("KeyS")) dy += 1;
8   if (input.keysDown.has("KeyA")) dx -= 1;
9   if (input.keysDown.has("KeyD")) dx += 1;
10
11  const speed = 220;
12
13  for (const { c1: tr } of w.query(Transform, PlayerTag)) {
14    tr.x += dx * speed * dt;
15    tr.y += dy * speed * dt;
16  }
17 }
```

4.2.6 5) Use `AssetCacheRes` in a render system (deduped async loads)

```
1 export function renderSpritesSystem(ctx: CanvasRenderingContext2D)
2 {
3   return (w: WorldApi, _dt: number): void => {
4     const assets = w.requireResource(AssetCacheRes);
5
6     for (const { c1: tr, c2: sp } of w.query(Transform, Sprite)) {
7       assets.getImage(sp.url).catch(() => {});
8       const img = assets.peekImage(sp.url);
9       if (!img) continue;
10
11       ctx.drawImage(img, tr.x, tr.y, sp.w, sp.h);
12     }
13   };
14 }
```

4.2.7 6) Run phases in order

Minimal schedule:

```
1 sched.add(world, "beginFrame", beginFrameSystem);
2 sched.add(world, "update", playerMoveSystem);
3 sched.add(world, "render", renderSpritesSystem(ctx));
```

Game loop:

```
1 sched.run(world, dt, ["beginFrame", "update", "render"]);
```

4.2.8 Common variations


Optional resource usage


If a resource is optional (debug/editor), use:

```
1  const dbg = w.getResource(DebugRes);  
2  if (dbg) dbg.enabled = true;
```

Preload assets (menu/loading screen)

```
1  await Promise.all(urls.map(u => world.requireResource(AssetCacheRes).getImage(u)));
```

 January 30, 2026

 January 9, 2026

4.3 How to add/remove components at runtime

1. Define your component types (classes):

```
1 class Position { constructor(public x = 0, public y = 0) {} }
2 class Velocity { constructor(public x = 0, public y = 0) {} }
```

1. Add/remove **immediately** when you are **not** iterating a query:

```
1 const e = world.spawn();
2 world.add(e, Position, new Position(0, 0));
3 world.add(e, Velocity, new Velocity(1, 0));
4
5 // Or add many at once
6 const z = world.spawn();
7 world.addMany(z, [Position, new Position(0, 0)], [Velocity, new Velocity(1, 0)])
8
9 world.remove(e, Velocity);
```

1. Add/remove **during a query/system** using **deferred commands**:

```
1 world.addSystem((w: any) => {
2   for (const { e, c1: pos } of w.query(Position)) {
3     if (pos.x > 10) w.cmd().add(e, Velocity, new Velocity(1, 0));
4     if (pos.x < 0) w.cmd().remove(e, Velocity);
5   }
6 });
7
8 // Or remove many at once
9 world.addSystem((w: any) => {
10  for (const { e, c1: pos } of w.query(Position, Velocity)) {
11    if (pos.x < 0) w.cmd().removeMany(e, Position, Velocity);
12  }
13 });
14
15 // apply queued structural changes
16 world.flush();
```

 March 16, 2026

 January 4, 2026

4.4 How to despawn entities safely

1. Despawn **immediately** when not iterating:


```
1 world.despawn(e);
```


1. Despawn **during a query/system** via `cmd()`:

```
1 world.addSystem((w: any) => {  
2   for (const { e, c1: pos } of w.query(Position)) {  
3     if (pos.x > 10) w.cmd().despawn(e);  
4   }  
5 });  
6  
7 // apply despawns  
8 world.flush();
```

1. Or rely on an end-of-frame flush:

```
1 world.update(dt); // runs systems, then flushes
```

 January 30, 2026

 January 4, 2026

4.5 How to filter queries with `with` and `without`

`QueryFilter` lets you narrow any query beyond the components you ask for by value.

```
1 type QueryFilter = {
2   readonly with?: ReadonlyArray<ComponentCtor<any>>;
3   readonly without?: ReadonlyArray<ComponentCtor<any>>;
4 };
```

Pass it as the **last argument** (or just before the callback for `queryEach`).

4.5.1 Exclude entities that have a component — `without`

Problem

You want to move every entity that has `Position` and `Velocity`, but not those that are frozen.

Solution

Add a `without: [FrozenTag]` filter. Archetypes that include `FrozenTag` are skipped entirely.

```
1 class FrozenTag {}
2
3 world.queryEach(Position, Velocity, { without: [FrozenTag] }, (e, pos, vel) => {
4   pos.x += vel.x * dt;
5   pos.y += vel.y * dt;
6 });
```

Entities that also carry `FrozenTag` are invisible to this query — no per-entity branch needed.

4.5.2 Require a marker without returning its value — `with`

Problem

You want to query only the player entity (tagged `PlayerTag`) but you don't need the tag value in your callback.

Solution

Use `with: [PlayerTag]`. The archetype must include `PlayerTag`, but the tag is not returned.

```
1 class PlayerTag {}
2
3 // query – row is { e, c1: Position, c2: Circle }, no c3 for PlayerTag
4 for (const { c1: pos, c2: circle } of world.query(Position, Circle, { with: [PlayerTag] })) {
5   // ...
6 }
7
8 // queryEach – callback receives (e, pos, circle), no PlayerTag argument
9 world.queryEach(Position, Circle, { with: [PlayerTag] }, (e, pos, circle) => {
10   // clamp player to viewport ...
11 });
```

This is cleaner than listing `PlayerTag` as a positional argument when you don't use it.

4.5.3 Combine with and without

Both fields can be used together:

```

1 // Entities that are Active and not Stunned
2 world.queryEach(Position, Velocity, { with: [Active], without: [Stunned] }, (e, pos, vel) => {
3     pos.x += vel.x * dt;
4     pos.y += vel.y * dt;
5 });

```

4.5.4 Apply a filter to queryTables

`queryTables` accepts the same filter as the last argument:

```

1 // Render active items (full colour)
2 for (const table of world.queryTables(Position, Circle, ItemKind, { without: [FrozenTag] })) {
3     for (let i = 0; i < table.entities.length; i++) {
4         draw(table.c1[i]!, table.c2[i]!, "full");
5     }
6 }
7
8 // Render frozen items (faded)
9 for (const table of world.queryTables(Position, Circle, ItemKind, { with: [FrozenTag] })) {
10    for (let i = 0; i < table.entities.length; i++) {
11        draw(table.c1[i]!, table.c2[i]!, "faded");
12    }
13 }

```

4.5.5 Freeze entities on a state transition

A common pattern: on game-over, mark all active falling items as frozen so movement systems skip them automatically.

```

1 // In a resolve system – called once when lives reach 0
2 if (game.lives === 0) {
3     game.running = false;
4
5     // `without: [FrozenTag]` prevents double-adding if this fires more than once
6     for (const { e } of world.query(FallingKindToken, { without: [FrozenTag] })) {
7         world.cmd().add(e, FrozenTag, new FrozenTag());
8     }
9 }
10
11 // In the movement system – frozen items are silently skipped
12 world.queryEach(Position, Velocity, { without: [FrozenTag] }, (e, pos, vel) => {
13     pos.x += vel.x * dt;
14     pos.y += vel.y * dt;
15 });

```

This removes the need for any `if (game.running)` guard inside movement loops.

4.5.6 What `queryFilter` does NOT do

- It does **not** return extra components — `with` is purely a filter, not a projection.
- It does **not** affect the `cN` numbering of returned components. Only positional constructor arguments count.
- It is checked at the **archetype level** — the filter is applied once per matching table, not once per entity, so there is no per-entity overhead.

🕒 March 11, 2026

🕒 March 11, 2026

4.6 How to have multiple Worlds (globe vs ground simulation)

1. Create two worlds:

```
1 const globeWorld = new World();
2 const groundWorld = new World();
```

1. Give each one its own schedule (recommended):

```
1 const globeSched = new Schedule();
2 const groundSched = new Schedule();
```

1. Run both each frame (same dt):

```
1 globeSched.run(globeWorld, dt, ["input", "sim", "render"]);
2 groundSched.run(groundWorld, dt, ["input", "sim", "render"]);
```

1. Share data **explicitly** between worlds (pick one):

2. copy values at a known point (end of `sim`, start of other `sim`)

3. or have a "bridge" step in your outer loop that reads from one world and writes into the other (via normal `add/set` or via `cmd()` + `flush()`)

🕒 January 14, 2026

🕒 January 4, 2026

4.7 How to integrate ECS into a game loop

4.7.1 Option A — Use `world.update(dt)`

1. Register systems with `addSystem(...)`
2. In your loop call:

```
1 function tick(dt: number) {  
2   world.update(dt); // runs systems, then flushes  
3 }
```

4.7.2 Option B — Use `Schedule` phases (recommended for games)

1. Build a schedule (`input`, `sim`, `render`)
2. In `requestAnimationFrame`:

```
1 let last = performance.now();  
2  
3 function frame(now: number) {  
4   const dt = (now - last) / 1000;  
5   last = now;  
6  
7   sched.run(world, dt, ["input", "sim", "render"]); // flush between phases  
8   renderer.render(scene, camera);  
9  
10  requestAnimationFrame(frame);  
11 }  
12  
13 requestAnimationFrame(frame);
```

 January 4, 2026

 January 4, 2026

4.8 How to run logic conditionally

4.8.1 Option A — Guard inside the system (simple)

1. Put a condition at the top:

```

1 let paused = false;
2
3 world.addSystem((w: any, dt: number) => {
4   if (paused) return;
5   for (const { c1: pos, c2: vel } of w.query(Position, Velocity)) {
6     pos.x += vel.x * dt;
7   }
8 });

```

4.8.2 Option B — Conditional phases (skip whole groups)

1. Maintain your phase list dynamically:

```

1 const base = ["input", "sim", "render"];
2
3 function getPhases(paused: boolean) {
4   return paused ? ["input", "render"] : base;
5 }
6
7 sched.run(world, dt, getPhases(paused));

```


4.8.3 Option C — Wrap systems (reuse predicates)

1. Make a helper:

```

1 const runIf = (pred: () => boolean, fn: (w: any, dt: number) => void) =>
2   (w: any, dt: number) => { if (pred()) fn(w, dt); };
3
4 world.addSystem(runIf(() => !paused, (w, dt) => {
5   for (const { c1: pos, c2: vel } of w.query(Position, Velocity)) {
6     pos.x += vel.x * dt;
7   }
8 }));

```

 January 4, 2026

 January 4, 2026

4.9 How to save and load a World

This guide shows a complete save/load flow:

1. register snapshot codecs
 2. create a snapshot
 3. persist to JSON
 4. restore later
-

4.9.1 1) Define data components/resources

```

1  class Position {
2      constructor(public x = 0, public y = 0) {}
3  }
4
5  class Health {
6      constructor(public hp = 100) {}
7  }
8
9  class GameState {
10     constructor(public wave = 1, public score = 0) {}
11 }
12
13 type SpawnConfig = { interval: number };
14 const SpawnConfigToken = (() => ({ interval: 1.0 })) as ComponentCtor<SpawnConfig>;

```

4.9.2 2) Register snapshot codecs once at boot

```

1  import { SnapshotCodec, World, type ComponentCtor } from "archetype-ecs-lib";
2
3  const world = new World();
4
5  const positionCodec: SnapshotCodec<Position, { x: number; y: number }> = {
6      key: "comp.position",
7      serialize: (v) => ({ x: v.x, y: v.y }),
8      deserialize: (d) => new Position(d.x, d.y),
9  };
10
11 const healthCodec: SnapshotCodec<Health, { hp: number }> = {
12     key: "comp.health",
13     serialize: (v) => ({ hp: v.hp }),
14     deserialize: (d) => new Health(d.hp),
15 };
16
17 const gameStateCodec: SnapshotCodec<GameState, { wave: number; score: number }> = {
18     key: "res.game-state",
19     serialize: (v) => ({ wave: v.wave, score: v.score }),
20     deserialize: (d) => new GameState(d.wave, d.score),
21 };
22
23 const spawnConfigCodec: SnapshotCodec<SpawnConfig, { interval: number }> = {
24     key: "res.spawn-config",
25     serialize: (v) => ({ interval: v.interval }),
26     deserialize: (d) => ({ interval: d.interval }),
27 };
28
29 world.registerComponentSnapshot(Position, positionCodec);
30 world.registerComponentSnapshot(Health, healthCodec);
31 world.registerResourceSnapshot(GameState, gameStateCodec);
32 world.registerResourceSnapshot(SpawnConfigToken, spawnConfigCodec);

```

Only registered types are persisted.

4.9.3 3) Create a snapshot and persist it

```

1  const snapshot = world.snapshot();
2  const json = JSON.stringify(snapshot);
3
4  localStorage.setItem("save-slot-1", json);

```

4.9.4 4) Restore from persisted JSON

```

1  const raw = localStorage.getItem("save-slot-1");
2  if (raw) {
3    const snapshot = JSON.parse(raw);
4    world.restore(snapshot);
5  }

```

`restore()` will clear runtime queues (pending commands/events) and reconstruct persisted entity/resource state.

4.9.5 5) Common pattern: baseline + quick save

```

1  const baseline = world.snapshot(); // after bootstrap
2  let quickSave: ReturnType<World["snapshot"]> | null = null;
3
4  function doQuickSave() {
5    quickSave = world.snapshot();
6  }
7
8  function doQuickLoad() {
9    if (quickSave) world.restore(quickSave);
10 }
11
12 function resetRun() {
13   world.restore(baseline);
14 }

```

4.9.6 6) Important checks

- Register codecs before calling `restore(...)`.
- Keep codec keys stable across versions.
- Ensure codec output is JSON-safe if you store saves as JSON.
- Re-register systems/schedules in app boot code (they are runtime behavior, not snapshot data).

🕒 February 8, 2026

🕒 February 8, 2026

4.10 How to split logic into multiple system phases

1. Create a `Schedule` and register systems by phase name:

```
1  const sched = new Schedule();
2
3  sched
4  .add(world, "input", (w: any) => { /* ... */ })
5  .add(world, "sim", (w: any, dt: number) => { /* ... */ })
6  .add(world, "render", (w: any) => { /* ... */ });
```

1. Define phase order:

```
1  const phases = ["input", "sim", "render"];
```

1. Run it each tick (flush happens after each phase):

```
1  sched.run(world, dt, phases);
```

🕒 January 30, 2026

🕒 January 4, 2026

4.11 How to use ECS alongside Three.js

4.11.1 Pattern: ECS owns state, Three.js owns objects

1. Keep Three.js objects in a map (outside ECS):

```
1 const meshes = new Map<number, THREE.Object3D>(); // key = entity.id
```

1. Add components for simulation and “render tag”:

```
1 class Position { constructor(public x=0, public y=0, public z=0) {} }
2 class Renderable { constructor(public kind: "cube" | "ship" = "cube") {} }
```

1. Spawn entities in ECS:

```
1 const e = world.spawnWith(
2   [Position, new Position(0, 0, 0)],
3   [Renderable, new Renderable("cube")],
4 )
```

1. Create a **render-sync** system in a `render` phase:

2. create missing meshes

3. update transforms

4. remove meshes for despawned entities (see step 5)

```
1 sched.add(world, "render", (w: any) => {
2   for (const { e, c1: pos, c2: rend } of w.query(Position, Renderable)) {
3     let obj = meshes.get(e.id);
4     if (!obj) {
5       obj = makeObjectFromKind(rend.kind); // your factory
6       scene.add(obj);
7       meshes.set(e.id, obj);
8     }
9     obj.position.set(pos.x, pos.y, pos.z);
10  }
11 });
```

1. Despawn visually **after flush**:

2. despawn in ECS via `cmd().despawn(e)`

3. after the flush boundary, remove from `meshes` if it's gone

A simple cleanup pass each frame:

```
1 for (const [id, obj] of meshes) {
2   // if you track alive entities externally, remove when not alive anymore.
3   // (One common approach: record seen IDs during the render query and remove the rest.)
4 }
```

🕒 March 23, 2026

🕒 January 4, 2026

4.12 How to use Events to decouple systems across phases

4.12.1 Goal

Emit events in one phase and consume them in a later phase, without coupling systems directly.

This guide assumes you already have a `Schedule` with multiple phases and that the schedule swaps events between phases.

4.12.2 1) Define event types

Use classes (recommended) or token keys.

```

1  export class DamageEvent {
2      constructor(public target: Entity, public amount: number) {}
3  }
4
5  export class PlaySoundEvent {
6      constructor(public id: string) {}
7  }

```

4.12.3 2) Emit events from a producer system

Example: gameplay system emits damage + sound.

```

1  function combatSystem(w: WorldApi, _dt: number) {
2      // ... detect hit
3      w.emit(DamageEvent, new DamageEvent(target, 10));
4      w.emit(PlaySoundEvent, new PlaySoundEvent("hit"));
5  }

```

4.12.4 3) Consume events in the next phase

Place a consumer in the **next** phase (phase-scoped delivery):

```

1  function applyDamageSystem(w: WorldApi, _dt: number) {
2      w.drainEvents(DamageEvent, (ev) => {
3          const hp = w.get(ev.target, Health);
4          if (!hp) return;
5          hp.value -= ev.amount;
6      });
7  }

```

Schedule order:

```

1  schedule.add(world, "update", combatSystem);
2  schedule.add(world, "afterUpdate", applyDamageSystem);

```

4.12.5 4) Deliver events to late phases (forwarding pattern)

With phase-scoped delivery, an event emitted in `update` is visible in `afterUpdate`. If you want it to reach `audio` several phases later, forward it:

```

1  function forwardSoundSystem(w: WorldApi, _dt: number) {
2    w.drainEvents(PlaySoundEvent, (ev) => {
3      w.emit(PlaySoundEvent, ev); // re-emit for the next phase
4    });
5  }
6
7  function audioSystem(w: WorldApi, _dt: number) {
8    w.drainEvents(PlaySoundEvent, (ev) => {
9      console.log("[audio] play:", ev.id);
10   });
11 }

```

Example pipeline:

```

1  schedule.add(world, "update", combatSystem); // emits PlaySoundEvent
2  schedule.add(world, "afterUpdate", forwardSoundSystem); // forwards -> render
3  schedule.add(world, "afterRender", forwardSoundSystem); // forwards -> audio
4  schedule.add(world, "audio", audioSystem); // consumes

```

4.12.6 5) Use `events(key).values()` for read-only inspection

If you need to check what's readable without consuming it:

```

1  const pending = w.events(DamageEvent).values();
2  if (pending.length > 0) {
3    // inspect (do not store array reference)
4  }

```

Prefer `drainEvents` for typical processing.

4.12.7 6) Clear events when resetting state

To clear one type:

```

1  w.clearEvents(DamageEvent);

```

To clear all readable event buffers:

```

1  w.clearEvents();

```

🕒 January 30, 2026

🕒 January 13, 2026

5. Reference

5.1 Archetypes

5.1.1 Purpose

An **archetype** is an internal storage “table” that groups together all entities sharing the **same set of component types**. Archetypes are the core performance mechanism of this ECS: queries match archetypes first, then iterate rows inside them.

5.1.2 Storage model

Table layout (SoA)

Archetypes store component data in **Structure of Arrays (SoA)** form:

- **one column per component type**
- each entity occupies a **row** across all columns

This is the reason queries are efficient: iteration is over dense arrays rather than scattered objects.

5.1.3 Archetype membership

Structural changes move entities between archetypes

When an entity’s component *set* changes, the entity moves to a different archetype:

- `add(e, Ctor, value)` is **structural** and *may move* the entity to another archetype
- `remove(e, Ctor)` is **structural** and *may move* the entity to another archetype

Non-structural updates do not change archetype membership:

- `set(e, Ctor, value)` updates the value but does not change the component set
-

5.1.4 Queries and archetypes

Archetype filtering

`query(...ctors)` only iterates archetypes that contain *all* required component columns, then yields matching entity rows.

Query row shape

For `query(A, B, C)`, the yielded row contains:

- `e` (entity handle)
 - `c1, c2, c3` component values in the same order as the ctor arguments
-

5.1.5 Safety constraints

Structural changes during iteration

While iterating queries (and generally while systems run), doing structural changes directly can throw. The recommended pattern is:

- enqueue structural changes via `world.cmd()`
- apply them via `world.flush()` (or at the end of `world.update(dt)`)

This matters because structural changes imply archetype moves.

5.1.6 Visibility / Public API

Archetypes are an **internal mechanism** (the public exports are `Types`, `TypeRegistry`, `Commands`, `World`, `Schedule`). Users interact with archetypes only indirectly through `World` operations and `query()`.

🕒 January 4, 2026

🕒 January 4, 2026

5.2 Commands

5.2.1 Purpose

`Commands` is a **deferred structural change buffer**. It lets you enqueue structural operations (spawn/despawn/add/remove) while iterating queries or running systems, then apply them later via `world.flush()` (or at the end of `world.update(dt)`).

5.2.2 How to obtain a `Commands` buffer

`world.cmd(): Commands`

`World.cmd()` returns a `Commands` instance you can use to enqueue operations.

Typical usage:

```

1  const cmd = world.cmd();
2
3  cmd.spawn((e) => {
4      cmd.add(e, Position, new Position(0, 0));
5  });
6
7  cmd.add(entity, Velocity, new Velocity(1, 0));
8  cmd.remove(entity, Velocity);
9  cmd.despawn(entity);
10
11 world.flush();

```

5.2.3 Supported operations

The command buffer supports these operations (as documented by the project):

`spawn(init?)`

Enqueues creation of a new entity.

- `init?: (e: Entity) => void` is an optional callback invoked with the spawned entity, typically used to enqueue `add()` calls for initial components.

`spawnWith(...items: ComponentCtorBundleItem[])`

Queues the creation of a new entity, along with its initial components, and applies everything on the next flush (within the same flush cycle).

- `...items: ComponentCtorBundleItem[]` is the list of components to add to the newly created entity.
- Internally, it iterates over the items and calls `add(e, ctor, value)` for each component.

```

1  world.cmd().spawnWith(
2      [Position, new Position(0, 0)],
3      [Velocity, new Velocity(1, 0)],
4  );

```

`despawn(e: Entity)`

Enqueues removal of an entity.

`despawnMany(entities: Entity[])`

Enqueues the destruction of multiple entities. The actual removals are applied when commands are flushed.

- `entities: Entity[]` is the list of entities to despawn.
- Internally, it iterates over the array and calls `despawn(e)` for each entity.

`add(e, ctor, value)`

Enqueues adding a component to an entity. This is a **structural** change (it may move the entity between archetypes), which is why it is commonly deferred.

`addMany(e: Entity, ...items: ComponentCtorBundleItem[])`

Enqueues adding multiple components to an existing entity as a **single batched command**. On flush, all components are applied in one archetype migration instead of one per component.

- `e: Entity` is the target entity.
- `...items: ComponentCtorBundleItem[]` is the list of components to add.
- No-op if `items` is empty.

`remove(e, ctor)`

Enqueues removing a component from an entity. This is also a **structural** change.

`removeMany(e: Entity, ...ctors: ComponentCtor<any>[])`

Enqueues removal of multiple component types from an entity as a **single batched command**. On flush, all components are removed in one archetype migration instead of one per component.

- `e: Entity` is the target entity.
- `...ctors: ComponentCtor<any>[]` is the list of component constructors (types) to remove.
- No-op if `ctors` is empty.

5.2.4 Reusable bundles

Use the `bundle()` helper to define a named, reusable group of components:

```

1  import { bundle } from "archetype-ecs-lib";
2
3  const PhysicsBundle = bundle(
4    [Position, new Position(0, 0)],
5    [Velocity, new Velocity(0, 0)],
6  );
7
8  // Spread into cmd or direct world calls
9  world.cmd().spawnWith(...PhysicsBundle);
10 world.spawnWith(...PhysicsBundle);
11 world.cmd().addMany(e, ...PhysicsBundle);

```

`bundle()` simply returns a typed `readonly ComponentCtorBundleItem[]`, avoiding the need for `as const` casts.

5.2.5 Applying commands

`world.flush(): void`

Applies all queued commands. `World.update(dt)` also flushes automatically at the end of the frame.

With `Schedule`

When using `Schedule`, `world.flush()` is called **after each phase**, creating deterministic “phase barriers” for command application.

5.2.6 Safety rule

Direct structural operations can throw while iterating queries or running systems. The intended pattern is:

- enqueue structural changes with `world.cmd()`
- apply them with `world.flush()` (or let `update()` do it)

 March 23, 2026

 January 4, 2026

5.3 Components

5.3.1 Purpose

A **component** is a unit of data attached to an `Entity`. In this ECS, components are stored in **archetypes (tables)** using a **Structure-of-Arrays (SoA)** layout: **one column per component type**.

5.3.2 Component “type” (key)

A component type is identified by a **constructor** (typically a class):

```
1 class Position { constructor(public x = 0, public y = 0) {} }
2 class Velocity { constructor(public x = 0, public y = 0) {} }
```

Any class used as a type key is considered a valid component type.

TypeId mapping

Internally, component constructors are mapped to a stable numeric `TypeId` via `typeId()`. `TypeId` assignment is **process-local** and based on **constructor identity** (via `WeakMap`).

5.3.3 Component “value”

The component value is the actual instance stored in the archetype column (e.g. `new Position(1,2)`).

- Values are stored per-archetype, per-column (SoA)
 - Queries return **direct references** to these values (you mutate them in place)
-

5.3.4 World operations on components

All component operations are done through `World` using the component constructor as the key.

Presence / access

- `has(e, Ctor): boolean`
- `get(e, Ctor): T | undefined`

Update (non-structural)

- `set(e, Ctor, value): void` Requires the component to exist; otherwise throws.

Structural changes

These may **move the entity between archetypes**:

- `add(e, Ctor, value): void`
 - `remove(e, Ctor): void`
-

5.3.5 Queries and component ordering

`world.query(A, B, C)` yields rows shaped like:

- `e`: the entity
- `c1`, `c2`, `c3`: component values in the **same order** as the ctor arguments

Example:


```
1 for (const { e, c1: pos, c2: vel } of world.query(Position, Velocity)) { }
```

5.3.6 Safety rules during iteration

While iterating a query (or while systems are running), **direct structural changes can throw**. Use deferred commands instead:

- enqueue via `world.cmd()`
- apply via `world.flush()`

 January 4, 2026

 January 4, 2026

5.4 Debugging & Profiling Reference

5.4.1 StatsOverlayOptions

Configuration options passed to `World` constructor.

```

1  type StatsOverlayOptions = Readonly<{
2    parent?: HTMLElement;
3    left?: number;
4    top?: number;
5    width?: number;
6    height?: number;
7    targetFrameMs?: number;
8    slowFrameMs?: number;
9    maxSamples?: number;
10 }>;

```

Option	Type	Default	Description
<code>parent</code>	<code>HTMLElement</code>	<code>document.body</code>	Container element for the overlay
<code>left</code>	<code>number</code>	<code>8</code>	Left offset in pixels
<code>top</code>	<code>number</code>	<code>8</code>	Top offset in pixels
<code>width</code>	<code>number</code>	<code>320</code>	Canvas width in pixels
<code>height</code>	<code>number</code>	<code>80</code>	Canvas height in pixels
<code>targetFrameMs</code>	<code>number</code>	<code>16.67</code>	Target frame time (white line)
<code>slowFrameMs</code>	<code>number</code>	<code>20</code>	Slow frame threshold (red bars)
<code>maxSamples</code>	<code>number</code>	<code>240</code>	Number of frames in history graph

5.4.2 WorldStats

Returned by `world.stats()`.

```

1  type WorldStats = Readonly<{
2    aliveEntities: number;
3    archetypes: number;
4    rows: number;
5    systems: number;
6    resources: number;
7    eventChannels: number;
8    pendingCommands: boolean;
9    frame: number;
10   dt: number;
11   frameMs: number;
12   phaseMs: Record<string, number>;
13   systemMs: Record<string, number>;
14 }>;

```

Field	Description
<code>aliveEntities</code>	Count of non-despawed entities
<code>archetypes</code>	Number of unique component combinations
<code>rows</code>	Total rows across all archetype tables
<code>systems</code>	Number of registered systems
<code>resources</code>	Number of stored resources
<code>eventChannels</code>	Number of active event channels
<code>pendingCommands</code>	Whether command queue has pending operations
<code>frame</code>	Current frame counter
<code>dt</code>	Delta time passed to last <code>update()</code> call
<code>frameMs</code>	Wall-clock time of last frame
<code>phaseMs</code>	Per-phase timing breakdown
<code>systemMs</code>	Per-system timing breakdown

5.4.3 WorldStatsHistory

Returned by `world.statsHistory()`.

```

1  type WorldStatsHistory = Readonly<{
2    capacity: number;
3    size: number;
4    dt: ReadonlyArray<number>;
5    frameMs: ReadonlyArray<number>;
6    phaseMs: Record<string, ReadonlyArray<number>>;
7    systemMs: Record<string, ReadonlyArray<number>>;
8  }>;

```

Field	Description
<code>capacity</code>	Maximum frames stored
<code>size</code>	Current number of frames stored
<code>dt</code>	Rolling history of delta times
<code>frameMs</code>	Rolling history of frame times
<code>phaseMs</code>	Per-phase timing history
<code>systemMs</code>	Per-system timing history

5.4.4 Profiling API

`world.setProfilingEnabled(enabled: boolean)`

Enable or disable a timing collection. When disabled, `frameMs`, `phaseMs`, and `systemMs` will be `0`.

`world.setProfilingHistorySize(frames: number)`

Set the rolling history capacity. Existing history is trimmed if reduced.

5.4.5 Overlay API

`world.setDebugging(enabled: boolean)`

Enable or disable the stats overlay.
When disabled, the method `world.destroyOverlay()` is called.

`world.destroyOverlay()`

Remove the stats overlay from the DOM and clean up event listeners.


`world.updateOverlay(stats, history)`

Manually trigger an overlay render. Called automatically by `world.update()` and `schedule.run()`.

5.4.6 Overlay UI Elements

Element	Function
Title bar	Drag to reposition overlay
/ button	Toggle console debug logging
-/+ button	Collapse/expand overlay content
Graph	Frame time history visualization
Legend	Target line, OK bars, Slow bars thresholds

 February 3, 2026

 February 3, 2026

5.5 Entity

5.5.1 Purpose

An **Entity** is a lightweight, opaque handle used to reference rows stored inside archetypes. It is **not** the data itself (components hold the data).

5.5.2 Type

```
1 type Entity = { id: number; gen: number };
```

- `id`: stable numeric slot identifier
 - `gen`: **generation counter** used to detect stale handles after despawn / reuse
-

5.5.3 Semantics

Identity

An entity handle is considered valid only if **both**:

- the `id` refers to an allocated slot
- the `gen` matches the current generation for that slot

Stale handles

If an entity is despawned and the `id` is later reused, the `gen` will differ. This prevents accidentally operating on “the new entity that reused the same id”.

5.5.4 Where entities come from

- `world.spawn()` returns an `Entity` handle
 - `world.query(...)` yields rows that include `e: Entity`
-

5.5.5 Where entities are used

Entities are passed into World operations (examples):

- lifecycle: `despawn(e)`
 - components: `add(e, Ctor, value)`, `remove(e, Ctor)`, `get(e, Ctor)`, `set(e, Ctor, value)`
 - commands (deferred): `cmd.despawn(e)`, `cmd.add(e, ...)`, `cmd.remove(e, ...)`
-

5.5.6 Related behavior

Safety during iteration

When iterating query results (which contain `e: Entity`), structural changes should be deferred via commands and applied with `flush()`.

🕒 January 4, 2026

🕒 January 4, 2026

5.6 Reference: Events API

5.6.1 Overview

Events are **typed, transient messages** used to decouple systems. They are stored per event type in **double-buffered channels**:

- `emit()` appends to the **write buffer** (current phase)
- `drain()` / `values()` read from the **read buffer** (previous phase)
- At each phase boundary, `world.swapEvents()` swaps buffers so events become visible to the next phase

Key type

Event channels are keyed by `ComponentCtor<T>` (same as components/resources). Keys are compared by identity.

5.6.2 `EventChannel<T>` (Events.ts)

`emit(ev: T): void`

Appends an event to the **write buffer** for the current phase.

Notes

- Emitted events are **not readable in the same phase**
 - They become readable after the next `swapBuffers()` / `world.swapEvents()`
-

`drain(fn: (ev: T) => void): void`

Iterates all **readable events** (read buffer) and then **clears** that buffer.

Semantics

- Reads only events emitted in the **previous phase**
- After `drain`, `count()` becomes `0`

Performance

- No iterator allocations; uses indexed loop
 - Clears with `length = 0`
-

`values(): readonly T[]`

Returns a read-only view of the **read buffer**.

Semantics

- Snapshot is valid until the next boundary swap
 - Do not store the returned array long-term
-

`count(): number`

Returns the number of readable events currently in the **read buffer**.

`clear(): void`

Clears the **read buffer** only.

`clearAll(): void`

Clears both **read** and **write** buffers.

`swapBuffers(): void (internal)`

Swaps read/write buffers and clears the new write buffer.

Semantics

- Makes events emitted in the previous phase readable now
 - Drops any undrained events from the prior read buffer at the next swap (phase-scoped delivery)
-

5.6.3 Delivery model summary (phase-scoped)

If you run phases:

`A -> B -> C`

Events emitted in **A** are readable in **B**. If not drained in **B**, they are dropped at `B -> C` swap.

🕒 January 13, 2026

🕒 January 13, 2026

5.7 Non goals

🕒 January 4, 2026

🕒 January 4, 2026

5.8 Query — Reference

5.8.1 Purpose

A **Query** iterates all entities that have **all required component types**, efficiently by scanning only the **matching archetypes (tables)**.

An optional `QueryFilter` can be appended to any query to refine which archetypes are visited:

- `with` — require additional component types to be present **without** returning their values
- `without` — exclude archetypes that have **any** of the listed component types

5.8.2 QueryFilter

```
1 type QueryFilter = {
2   readonly with?: ReadonlyArray<ComponentCtor<any>>;
3   readonly without?: ReadonlyArray<ComponentCtor<any>>;
4 };
```

Pass it as the **last argument** before the callback (for `queryEach`) or as the last argument overall (for `query` / `queryTables`):

```
1 // require Active, exclude Frozen – neither value is returned
2 world.query(Position, Velocity, { with: [Active], without: [Frozen] })
3
4 // exclude Dead without returning any extra component
5 world.queryEach(Position, { without: [Dead] }, (e, pos) => { ... })
```

Both fields are optional; omit the filter entirely when no filtering is needed.

5.8.3 API

`world.query(...ctors: ComponentCtor<any>[], filter?: QueryFilter): Iterable<any>`

`ctors` is a list of component constructors (types) you want to require.

```
1 for (const row of world.query(Position, Velocity)) {
2   // ...
3 }
```

Queries yield rows shaped like:

- `e`: the `Entity`
- `c1`, `c2`, `c3`, ...: component values **in the same order** as the `ctors` arguments

So `query(A, B, C)` yields `{ e, c1: A, c2: B, c3: C }`.

Row mapping and ordering

DETERMINISTIC COMPONENT FIELDS

The mapping is positional:

- `query(A) → { e, c1 }`
- `query(A, B) → { e, c1, c2 }`
- `query(A, B, C) → { e, c1, c2, c3 }`

`cN` always corresponds to the `N`th constructor you passed. Components listed inside a `QueryFilter` do **not** count toward `cN` indices and are not returned in the row.

Examples

Basic query:

```
1 for (const { e, c1: pos, c2: vel } of world.query(Position, Velocity)) {
2   pos.x += vel.x;
3   pos.y += vel.y;
4
5   // Safe structural change: defer it
6   if (pos.x > 10) world.cmd().despawn(e);
7 }
```

With a filter — require `Active`, exclude `Frozen`:

```
1 for (const { c1: pos, c2: vel } of world.query(Position, Velocity, { with: [Active], without: [Frozen] })) {
2   pos.x += vel.x;
3   pos.y += vel.y;
4 }
```

The row shape is still `{ e, c1: Position, c2: Velocity }` — the filter components are invisible in the result.

`world.queryTables(...ctors: ComponentCtor<any>[], filter?: QueryFilter): Iterable<any>`

`ctors` is a list of component constructors (types) you want to require.

```
1 for (const table of world.queryTables(Position, Velocity)) {
2   // ...
3 }
```

Queries yield **one item per matching archetype (table)** (SoA columns + entity array):

- `entities: Entity[]` (row-aligned with all columns)
- `c1, c2, c3, ...`: component **columns** (`T[]`) in the same order as `ctors`

So `queryTables(A, B, C)` yields `{ entities, c1: A[], c2: B[], c3: C[] }`.

WHY `queryTables`?

Use this when you want fewer allocations and more cache-friendly loops:

- you iterate **columns + indices** instead of creating one `{e, c1, ...}` object per entity
- you can batch work per archetype

Examples:

```
1 for (const { entities, c1: pos, c2: vel } of world.queryTables(Position, Velocity)) {
2   for (let i = 0; i < entities.length; i++) {
3     pos[i]!.x += vel[i]!.dx;
4     pos[i]!.y += vel[i]!.dy;
5   }
6 }
```

With a filter:

```
1 for (const { entities, c1: pos } of world.queryTables(Position, { without: [Frozen] })) {
2   for (let i = 0; i < entities.length; i++) {
3     pos[i]!.x += 1;
4   }
5 }
```

```
world.queryEach(...ctors, filter?: QueryFilter, fn): void
```

Callback-based query: no generator, no yielded row objects.

```
1 world.queryEach(Position, (e, pos) => {
2   // ...
3 });
4
5 world.queryEach(Position, Velocity, (e, pos, vel) => {
6   // ...
7 });
```

`queryEach(A, B, ...)` calls `fn(e, c1, c2, ...)` where `cN` matches the Nth constructor argument.

So `queryEach(A, B, C, fn)` calls `fn(e, A_value, B_value, C_value)` in that order.

WITH A FILTER

Place the `QueryFilter` **between the last component constructor and the callback**:

```
1 // Exclude Frozen – callback receives only Position and Velocity
2 world.queryEach(Position, Velocity, { without: [Frozen] }, (e, pos, vel) => {
3   pos.x += vel.dx;
4   pos.y += vel.dy;
5 });
6
7 // Require Active marker – callback still receives only Position
8 world.queryEach(Position, { with: [Active] }, (e, pos) => {
9   pos.x += 1;
10 });
```

Components inside the filter are **not** added as callback parameters.

WHY `queryEach` ?

Use this when you want the simplest “do work per entity” loop without generator overhead:

```
1 world.queryEach(Position, Velocity, (e, pos, vel) => {
2   pos.x += vel.dx;
3   pos.y += vel.dy;
4
5   // Safe structural change: defer it
6   if (pos.x > 10) world.cmd().despawn(e);
7 });
```

5.8.4 Safety rules during iteration

While iterating a query (or while systems are running), **structural changes** (spawn/despawn/add/remove) can throw.

Use:

- `world.cmd()` to defer changes
- `world.flush()` (or `world.update()`) to apply them safely

5.8.5 Performance

Throughput summary

Measured on a single machine with 50 000 entities (warmup excluded):

Query style	~ops/s
<code>queryTables</code>	~40M
<code>queryEach</code>	~21M
<code>query</code>	~4M

These numbers reflect the structural overhead each style carries, not just your application logic.

Why the gap?

`query` allocates one `{ e, c1, c2, ... }` object per entity per iteration. Even with stable V8 hidden classes (achieved via per-arity `switch` branches), GC pressure accumulates at scale.

`queryEach` skips the yield object entirely. The callback receives component values as direct arguments, so no heap allocation occurs in the inner loop. Generator protocol overhead is also eliminated.

`queryTables` yields one object per **archetype**, not per entity. You receive raw SoA columns (`c1: T[]`) and index into them yourself. This is the lowest-overhead loop possible and most cache-friendly for bulk reads.

Internal optimisations that benefit all three

- **Query result cache** — `_matchingArchetypes()` caches the list of matching archetypes per query signature. Because archetypes are append-only, only newly-created archetypes are re-checked on subsequent calls. The key is the sorted `TypeId[]` stringified by `signatureKey`.
- **Archetype edge cache** — `world.add()` / `world.remove()` cache the destination archetype on `addEdges` / `removeEdges` maps, skipping `signatureKey` lookups after the first transition. This reduces structural-change cost, not query cost, but it lowers the overhead of entity moves that feed into queries.

Decision guide

Use `queryTables` when: - you are processing large numbers of entities per frame (physics, particles, path-following) - you want to batch operations across all entities in an archetype (SIMD-style) - you do not need the entity handle (`e`) on every row

```
1 for (const { entities, c1: pos, c2: vel } of world.queryTables(Position, Velocity)) {
2   for (let i = 0; i < entities.length; i++) {
3     pos[i]!.x += vel[i]!.dx;
4     pos[i]!.y += vel[i]!.dy;
5   }
6 }
```

Use `queryEach` when: - you want straightforward per-entity logic without managing array indices - you need the entity handle on every row (e.g. to enqueue a deferred command) - you want better throughput than `query` without the `queryTables` index bookkeeping

```
1 world.queryEach(Position, Velocity, (e, pos, vel) => {
2   pos.x += vel.dx;
3   pos.y += vel.dy;
4   if (pos.x > bounds) world.cmd().despawn(e);
5 });
```

Use `query` when: - entity count is low (< ~1 000 entities matching) - you want to break early from iteration (`for...of` + `break` / `return`) - readability matters more than peak throughput

```
1 for (const { e, c1: pos } of world.query(Position)) {
2   if (someCondition(pos)) break; // generator supports early exit
3 }
```

Iteration count vs archetype count

All three styles iterate over **archetypes first**, then entities within each archetype. If your entities are spread across many archetypes (high fragmentation), each query pays a fixed per-archetype cost regardless of style. Keeping component combinations consolidated reduces this overhead.

🕒 March 12, 2026

🕒 January 4, 2026

5.9 Resources (Singletons / World Globals)

Resources are **typed singleton values stored on the world**, keyed by a `ComponentCtor<T>` (same “key shape” as components). They are **not attached to entities**.

They’re ideal for global state like **Time, Input, Asset caches, Config, RNG, Selection**, etc.

5.9.1 Concepts

What is a Resource?

A resource is a **single instance of data** stored globally in the ECS `World`.

- **Components** → many per world, attached to entities
- **Resources** → one per key, stored in the world

Key type: `ComponentCtor<T>`

All resource APIs use:

```
1 ComponentCtor<T>
```

This usually means:

- a **class constructor** (e.g. `class TimeRes { ... }`)
- or a **token function** (unique function used as a key)

Keys are compared by **identity** (reference equality), not by name.

5.9.2 API summary

All methods live on `World / WorldApi`.

```
1 setResource<T>(key: ComponentCtor<T>, value: T): void
2 getResource<T>(key: ComponentCtor<T>): T | undefined
3 requireResource<T>(key: ComponentCtor<T>): T
4 hasResource<T>(key: ComponentCtor<T>): boolean
5 removeResource<T>(key: ComponentCtor<T>): boolean
6 initResource<T>(key: ComponentCtor<T>, factory: () => T): T
```

Structural safety: resource operations are **not structural changes** (unlike spawn/despawn/add/remove). They do not require flushing and are safe to call during system execution.

5.9.3 Method reference

setResource<T>(key, value): void

Stores (or replaces) the resource value for `key`.

Behavior

- Overwrites any existing value.
- Does not flush and does not affect archetypes.

Example

```

1 class ConfigRes { constructor(public difficulty: "easy" | "hard") {} }
2
3 world.setResource(ConfigRes, new ConfigRes("hard"));

```

getResource<T>(key): T | undefined

Returns the resource value if present, otherwise `undefined`.

Use when

- the resource is **optional** (debug tools, plugins, editor-only state)

Important note

- If you **explicitly store** `undefined` as the value, this also returns `undefined`.
- Use `hasResource(key)` to distinguish:
 - “missing”
 - vs “present but undefined”

Example

```

1 const debug = world.getResource(DebugRes);
2 if (debug) debug.enabled = true;

```

requireResource<T>(key): T

Returns the resource value if present, otherwise **throws**.

Use when

- the resource is **required** for correct operation (Time, Input, AssetCache, Config)

Throws

- Error if missing

Example

```

1 const input = w.requireResource(InputStateRes);
2 if (input.keysDown.has("KeyW")) { /* ... */ }

```

hasResource<T>(key): boolean

Checks whether an entry exists for `key`.

Use when

- you need to distinguish missing vs present-but-undefined
- you want conditional initialization

Example

```

1 if (!world.hasResource(TimeRes)) {
2   world.setResource(TimeRes, new TimeRes());
3 }

```

```
removeResource<T>(key): boolean
```

Removes the resource entry for `key`.

Returns

- `true` if the entry existed and was removed
- `false` otherwise

Example

```
1 world.removeResource(DebugRes);
```

```
initResource<T>(key, factory): T
```

Insert-once helper.

Behavior

- If resource exists → returns existing value (factory is not called)
- If missing → calls `factory()`, stores, returns the new value

Use when

- bootstrapping default resources without double-init

Example

```
1 class TimeRes { dt = 0; elapsed = 0; }
2
3 world.initResource(TimeRes, () => new TimeRes());
```

5.9.4 Usage patterns

Pattern: “bootstrap required resources once”

```
1 class TimeRes { dt = 0; elapsed = 0; }
2 class InputStateRes { keysDown = new Set<string>(); }
3
4 world.initResource(TimeRes, () => new TimeRes());
5 world.initResource(InputStateRes, () => new InputStateRes());
```

Pattern: “systems read required resources”

```
1 function timeSystem(w: WorldApi, dt: number) {
2   const time = w.requireResource(TimeRes);
3   time.dt = dt;
4   time.elapsed += dt;
5 }
```

Pattern: “asset cache resource”

```
1 class AssetCacheRes {
2   images = new Map<string, HTMLImageElement>();
3 }
4
5 world.initResource(AssetCacheRes, () => new AssetCacheRes());
```

5.9.5 Gotchas

1) Keys must be stable and unique

Because keys are identity-based:

- `class TimeRes {}` used as key is stable
- a top-level `const TOKEN = (() => {})` as `ComponentCtor<T>` is stable
- creating a new token function inline each time won't match previous entries


2) Prefer `requireResource()` in gameplay systems

It keeps systems clean and fails fast when initialization is missing.

3) Resources are not entities

Do not use resources for data that should exist per-entity (that's components).

 January 14, 2026

 January 9, 2026

5.10 Save and Load (World Snapshot) Reference

This page documents the snapshot API on `World`.

5.10.1 Methods

`registerComponentSnapshot<T, D>(key, codec): this`

Registers a component snapshot codec.

- `key: ComponentCtor<T>` is the component token/class used in ECS storage.
- `codec: SnapshotCodec<T, D>` defines:
 - `key: string` stable snapshot type key (for example `"comp.position"`)
 - `serialize(value: T): D`
 - `deserialize(data: D): T`

Notes:

- `codec.key` must be non-empty.
 - Snapshot type keys must be unique per registry.
 - Calling again with the same ctor replaces the previous codec.
-

`unregisterComponentSnapshot<T>(key): boolean`

Removes a registered component snapshot codec.

Returns:

- `true` if removed
 - `false` if none existed
-

`registerResourceSnapshot<T, D>(key, codec): this`

Registers a resource snapshot codec.

Same rules as component registration.

`unregisterResourceSnapshot<T>(key): boolean`

Removes a registered resource snapshot codec.

`snapshot(): WorldSnapshot`

Exports world state.

Behavior:

- throws if called during iteration/structural lock
- flushes pending commands first
- serializes only registered components/resources
- returns deterministic ordering for entities/components/resources

`restore(snapshot: WorldSnapshot): void`

Loads world state.

Behavior:

- throws if called during iteration/structural lock
- validates snapshot format string
- clears pending commands
- clears event channel buffers
- clears current resources
- restores allocator + entities + registered data

`restore()` does not restore systems/schedules.

5.10.2 SnapshotCodec<T, D>

```
1 type SnapshotCodec<T, D = unknown> = {
2   key: string;
3   serialize(value: T): D;
4   deserialize(data: D): T;
5 };
```

Guidelines:

- keep `key` stable across versions
- output JSON-safe data if you persist via `JSON.stringify`
- avoid embedding runtime handles/functions in codec data

5.10.3 WorldSnapshot shape

```
1 type WorldSnapshot = {
2   format: "archetype-ecs/world-snapshot@1";
3   allocator: {
4     nextId: number;
5     free: number[];
6     generations: Array<readonly [id: number, gen: number]>;
7   };
8   entities: Array<{
9     id: number;
10    gen: number;
11    components: Array<{ type: string; data: unknown }>;
12  }>;
13   resources: Array<{ type: string; data: unknown }>;
14 };
```

5.10.4 Persisted vs non-persisted

Persisted:

- alive entities
- registered components
- registered resources
- allocator state

Not persisted:

- systems
 - schedules
 - events
 - command queue
-

5.10.5 Error conditions (common)

- unsupported snapshot format
 - duplicate snapshot type entries in payload
 - missing codec for a component/resource type during restore
 - invalid allocator/entity ids or generations
 - entity id marked both alive and free
-

5.10.6 Determinism guarantees

For the same world state and codec outputs:

- `entities` are ordered by entity id
- each entity `components` list is ordered by codec key
- `resources` list is ordered by codec key

🕒 February 8, 2026

🕒 February 8, 2026

5.11 Schedule

5.11.1 Purpose

`Schedule` is a **multiphase runner**.

It groups system functions under **named phases** (e.g. "input", "update", "render"), then executes those phases in a chosen order. At **phase boundaries**, it can automatically:

- flush deferred structural commands (`world.flush()`), only if commands are pending)
- deliver events to the next phase (`world.swapEvents()`)

This lets you build a deterministic pipeline (input → simulation → rendering → audio, etc.) without running into "structural change during iteration" problems.

5.11.2 Construction

```
1 ts
2 const schedule = new Schedule();
```

A `Schedule` is independent from `World`, you pass the `world` at run time.

5.11.3 Adding systems to phases

```
add(world: WorldApi, phase: string, fn: SystemFn): { after, before }
```

Registers `fn` under `phase` for the given `world`.

- You can register multiple systems under the same phase (they run in insertion order).
- Returns an object with `after()` and `before()` methods for chaining phase constraints.

Example:

```
1 ts
2 schedule
3   .add(world, "input", inputSystem)
4   .add(world, "update", updateSystem)
5   .add(world, "render", renderSystem);
```

5.11.4 Phase ordering constraints

Constraints are **phase-level** (not system-level): they affect the relative order of phases, not the order of systems within a phase.

```
after(otherPhase: string): this
```

Constrain the **most recently added phase** to run after `otherPhase`.

```
1 ts
2 schedule.add(world, "sim", simSystem).after("input"); // input -> sim
```

You can chain multiple constraints:

```
1 ts
2 schedule.add(world, "sim", simSystem).after("beginFrame").after("input");
```

`after()` must be called after `add(...)`. Calling it before any `add(...)` throws an error.

`before(otherPhase: string): this`

Constrain the **most recently added phase** to run before `otherPhase`.

```
1 ts
2 schedule.add(world, "input", inputSystem).before("sim"); // input -> sim
```

`before()` must be called after `add(...)`. Calling it before any `add(...)` throws an error.

5.11.5 Selecting a phase order

`Schedule.run()` chooses a phase order using the following precedence:

1. If `run(..., phaseOrder)` is provided, it is used as-is.
2. Else, if `setOrder([...])` was called, the stored order is used.
3. Else, an order is computed from `.after()` / `.before()` constraints (stable topological sort).

`setOrder(phases: string[]): this`

Set a default phase order used by `run(world, dt)` when no `phaseOrder` is passed.

```
1 ts
2 schedule.setOrder(["input", "sim", "render"]);
```

5.11.6 Phase boundary behavior

`setBoundaryMode(mode: "auto" | "manual"): this`

Controls what happens after each phase:

- "auto" (default):
 - if `world.cmd().hasPending()` → `world.flush()`
 - always `world.swapEvents()`
- "manual":
 - do nothing automatically; the caller is responsible for `world.flush()` / `world.swapEvents()`

```
1 ts
2 schedule.setBoundaryMode("auto"); // default
3 schedule.setBoundaryMode("manual"); // advanced usage
```

5.11.7 Running phases

`run(world: WorldApi, dt: number, phaseOrder?: string[]): void`

Runs the schedule for a single tick:

- Executes phases in the chosen order.
- Executes all systems registered under each phase.
- Applies phase boundary behavior according to `setBoundaryMode()`.

Example (explicit order):

```
1 ts
2 schedule.run(world, 1/60, ["input", "sim", "render"]);
```

Example (computed order from constraints):

```
1 ts
2 schedule
3   .add(world, "input", inputSystem)
4   .add(world, "sim", simSystem).after("input")
5   .add(world, "render", renderSystem).after("sim");
6
7 schedule.run(world, 1 / 60);
```

5.11.8 Errors and lifecycle notes

- **System errors:** If a system throws, `Schedule` rethrows a wrapped error with context:
 - "[phase=<phase> system=<name>] <message>"
- **Cyclic constraints:** If constraints contain a cycle and no explicit order is provided, `run()` throws.
- **No phase order:** If no phase order can be determined (no explicit order, no stored order, and nothing scheduled), `run()` throws:
 - `Schedule.run` requires a phase order (pass it as an argument or call `schedule.setOrder([...])`)

Lifecycle conflict detection

`Schedule.run()` and `World.update()` are **mutually exclusive** on the same `World` instance:

- If you register systems via `world.addSystem()` and then call `schedule.run()`, an error is thrown.
- If you register systems via `schedule.add()` and then call `world.update()`, an error is thrown.

This prevents confusing behavior from mixing two different system execution models in the same world.

Choose ONE approach:

Approach	Register systems with	Run with
Simple (single-phase)	<code>world.addSystem(fn)</code>	<code>world.update(dt)</code>
Multi-phase	<code>schedule.add(world, phase, fn)</code>	<code>schedule.run(world, dt, phases)</code>

5.11.9 Relationship to `World.update(dt)`

Feature	<code>World.update(dt)</code>	<code>Schedule.run(world, dt, phases)</code>
System registration	<code>world.addSystem(fn)</code>	<code>schedule.add(world, phase, fn)</code>
Phase support	Single implicit phase	Multiple named phases
Phase ordering	N/A	Via <code>after()/before()</code> or explicit order
Command flush	Once at end	After each phase (if pending)
Event swap	Once at end	After each phase
Best for	Simple game loops, prototyping	Complex pipelines, deterministic ordering

🕒 January 30, 2026

🕒 January 4, 2026

5.12 Systems

5.12.1 Purpose

A **system** is a function executed by the ECS to update simulation state (usually by iterating queries and mutating component values). Systems are registered on the `World`, and executed during `world.update(dt)`.

5.12.2 System function type

SystemFn

A system is a function with the signature:

- `(world: WorldApi, dt: number) => void`

In practice, examples call `query()` and `cmd()` inside systems, which are available through `WorldApi`.

5.12.3 Registering systems

`world.addSystem(fn): this`

Adds a system to the world.

- Systems run in the **order they were added** (as described by “runs systems in order”).

Example:

```

1 world.addSystem((w: any, dt: number) => {
2   for (const { e, c1: pos, c2: vel } of w.query(Position, Velocity)) {
3     pos.x += vel.x * dt;
4     pos.y += vel.y * dt;
5
6     if (pos.x > 10) w.cmd().despawn(e);
7   }
8 });

```

5.12.4 Running systems (frame execution)

`world.update(dt): void`

Runs one ECS frame:

1. Runs all registered systems (in order)
2. Flushes queued commands at the end

The reference summary explicitly lists:

- `addSystem(fn): this`
 - `update(dt): void` (*runs systems in order, then flushes*)
-

5.12.5 Structural changes inside systems

While systems are running (and while iterating queries), doing structural changes directly can throw. The recommended pattern is:

- enqueue structural changes with `world.cmd()`
 - apply them with `world.flush()` (or let `update()` do it at the end)
-

5.12.6 Systems in phases (Schedule)

If you need explicit ordering across *groups* of systems, use `Schedule` :

- `sched.add(phase, systemFn)`
- `sched.run(world, dt, phases)` runs phases in order and calls `world.flush()` after each phase

This provides deterministic “phase barriers” where deferred commands are applied.

🕒 January 6, 2026

🕒 January 4, 2026

5.13 World

5.13.1 Purpose

`World` is the **central authority** of the ECS. It owns and coordinates:

- entity lifecycle
- archetypes and component storage
- queries
- deferred structural commands
- system execution

There is **exactly one** `World` **instance per ECS context**.

5.13.2 Construction

```
1 const world = new World();
```

Side effects

- Initializes an empty entity pool
 - Initializes archetype storage
 - Initializes command buffer
 - Initializes system list
-

5.13.3 Entity Lifecycle API

`spawn(): Entity`

Creates a new entity immediately.

- Allocates a new entity id
- Marks entity as alive
- Places entity in the empty archetype

```
1 const e = world.spawn();
```

`spawnWith(...items: ComponentCtorBundleItem[]): Entity`

Creates a new entity along with its initial components immediately.

- `...items: ComponentCtorBundleItem[]` is the list of components to add to the newly created entity.
 - Internally, it iterates over the items and calls `add` for each component.
-

despawn(e: Entity): void

Immediately removes an entity.

- Invalidates the entity handle (`gen` mismatch)
- Removes the entity from its archetype
- Frees the slot for reuse

Throws if:

- entity is stale or not alive

despawnMany(entities: Entity[]): void

Immediately removes multiple entities.

- `entities: Entity[]` is the list of entities to despawn.
- Internally, it iterates over the array and calls `despawn(e)` for each entity.

isAlive(e: Entity): boolean

Checks whether an entity handle is still valid.

```
1 if (world.isAlive(e)) { ... }
```

destroy(): void

Releases all world state: archetypes, entities, resources, event channels, and systems. Any further structural use of the world after this call will throw.

```
1 world.destroy();
```

Intended for multi-world setups and hot-reload scenarios where you need to fully tear down a world instance.

5.13.4 Component API

All component types are identified by **constructor identity**.

has<T>(e: Entity, ctor: ComponentCtor<T>): boolean

Checks if an entity has a component.

get<T>(e: Entity, ctor: ComponentCtor<T>): T | undefined

Returns the component value or `undefined`.

- Does **not** throw if missing
- Returns `undefined` for stale entities

```
add<T>(e: Entity, ctor: ComponentCtor<T>, value: T): void
```

Adds a component to an entity.

- **Structural change**

- Moves the entity to a different archetype

Throws if:

- entity is stale
 - component already exists
 - structural changes are forbidden (see iteration rules)
-

```
addMany(e: Entity, ...items: ComponentCtorBundleItem[]): void
```

Adds multiple components to an existing entity in a **single archetype migration**.

- `e: Entity` is the target entity.
 - `...items: ComponentCtorBundleItem[]` is the list of components to add.
 - No-op if `items` is empty.
 - Prefer over multiple `add()` calls when adding several components at once.
-

```
remove<T>(e: Entity, ctor: ComponentCtor<T>): void
```

Removes a component.

- **Structural change**

- Moves the entity to a different archetype

Throws if:

- entity is stale
 - component does not exist
 - structural changes are forbidden
-

```
removeMany(e: Entity, ...ctors: ComponentCtor<any>[]): void
```

Removes multiple component types from an entity in a **single archetype migration**.

- `e: Entity` is the target entity.
 - `...ctors: ComponentCtor<any>[]` is the list of component constructors (types) to remove.
 - No-op if `ctors` is empty.
 - Prefer over multiple `remove()` calls when removing several components at once.
-

```
set<T>(e: Entity, ctor: ComponentCtor<T>, value: T): void
```

Updates an existing component value.

- **Non-structural**

- Does not change archetypes

Throws if:

- entity is stale
- component does not exist

5.13.5 Query API

`query(...ctors): Iterable<QueryRow>`

Iterates entities that contain **all requested components**.

```
1 for (const { e, c1, c2 } of world.query(A, B)) {
2   // e -> Entity
3   // c1 -> A
4   // c2 -> B
5 }
```

PROPERTIES

- Iterates archetypes, not entities
- Components are returned as `c1`, `c2`, ... in **argument order**
- Query iteration **locks structural changes**

5.13.6 Structural Change Rules

While iterating a query or running systems:

- ❌ `spawn`, `despawn`, `add`, `remove` are forbidden
- ✅ `get`, `set`, `has` are allowed

Violations throw a runtime error.

5.13.7 Command Buffer API

`cmd(): Commands`

Returns a command buffer for **deferred structural changes**.

```
1 world.cmd().despawn(e);
```

Commands are **queued**, not applied immediately.

`flush(): void`

Applies all queued commands.

- Safe to call after queries
- Automatically called by `update()` and `Schedule`

5.13.8 Snapshot / Restore API

`registerComponentSnapshot<T, D>(key: ComponentCtor<T>, codec: SnapshotCodec<T, D>): this`

Registers a component serializer/deserializer for snapshots.

```
unregisterComponentSnapshot<T>(key: ComponentCtor<T>): boolean
```

Removes a component snapshot registration.

```
registerResourceSnapshot<T, D>(key: ComponentCtor<T>, codec: SnapshotCodec<T, D>): this
```

Registers a resource serializer/deserializer for snapshots.

```
unregisterResourceSnapshot<T>(key: ComponentCtor<T>): boolean
```

Removes a resource snapshot registration.

```
snapshot(): WorldSnapshot
```

Exports world data state (entities + registered components/resources + allocator).

```
restore(snapshot: WorldSnapshot): void
```

Loads snapshot data into the world.

Runtime behavior (systems/schedule/events/queued commands) is not persisted.

5.13.9 System API

```
addSystem(fn: SystemFn): this
```

Registers a system.

```
1 world.addSystem((w, dt) => { ... });
```

Systems are executed **in insertion order**.

```
update(dt: number): void
```

Runs one ECS frame.

Execution order:

1. Run all systems
2. Flush deferred commands

```
1 world.update(1 / 60);
```

5.13.10 Events API

```
emit<T>(key: ComponentCtor<T>, ev: T): void
```

Emits an event of type T into the current phase write buffer.

`events<T>(key: ComponentCtor<T>): EventChannel<T>`

Returns the event channel for `key`, creating it if missing.

`drainEvents<T>(key: ComponentCtor<T>, fn: (ev: T) => void): void`

Drains readable events for the given type.

Behavior

- If the channel doesn't exist yet, it's a **no-op** (does not allocate/create)

`clearEvents<T>(key?: ComponentCtor<T>): void`

Clears readable events.

- If `key` is provided: clears that event type's **read buffer**
- If omitted: clears the **read buffers of all** event types

`swapEvents(): void (internal / schedule boundary)`

Swaps all event channels' buffers. Called by `schedule` at phase boundaries.

Required schedule behavior At each phase boundary:

```
1 world.flush();
2 world.swapEvents();
```

5.13.11 Internal Guarantees

- Archetypes use **Structure of Arrays (SoA)**
- Entity handles are generation-safe
- Component lookups are O(1) per archetype row
- Queries are archetype-filtered, not entity-scanned

5.13.12 Error Conditions (Summary)

Operation	Error Condition
add / remove	during query iteration
add	component already exists
remove	component missing
set	component missing
any	stale entity

5.13.13 Design Constraints

- Single-threaded

- No automatic conflict detection
- No parallel systems
- No borrowing model

These are **intentional** for simplicity and predictability.

🕒 March 23, 2026

🕒 January 4, 2026

6. Tutorials

6.1 Debugging Your First ECS Application

In this tutorial, you'll learn how to use the built-in debugging tools to understand and optimize your ECS application.

6.1.1 Prerequisites

- Basic familiarity with the ECS library
- A working browser-based project
- About 15 minutes

6.1.2 What You'll Build

A simple simulation with intentional performance issues that you'll identify and fix using the debugging tools.

6.1.3 Step 1: Create a World with Stats Overlay

Create a new file `debug-tutorial.ts`:

```

1  import { World } from "archetype-ecs-lib";
2
3  // Create world - overlay appears automatically
4  const world = new World({
5    statsOverlayOptions: {
6      width: 400,
7      height: 100
8    }
9  });
10 world.setDebugging(true);
11
12 console.log("World created! Look for the stats overlay in the top-left corner.");

```

Run your application. You should see the ECS Stats overlay appear.

Checkpoint: The overlay should show: - Frame 0 - Archetypes: 1 - Alive entities: 0

6.1.4 Step 2: Add Some Entities

Let's spawn some entities and watch the stats update:

```

1  // Components
2  class Position { constructor(public x = 0, public y = 0) {} }
3  class Velocity { constructor(public vx = 0, public vy = 0) {} }
4
5  // Spawn 100 moving entities
6  for (let i = 0; i < 100; i++) {
7    const e = world.spawn();
8    world.add(e, Position, new Position(Math.random() * 800, Math.random() * 600));
9    world.add(e, Velocity, new Velocity(Math.random() * 2 - 1, Math.random() * 2 - 1));
10 }
11
12 console.log("Spawned 100 entities");

```

Checkpoint: The overlay should now show: - Archetypes: 2 (empty + Position+Velocity) - Alive entities: 100 - Rows: 100

6.1.5 Step 3: Add a Movement System

```

1  function movementSystem(world: World, dt: number) {
2      for (const { e, c1: pos, c2: vel } of world.query(Position, Velocity)) {
3          pos.x += vel.vx * dt * 60;
4          pos.y += vel.vy * dt * 60;
5      }
6  }
7
8  world.addSystem(movementSystem);

```

6.1.6 Step 4: Create the Game Loop

```

1  let lastTime = performance.now();
2
3  function gameLoop() {
4      const now = performance.now();
5      const dt = (now - lastTime) / 1000;
6      lastTime = now;
7
8      world.update(dt);
9
10     requestAnimationFrame(gameLoop);
11 }
12
13 gameLoop();

```

Checkpoint: - Watch the frame counter increment - The graph should show blue bars (fast frames) - `frame=` should be very low (< 1ms)

6.1.7 Step 5: Introduce a Performance Problem

Let's add a "bad" system that creates performance issues:

```

1  function expensiveSystem(world: World, dt: number) {
2      // Simulate expensive computation
3      for (const { e, c1: pos } of world.query(Position)) {
4          // Intentionally slow: nested loop
5          let sum = 0;
6          for (let i = 0; i < 10000; i++) {
7              sum += Math.sin(pos.x + i) * Math.cos(pos.y + i);
8          }
9          (pos as any)._temp = sum; // Store result
10     }
11 }
12
13 world.addSystem(expensiveSystem);

```

Checkpoint: - Watch `frame=` time increase dramatically - Red bars should appear in the graph - Click to see phase timings in the console

6.1.8 Step 6: Identify the Slow System

Open your browser's developer console and click the button on the overlay to enable debug logging.

You'll see the output something like:

```

1  Phases: update=45.23ms

```

But we need per-system timing! The overlay shows this when using named functions:

Look at the overlay carefully. With profiling enabled, you'll see which system takes the most time.

6.1.9 Step 7: Fix the Performance Issue

Now that we've identified `expensiveSystem` as the problem, let's optimize it:

```

1 // Option 1: Process fewer entities per frame
2 let processedThisFrame = 0;
3 function optimizedExpensiveSystem(world: World, dt: number) {
4   processedThisFrame = 0;
5   for (const { e, c1: pos } of world.query(Position)) {
6     if (processedThisFrame++ > 10) break; // Only process 10 per frame
7
8     let sum = 0;
9     for (let i = 0; i < 10000; i++) {
10      sum += Math.sin(pos.x + i) * Math.cos(pos.y + i);
11    }
12    (pos as any)._temp = sum;
13  }
14 }
15
16 // Replace the bad system (in real code, you'd remove and re-add)

```

6.1.10 Step 8: Verify the Fix

After applying the optimization:

Checkpoint: - `frame=` time should drop significantly - Graph should return to mostly blue bars - Console logging should show lower phase times

6.1.11 Step 9: Explore Archetype Fragmentation

Let's intentionally cause archetype fragmentation:

```

1 class TempMarker {}
2
3 function fragmentingSystem(world: World, dt: number) {
4   for (const { e } of world.query(Position)) {
5     // Bad pattern: constantly adding/removing components
6     if (Math.random() > 0.5) {
7       if (!world.has(e, TempMarker)) {
8         world.add(e, TempMarker, new TempMarker());
9       }
10    } else {
11      if (world.has(e, TempMarker)) {
12        world.remove(e, TempMarker);
13      }
14    }
15  }
16 }
17
18 world.addSystem(fragmentingSystem);

```

Checkpoint: - Watch the "Archetypes" counter grow - This indicates archetype fragmentation

6.1.12 Step 10: Clean Up

Remove the debugging artifacts when going to production:

```

1 // Disable profiling for release builds
2 if (process.env.NODE_ENV === "production") {
3   world.setProfilingEnabled(false);
4   world.destroyOverlay();
5 }

```


6.1.13 Summary


In this tutorial, you learned how to:

1. Create a World with the stats overlay
2. Read entity and archetype counts
3. Identify slow systems using frame timing
4. Enable console debug logging
5. Recognize archetype fragmentation
6. Disable debugging tools for production

6.1.14 Next Steps

- Read the [Understanding ECS Debugging & Profiling](#) to understand what the metrics mean
- Check the [Debugging & Profiling Reference](#) for all available options
- Explore the [How to Debug Your ECS Application](#) for optimization techniques

 February 3, 2026

 February 3, 2026

6.2 Tutorial 1 - Your first ECS World

Outcome: you'll run a tiny simulation loop where entities with `Position` + `Velocity` move over time, using `World`, `spawn`, `add`, `query`, `addSystem`, and `update(dt)`.

6.2.1 1) What is an ECS? (one sentence)

ECS is a way to build simulations where **entities are IDs**, **components are data**, and **systems are functions that iterate entities with specific components**.

6.2.2 2) Create a tiny project

```
1 mkdir ecs-tutorial-1
2 cd ecs-tutorial-1
3 npm init -y
4 npm i archetype-ecs-lib
5 npm i -D typescript tsx
```

Install is `npm i archetype-ecs-lib`.

6.2.3 3) Create `tutorial1.ts`

Create a file named `tutorial1.ts` with this code:

```
1 import { World, WorldApi } from "archetype-ecs-lib";
2
3 // 1) Components = data (any class can be a component type)
4 class Position { constructor(public x = 0, public y = 0) {} }
5 class Velocity { constructor(public x = 0, public y = 0) {} }
6
7 // 2) Create a World (owns entities, components, systems)
8 const world = new World();
9
10 // 3) Spawn an entity and add components
11 const e = world.spawnWith(
12   [Position, new Position(0, 0)],
13   [Velocity, new Velocity(2, 0)], // 2 units/sec along x
14 )
15
16 // 4) Add a system (runs each update)
17 world.addSystem((w, dt) => {
18   for (const { e, c1: pos, c2: vel } of w.query(Position, Velocity)) {
19     pos.x += vel.x * dt;
20     pos.y += vel.y * dt;
21   }
22 });
23
24 // 5) Run a small simulation loop (60 frames)
25 const dt = 1 / 60;
26
27 for (let frame = 1; frame <= 60; frame++) {
28   world.update(dt);
29
30   // Read back Position and print it
31   const pos = world.get(e, Position!);
32   if (frame % 10 === 0) {
33     console.log(`frame ${frame}: x=${pos.x.toFixed(2)} y=${pos.y.toFixed(2)}`);
34   }
35 }
```

This uses the documented API:

- `spawn()`, `add(e, Ctor, value)`
- `addSystem(fn)`
- `query(Position, Velocity)` yielding `{ e, c1, c2 }`
- `update(dt)` to run systems each tick

6.2.4 4) Run it

```
1 npx tsx tutorial1.ts
```

You should see something like:

- frame 10: x=0.33 ...
- frame 60: x=2.00 ...

(Your exact decimals may differ slightly depending on rounding.)

6.2.5 5) You've built the core loop

You now have:

- a `World`
- entities created with `spawn()`
- components added with `add()`
- a system iterating `query(...)`
- a running simulation driven by `update(dt)`

🕒 March 23, 2026

🕒 January 4, 2026

6.3 Tutorial 2 - Components & archetypes

Outcome: you'll see how component sets automatically form **archetypes (tables)**, and how entities “move” between them when you `add()` / `remove()` components—without digging into internals. Archetypes store data in **SoA** (one column per component type).

6.3.1 1) Define a few component types

Create `tutorial2.ts`:

```
1 import { World } from "archetype-ecs-lib";
2
3 // Components are just data classes
4 class Position { constructor(public x = 0, public y = 0) {} }
5 class Velocity { constructor(public x = 0, public y = 0) {} }
6 class Health { constructor(public hp = 100) {} }
```

The ECS uses component constructors as the “type key”, and archetypes store entities in SoA tables.

6.3.2 2) Create a World and spawn entities with different component sets

```
1 const world = new World();
2
3 // e1 has: Position
4 const e1 = world.spawn();
5 world.add(e1, Position, new Position(1, 1));
6
7 // e2 has: Position + Velocity
8 const e2 = world.spawn();
9 world.add(e2, Position, new Position(0, 0));
10 world.add(e2, Velocity, new Velocity(1, 0));
11
12 // e3 has: Health
13 const e3 = world.spawn();
14 world.add(e3, Health, new Health(50));
```

6.3.3 3) Add a tiny helper to “see” matches

We can't (and don't need to) access archetype tables directly. Instead, we observe *which queries match*, before and after structural changes.

```
1 function ids(iter: Iterable<{ e: { id: number } }>): number[] {
2   const out: number[] = [];
3   for (const row of iter) out.push(row.e.id);
4   return out.sort((a, b) => a - b);
5 }
6
7 function dump(label: string) {
8   console.log(`\n=== ${label} ===`);
9   console.log("Position:", ids(world.query(Position)));
10  console.log("Velocity:", ids(world.query(Velocity)));
11  console.log("Health: ", ids(world.query(Health)));
12  console.log("Pos+Vel: ", ids(world.query(Position, Velocity)));
13  console.log("Pos+HP: ", ids(world.query(Position, Health)));
14 }
```

The query API yields `{ e, c1, c2, ... }` rows in the order you request components.

6.3.4 4) Observe the “automatic archetypes” effect

Add this and run once:

```
1 dump("initial");
```

You’ll see (by IDs) that:

- e1 matches `Position` only
- e2 matches both `Position` and `Pos+Vel`
- e3 matches `Health` only

What this demonstrates: entities with the **same component set** are stored together (same archetype). Archetypes are created implicitly as you introduce new component combinations.

6.3.5 5) Make an entity “move” between archetypes (add)

Now **add** a component to e1 :

```
1 world.add(e1, Velocity, new Velocity(0, 2));
2 dump("after: add Velocity to e1");
```

You should see:

- e1 now appears in `Velocity`
- and also in `Pos+Vel`

Why: `add()` is a **structural change** that can move an entity into a different archetype table (because its component set changed).

6.3.6 6) Make an entity “move” between archetypes (remove)

Now **remove** `Position` from e2 :

```
1 world.remove(e2, Position);
2 dump("after: remove Position from e2");
```

You should see:

- e2 disappears from `Position` and `Pos+Vel`
- e2 still appears in `Velocity`

Again: `remove()` is structural and can move the entity to a new archetype.

6.3.7 7) Run it

```
1 npx tsx tutorial2.ts
```

6.3.8 What you just learned (by doing)

- Components are plain data types (classes).
- Archetypes (tables) are created automatically for each distinct component set, stored in **SoA** layout.

- When you `add()` / `remove()` components, entities “move” because their component set changes (structural change).

Note for later tutorials: structural changes can be unsafe while iterating; that’s why `cmd()` + `flush()` exist.

 January 14, 2026

 January 4, 2026

6.4 Tutorial 3 - Deferred structural changes

Outcome: you'll learn the one rule that prevents most ECS bugs: **don't change entity structure while iterating**. You'll reproduce the problem safely, then fix it using **Commands** and **flush points** (via `Schedule`). The library explicitly supports this workflow: defer structural operations with `world.cmd()` and apply them with `world.flush()` / `Schedule` phase boundaries.

6.4.1 1) Create `tutorial4.ts`

```
1 import { World, WorldApi, Schedule } from "archetype-ecs-lib";
```

6.4.2 2) Define simple components

```
1 class Position { constructor(public x = 0) {} }
2 class Velocity { constructor(public x = 0) {} }
```

6.4.3 3) Setup: spawn a few movers

```
1 const world = new World();
2
3 function spawnMover(x: number, vx: number) {
4   const e = world.spawn();
5   world.add(e, Position, new Position(x));
6   world.add(e, Velocity, new Velocity(vx));
7   return e;
8 }
9
10 spawnMover(0, 2);
11 spawnMover(5, -3);
12 spawnMover(9, 1);
```

This is standard structural usage: `spawn()` + `add()`.

6.4.4 4) The unsafe thing (don't do this)

Add this function:

```
1 const unsafeDespawnInsideQuery: SystemFn = (w: WorldApi) => {
2   for (const { e, c1: pos } of w.query(Position)) {
3     if (pos.x > 8) {
4       // ❌ Structural change during iteration (may throw)
5       w.despawn(e);
6     }
7   }
8 }
```

Now call it once (inside a try/catch so the tutorial keeps going):

```
1 try {
2   unsafeDespawnInsideQuery(world);
3   console.log("unsafe: no error (but still not safe)");
4 } catch (err: any) {
5   console.log("unsafe: error as expected ->", String(err.message ?? err));
6 }
```

The lib will warn that structural changes during query iteration can throw and instructs to use `cmd()` + `flush()` instead.

6.4.5 5) The safe fix: use Commands

Replace the unsafe function with a safe one:

```

1  const safeDespawnInsideQuery: SystemFn = (w: WorldApi) => {
2    for (const { e, c1: pos } of w.query(RenderContextComponent)) {
3      if (pos.x > 8) {
4        // Defer structural change
5        w.cmd().despawn(e);
6      }
7    }
8  }

```

Commands let you queue:

- `spawn`, `despawn`, `add`, `remove`

6.4.6 6) Apply commands at a flush point

Option A — Manual flush

```

1  safeDespawnInsideQuery(world);
2  world.flush(); // apply queued despawns

```

`flush()` applies queued commands (and `update()` also flushes automatically at the end).

Option B — Flush at phase boundaries (recommended)

Use `Schedule`, which flushes after each phase:

```

1  const sched = new Schedule();
2
3  sched.add(world, "sim", (w: WorldApi) => {
4    // move
5    for (const { c1: pos, c2: vel } of w.query(Position, Velocity)) {
6      pos.x += vel.x;
7    }
8  });
9
10 sched.add(world, "cleanup", (w: WorldApi) => {
11   // safely despawn based on updated positions
12   safeDespawnInsideQuery(w);
13 });
14
15 // Flush happens after each phase automatically
16 const phases = ["sim", "cleanup"];

```

`Schedule.run(world, dt, phases)` runs phases and calls `world.flush()` after each phase.

6.4.7 7) Run a few ticks and print what's left

Add a small logger:

```

1  function logPositions(w: WorldApi, label: string) {
2    const items: string[] = [];
3    for (const { e, c1: pos } of w.query(Position)) {
4      items.push(`e${e.id}:${pos.x.toFixed(1)}`);
5    }
6    console.log(label, items.join(" | ") || "(none)");
7  }

```

Now run:

```

1  logPositions(world, "before");
2
3  for (let i = 0; i < 5; i++) {
4    sched.run(world, 0, phases);
5    logPositions(world, `after tick ${i + 1}`);
6  }

```

6.4.8 8) Full file (copy/paste)

```

1  import { World, Schedule } from "archetype-ecs-lib";
2
3  class Position { constructor(public x = 0) {} }
4  class Velocity { constructor(public x = 0) {} }
5
6  const world = new World();
7
8  function spawnMover(x: number, vx: number) {
9    const e = world.spawn();
10   world.add(e, Position, new Position(x));
11   world.add(e, Velocity, new Velocity(vx));
12   return e;
13 }
14
15 spawnMover(0, 2);
16 spawnMover(5, -3);
17 spawnMover(9, 1);
18
19 const unsafeDespawnInsideQuery: SystemFn = (w) => {
20   for (const { e, c1: pos } of w.query(Position)) {
21     if (pos.x > 8) {
22       w.despawn(e); // ❌ may throw
23     }
24   }
25 }
26
27 try {
28   unsafeDespawnInsideQuery(world as any);
29   console.log("unsafe: no error (but still not safe)");
30 } catch (err: any) {
31   console.log("unsafe: error as expected ->", String(err.message ?? err));
32 }
33
34 const safeDespawnInsideQuery: SystemFn = (w) => {
35   for (const { e, c1: pos } of w.query(Position)) {
36     if (pos.x > 8) w.cmd().despawn(e); // ✅ deferred
37   }
38 }
39
40 function logPositions(w: WorldApi, label: string) {
41   const items: string[] = [];
42   for (const { e, c1: pos } of w.query(Position)) {
43     items.push(`e${e.id}:${pos.x.toFixed(1)}`);
44   }
45   console.log(label, items.join(" | ") || "(none)");
46 }
47
48 const sched = new Schedule();
49
50 sched.add(world, "sim", (w: WorldApi) => {
51   for (const { c1: pos, c2: vel } of w.query(Position, Velocity)) {
52     pos.x += vel.x;
53   }
54 });
55
56 sched.add(world, "cleanup", (w: WorldApi) => {
57   safeDespawnInsideQuery(w);
58 });
59
60 const phases = ["sim", "cleanup"];
61
62 logPositions(world, "before");
63 for (let i = 0; i < 5; i++) {
64   sched.run(world, 0, phases); // flush after each phase
65   logPositions(world, `after tick ${i + 1}`);
66 }

```

6.4.9 9) Run it

```
1 npx tsx tutorial4.ts
```

You'll see:

- the unsafe version may throw (depending on timing/guarding)
- the safe version consistently despawns entities after they cross the threshold
- phase flush points make the timing predictable

🕒 January 30, 2026

🕒 January 4, 2026

6.5 Tutorial 4 - Writing systems

Outcome: you'll write real gameplay logic as **systems**: query components, mutate data safely, and run everything through a **Schedule** (input → sim → cleanup) with automatic `flush()` between phases.

6.5.1 1) Create `tutorial3.ts`

```
1 import { World, WorldApi, Schedule, SystemFn } from "archetype-ecs-lib";
```

The lib exports `World` and `Schedule`.

6.5.2 2) Define components (data only)

```
1 class Position { constructor(public x = 0, public y = 0) {} }
2 class Velocity { constructor(public x = 0, public y = 0) {} }
3 class Lifetime { constructor(public seconds = 1.0) {} } // despawn when <= 0
```

6.5.3 3) Create a `World` and spawn a few entities

```
1 const world = new World();
2
3 function spawnMover(x: number, y: number, vx: number, vy: number, life = 2.0) {
4   const e = world.spawn();
5   world.add(e, Position, new Position(x, y));
6   world.add(e, Velocity, new Velocity(vx, vy));
7   world.add(e, Lifetime, new Lifetime(life));
8   return e;
9 }
10
11 spawnMover(0, 0, 2, 0, 1.2);
12 spawnMover(0, 1, 1, 0, 2.5);
13 spawnMover(0, 2, -1, 0, 0.8);
```

This uses the documented structural ops: `spawn()` and `add()`.

6.5.4 4) System function signature (what you write)

A system is a function called like:

- `(world, dt) => void`

Systems are added using `world.addSystem()` like `world.addSystem((w: WorldApi, dt: number) => ...)`.

In this tutorial we'll register systems on a `Schedule` (phases), but the function shape is the same.

6.5.5 5) Write your first real system: movement

This system queries `Position + Velocity` and updates positions.

```
1 const movementSystem: SystemFn = (w: WorldApi, dt: number) => {
2   for (const { c1: pos, c2: vel } of w.query(Position, Velocity)) {
3     pos.x += vel.x * dt;
4     pos.y += vel.y * dt;
5   }
6 }
```

Query rows provide `{ e, c1, c2, ... }` in the same order as the query arguments.

6.5.6 6) Mutating data safely: despawn using commands

Despawning is a **structural change**, so do it through `cmd()` inside systems.

```
1 const lifetimeSystem: SystemFn = (w: WorldApi, dt: number) => {
2   for (const { e, c1: life } of w.query(Lifetime)) {
3     life.seconds -= dt;
4     if (life.seconds <= 0) {
5       w.cmd().despawn(e); // safe: deferred
6     }
7   }
8 }
```

6.5.7 7) Add a small “cleanup / log” system

We’ll print positions so you can see it running. This does not do structural changes.

```
1 const logSystem: SystemFn = (w: WorldApi, dt: number) => {
2   const lines: string[] = [];
3   for (const { e, c1: pos } of w.query(Position)) {
4     lines.push(`e${e.id} @ (${pos.x.toFixed(2)}, ${pos.y.toFixed(2)})`);
5   }
6   console.log(`frame ${frame}: ${lines.join(" | ")}`);
7 }
```

6.5.8 8) Run systems via world, ule (phases)

1. Create a schedule
2. Register systems under phases
3. Run phases each tick

```
1 const sched = new Schedule();
2
3 sched.add(world, "sim", movementSystem);
4 sched.add(world, "sim", lifetimeSystem);
5
6 // log in a separate phase so structural changes are already flushed
7 let frameNo = 0;
8 sched.add(world, "cleanup", (w: WorldApi) => {
9   frameNo++;
10  logSystem(w, frameNo);
11 });
12
13 const phases = ["sim", "cleanup"];
```

`Schedule.run(world, dt, phases)` runs phases in order and calls `world.flush()` after each phase.

6.5.9 9) Run the loop

```
1 const dt = 1 / 10; // bigger dt so it's easy to see
2 for (let i = 0; i < 20; i++) {
3   sched.run(world, dt, phases);
4 }
```

6.5.10 10) Full file (copy/paste)

```

1 import { World, WorldApi, Schedule, SystemFn } from "archetype-ecs-lib";
2
3 class Position { constructor(public x = 0, public y = 0) {} }
4 class Velocity { constructor(public x = 0, public y = 0) {} }
5 class Lifetime { constructor(public seconds = 1.0) {} }
6
7 const world = new World();
8
9 function spawnMover(x: number, y: number, vx: number, vy: number, life = 2.0) {
10   const e = world.spawn();
11   world.add(e, Position, new Position(x, y));
12   world.add(e, Velocity, new Velocity(vx, vy));
13   world.add(e, Lifetime, new Lifetime(life));
14   return e;
15 }
16
17 spawnMover(0, 0, 2, 0, 1.2);
18 spawnMover(0, 1, 1, 0, 2.5);
19 spawnMover(0, 2, -1, 0, 0.8);
20
21 const movementSystem: SystemFn = (w: WorldApi, dt: number) => {
22   for (const { c1: pos, c2: vel } of w.query(Position, Velocity)) {
23     pos.x += vel.x * dt;
24     pos.y += vel.y * dt;
25   }
26 }
27
28 const lifetimeSystem: SystemFn = (w: WorldApi, dt: number) => {
29   for (const { e, c1: life } of w.query(Lifetime)) {
30     life.seconds -= dt;
31     if (life.seconds <= 0) w.cmd().despawn(e);
32   }
33 }
34
35 const logSystem: SystemFn = (w: WorldApi, dt: number) => {
36   const lines: string[] = [];
37   for (const { e, c1: pos } of w.query(Position)) {
38     lines.push(`e${e.id} @ (${pos.x.toFixed(2)}, ${pos.y.toFixed(2)})`);
39   }
40   console.log(`frame ${frame}: ${lines.join(" | ")}`);
41 }
42
43 const sched = new Schedule();
44 sched.add(world, "sim", movementSystem);
45 sched.add(world, "sim", lifetimeSystem);
46
47 let frameNo = 0;
48 sched.add(world, "cleanup", (w: WorldApi) => {
49   frameNo++;
50   logSystem(w, frameNo);
51 });
52
53 const phases = ["sim", "cleanup"];
54
55 const dt = 1 / 10;
56 for (let i = 0; i < 20; i++) {
57   sched.run(world, dt, phases);
58 }

```

6.5.11 11) Run it

```
1 npx tsx tutorial3.ts
```

You'll see entities moving, then disappearing as their `Lifetime` reaches 0 (despawned safely via commands + phase flush).

🕒 January 30, 2026

🕒 January 4, 2026

6.6 Tutorial 5 - ECS + Three.js (render-sync + safe spawn/despawn)

Outcome: you'll see **moving cubes** in Three.js. You'll also **spawn** new cubes on click and **despawn** them safely using `cmd()` + **phase flush boundaries** (via `Schedule`).

6.6.1 1) Create a new project

```
1 mkdir ecs-threejs-tutorial
2 cd ecs-threejs-tutorial
3 npm init -y
4
5 npm i archetype-ecs-lib three
6 npm i -D vite typescript
```

The ECS package is installed as `archetype-ecs-lib`.

6.6.2 2) Add `index.html`

Create `index.html`:

```
1 <!doctype html>
2 <html lang="en">
3   <head>
4     <meta charset="UTF-8" />
5     <meta name="viewport" content="width=device-width, initial-scale=1.0" />
6     <title>ECS + Three.js Tutorial</title>
7     <style>
8       html, body { margin: 0; height: 100%; overflow: hidden; }
9       #hud {
10        position: fixed; left: 12px; top: 12px;
11        padding: 8px 10px; border-radius: 8px;
12        background: rgba(0,0,0,0.55); color: #fff;
13        font-family: system-ui, sans-serif; font-size: 13px;
14        user-select: none;
15      }
16     </style>
17   </head>
18   <body>
19     <div id="hud">Click to spawn cubes</div>
20     <script type="module" src="/src/main.ts"></script>
21   </body>
22 </html>
```

6.6.3 3) Add src/main.ts

Create src/main.ts:

```

1  import * as THREE from "three";
2  import { World, WorldApi, Schedule, SystemFn } from "archetype-ecs-lib";
3
4  // -----
5  // Components (data only)
6  // -----
7  class Position { constructor(public x = 0, public y = 0, public z = 0) {} }
8  class Velocity { constructor(public x = 0, public y = 0, public z = 0) {} }
9  class Lifetime { constructor(public seconds = 2.0) {} }
10 class Renderable { constructor(public kind: "cube" = "cube") {} }
11
12 // -----
13 // Three.js setup
14 // -----
15 const renderer = new THREE.WebGLRenderer({ antialias: true });
16 renderer.setSize(window.innerWidth, window.innerHeight);
17 document.body.appendChild(renderer.domElement);
18
19 const scene = new THREE.Scene();
20 scene.background = new THREE.Color(0x101018);
21
22 const camera = new THREE.PerspectiveCamera(
23   60,
24   window.innerWidth / window.innerHeight,
25   0.1,
26   1000
27 );
28 camera.position.set(0, 6, 14);
29 camera.lookAt(0, 0, 0);
30
31 const grid = new THREE.GridHelper(40, 40);
32 scene.add(grid);
33
34 window.addEventListener("resize", () => {
35   camera.aspect = window.innerWidth / window.innerHeight;
36   camera.updateProjectionMatrix();
37   renderer.setSize(window.innerWidth, window.innerHeight);
38 });
39
40 // -----
41 // ECS setup
42 // -----
43 const world = new World();
44 const sched = new Schedule();
45
46 // Map ECS entities -> Three.js objects
47 // Use id+gen so reuse of ids never points to the wrong mesh.
48 const entityKey = (e: { id: number; gen: number }) => `${e.id}:${e.gen}`;
49 const objects = new Map<string, THREE.Object3D>();
50
51 function makeObject(r: Renderable): THREE.Object3D {
52   // MeshNormalMaterial doesn't need lights
53   const geo = new THREE.BoxGeometry(1, 1, 1);
54   const mat = new THREE.MeshNormalMaterial();
55   const mesh = new THREE.Mesh(geo, mat);
56   mesh.castShadow = false;
57   mesh.receiveShadow = false;
58   return mesh;
59 }
60
61 // -----
62 // Input: click to spawn
63 // -----
64 let pendingClicks = 0;
65 window.addEventListener("pointerdown", () => pendingClicks++);
66
67 const inputPhase: SystemFn = (w: WorldApi) => {
68   if (pendingClicks <= 0) return;
69
70   const cmd = w.cmd();
71   for (let i = 0; i < pendingClicks; i++) {
72     cmd.spawn((e: any) => {
73       // Spawn near origin with random velocity and short lifetime
74       const x = (Math.random() - 0.5) * 8;
75       const z = (Math.random() - 0.5) * 8;
76
77       const vx = (Math.random() - 0.5) * 6;
78       const vz = (Math.random() - 0.5) * 6;
79
80       cmd.addMany(
81         e,
82         [Position, new Position(x, 0.5, z)],
83         [Velocity, new Velocity(vx, 0, vz)],
84         [Lifetime, new Lifetime(2.0 + Math.random() * 2.0)],
85         [Renderable, new Renderable("cube")],
86       );
87     });
88   }
89   pendingClicks = 0;
90 }
91
92 // -----
93 // Simulation: movement
94

```

```

95 // -----
96 const movementSystem: SystemFn = (w: WorldApi, dt: number) => {
97   for (const { c1: pos, c2: vel } of w.query(Position, Velocity)) {
98     pos.x += vel.x * dt;
99     pos.y += vel.y * dt;
100    pos.z += vel.z * dt;
101
102    // simple bounds bounce
103    const limit = 12;
104    if (pos.x < -limit || pos.x > limit) vel.x *= -1;
105    if (pos.z < -limit || pos.z > limit) vel.z *= -1;
106   }
107 }
108
109 // -----
110 // Simulation: lifetime -> despawn (deferred)
111 // -----
112 const lifetimeSystem: SystemFn = (w: WorldApi, dt: number) => {
113   for (const { e, c1: life } of w.query(Lifetime)) {
114     life.seconds -= dt;
115     if (life.seconds <= 0) {
116       w.cmd().despawn(e); // structural change deferred
117     }
118   }
119 }
120
121 // -----
122 // Render phase: ECS -> Three.js sync + remove despawned
123 // -----
124 const renderSync: SystemFn = (w: WorldApi) => {
125   const alive = new Set<string>();
126
127   // Create/update objects for all renderables
128   for (const { e, c1: pos, c2: rend } of w.query(Position, Renderable)) {
129     const key = entityKey(e);
130     alive.add(key);
131
132     let obj = objects.get(key);
133     if (!obj) {
134       obj = makeObject(rend);
135       scene.add(obj);
136       objects.set(key, obj);
137     }
138
139     obj.position.set(pos.x, pos.y, pos.z);
140   }
141
142   // Remove objects whose entities are gone
143   for (const [key, obj] of objects) {
144     if (!alive.has(key)) {
145       scene.remove(obj);
146       objects.delete(key);
147     }
148   }
149 }
150
151 // -----
152 // Schedule: phase ordering + flush boundaries
153 // -----
154 // Schedule runs phases in order and calls world.flush() after each phase. :contentReference[oaicite:2]{index=2}
155 sched.add(world, "input", inputPhase);
156 sched.add(world, "sim", movementSystem);
157 sched.add(world, "sim", lifetimeSystem);
158 sched.add(world, "render", renderSync);
159
160 const phases = ["input", "sim", "render"];
161
162 // -----
163 // Animation loop
164 // -----
165 let last = performance.now();
166
167 function frame(now: number) {
168   const dt = Math.min(0.05, (now - last) / 1000);
169   last = now;
170
171   // Run ECS phases (flush after each phase)
172   sched.run(world, dt, phases);
173
174   // Render Three.js
175   renderer.render(scene, camera);
176
177   requestAnimationFrame(frame);
178 }
179
180 requestAnimationFrame(frame);
181
182 // Spawn a few cubes at start (via commands so it behaves like gameplay)
183 pendingClicks = 6;

```

This uses:

- `Schedule.add(phase, fn)` and `Schedule.run(world, dt, phases)` with flush between phases
- `world.query(A, B)` yielding `{ e, c1, c2 }`
- deferred structural ops via `cmd().spawn/add/despawn`

6.6.4 4) Run it

```
1 npx vite
```

Open the local URL Vite prints.

6.6.5 5) What to try (hands-on)

1. **Click** to spawn cubes
2. Watch them move (ECS `Position + Velocity`)
3. Watch them disappear after a few seconds (`Lifetime` → `cmd().despawn(e)`), and see the Three.js mesh removed automatically in `renderSync()`.

That's the "safe loop":

- structural requests happen inside systems using `cmd()`
- structural changes become real at flush boundaries (Schedule phases)
- render-sync runs on a stable world snapshot

🕒 March 16, 2026

🕒 January 4, 2026

6.7 Tutorial 6 - Save and load

Outcome: you will build a tiny save/load loop with `World.snapshot()` and `World.restore()`.

6.7.1 1) Create `tutorial6.ts`

```

1  import {
2    World,
3    type ComponentCtor,
4    type SnapshotCodec,
5    type WorldSnapshot
6  } from "archetype-ecs-lib";

```

6.7.2 2) Define components and resources

```

1  class Position {
2    constructor(public x = 0, public y = 0) {}
3  }
4
5  class Health {
6    constructor(public hp = 100) {}
7  }
8
9  class RunState {
10   constructor(public wave = 1, public score = 0) {}
11 }
12
13 type SpawnRules = { cooldown: number };
14 const SpawnRulesToken = (( ) => ( { cooldown: 1.0 } )) as ComponentCtor<SpawnRules>;

```

6.7.3 3) Create world and register snapshot codecs

```

1  const world = new World();
2
3  const positionCodec: SnapshotCodec<Position, { x: number; y: number }> = {
4    key: "comp.position",
5    serialize: (v) => ( { x: v.x, y: v.y } ),
6    deserialize: (d) => new Position(d.x, d.y),
7  };
8
9  const healthCodec: SnapshotCodec<Health, { hp: number }> = {
10   key: "comp.health",
11   serialize: (v) => ( { hp: v.hp } ),
12   deserialize: (d) => new Health(d.hp),
13 };
14
15 const runStateCodec: SnapshotCodec<RunState, { wave: number; score: number }> = {
16   key: "res.run-state",
17   serialize: (v) => ( { wave: v.wave, score: v.score } ),
18   deserialize: (d) => new RunState(d.wave, d.score),
19 };
20
21 const spawnRulesCodec: SnapshotCodec<SpawnRules, { cooldown: number }> = {
22   key: "res.spawn-rules",
23   serialize: (v) => ( { cooldown: v.cooldown } ),
24   deserialize: (d) => ( { cooldown: d.cooldown } ),
25 };
26
27 world.registerComponentSnapshot(Position, positionCodec);
28 world.registerComponentSnapshot(Health, healthCodec);
29 world.registerResourceSnapshot(RunState, runStateCodec);
30 world.registerResourceSnapshot(SpawnRulesToken, spawnRulesCodec);

```

6.7.4 4) Bootstrap some game state

```

1 world.setResource(RunState, new RunState(3, 1200));
2 world.setResource(SpawnRulesToken, { cooldown: 0.5 });
3
4 const e = world.spawn();
5 world.add(e, Position, new Position(10, 20));
6 world.add(e, Health, new Health(75));

```

6.7.5 5) Save

```

1 const snapshotA = world.snapshot();
2 const json = JSON.stringify(snapshotA);
3 console.log("Saved bytes:", json.length);

```

6.7.6 6) Mutate state (simulate gameplay)

```

1 const pos = world.get(e, Position!);
2 pos.x = 99;
3 pos.y = 42;
4
5 world.requireResource(RunState).score += 300;

```

6.7.7 7) Load

```

1 const loaded = JSON.parse(json) as WorldSnapshot;
2 world.restore(loaded);

```

After restore, you get the original values from step 4:

```

1 console.log(world.get(e, Position)); // Position { x: 10, y: 20 }
2 console.log(world.requireResource(RunState)); // RunState { wave: 3, score: 1200 }

```

6.7.8 8) Useful pattern: baseline + quick save


```


1 const baseline = world.snapshot();
2 let quickSave: WorldSnapshot | null = null;
3
4 function saveNow() {
5     quickSave = world.snapshot();
6 }
7
8 function loadNow() {
9     if (quickSave) world.restore(quickSave);
10 }
11
12 function resetRun() {
13     world.restore(baseline);
14 }

```

6.7.9 9) What to remember

- Register codecs before calling `restore`.
- Save files are plain data; systems/schedules are rebuilt by normal app boot code.
- Keep codec keys stable (`"comp.*"` / `"res.*"`) for long-term compatibility.

 February 8, 2026

 February 8, 2026