RAM Simulation — availability, failures and downtime of the drawn system, with repairs and PM

Menu: System ▸ RAM Simulation...

Part 1 — What the function does

What it calculates. RAM stands for reliability, availability and maintainability. You have drawn a system on the reliability block diagram (RBD) canvas: blocks in series, in parallel, in k-of-n groups, in standby groups or in a bridge, each block with a life distribution. RAM Simulation plays the life of that system through a mission many times - the replications. In each replication it draws a life for every block, repairs failed blocks, renews blocks at their preventive maintenance (PM) age, and records when the system is up and when it is down. From all replications together it calculates:

Result	How ReliaEn sets it
Availability — mean (SE)	The up-time of each replication divided by the mission time T, averaged; SE is its standard error. An empirical band at your confidence level shows how far one mission can fall from the average
Mission reliability R̂(T)	The fraction of replications with no system failure, with exact Clopper–Pearson bounds
System failures — mean, System downtime — mean (h), MTBF_sys	The system's failures and down hours per mission, and its up-time per failure
PER-UNIT RESULTS	Each unit's failures, PM renewals, downtime and share of the system's downtime
ECONOMICS	The mission's cost, when you type at least one cost
BUILT-IN IDENTITY CHECKS — THE SIMULATION AUDITED AGAINST THE VERIFIED ANALYTIC ENGINES	The run compared with an exact analytic result wherever one exists, each check agree or OUTSIDE — do not use this run

Every number reproduces from the seed: the same diagram, form entries and seed give the same report, digit for digit.

From which data. The diagram on the RBD canvas, with each block's life distribution. The form adds the mission time, the number of replications, the seed, the confidence, a repair model and a PM age for each block, the number of repair crews and, optionally, three costs.

When an engineer uses it. When repairs, a limited number of repair crews, age-based PM or costs change the answer - anything the analytic RBD evaluation leaves out.

When not to use it.

Situation	Use instead
Only the reliability of a non-repairable system at one mission time	Evaluate the diagram on System ▸ RBD Canvas (chapter 4.1)
A set of states with constant failure and repair rates	System ▸ Markov Availability... (chapter 4.3)
How many redundant copies of each block to buy	System ▸ Reliability Optimization... (chapter 4.6)
The size of a spares stock	System ▸ Spare Parts Optimization... (chapter 5.3)

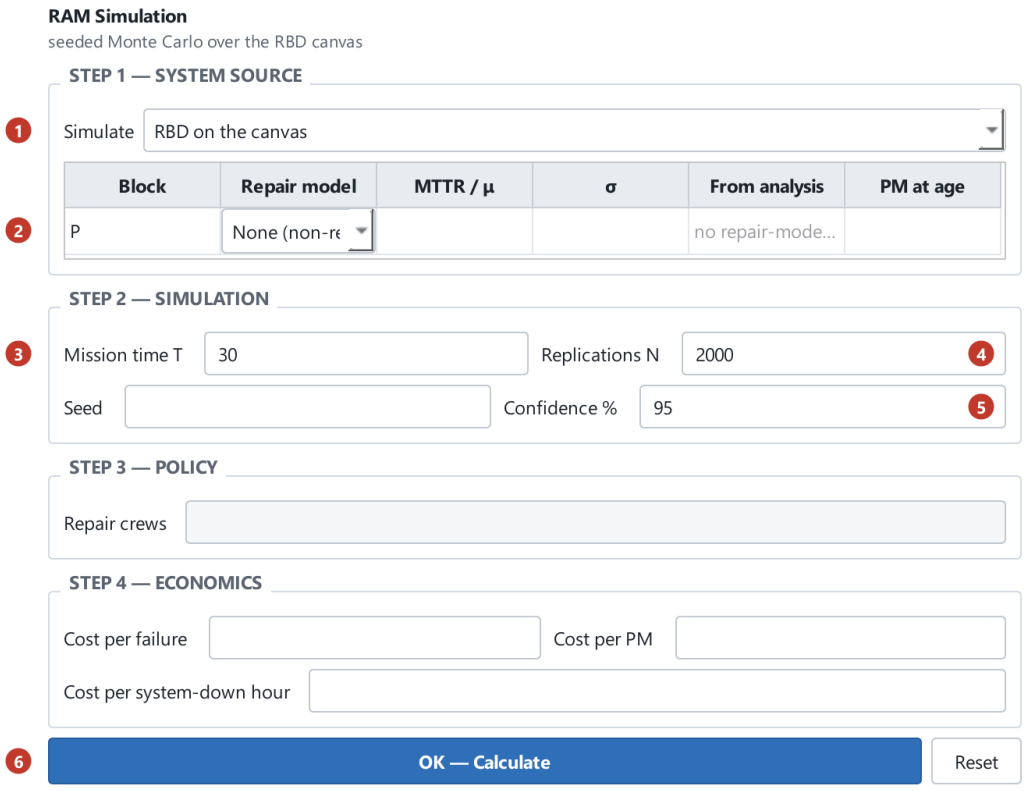
Part 2 — The example's input data

Setting. A cooling water station has three identical pumps; the process needs two. Each pump has a constant failure rate of 0.02 per hour. There is no repair during a 30-hour run. The engineer wants the chance that the station keeps two pumps for the whole run, and checks the simulation against the exact answer before adding repairs.

Input	Value
RBD canvas	one k-of-n block P: 3 units, 2 needed, Exponential (λ) 0.02, connected between start and end
Mission time T	30 (hours)
Replications N	2000
Seed	blank (the house seed 20260825)
Confidence %	95
Repair model, row P	None (non-repairable)
PM at age, Repair crews, the three costs	blank

No data file: the system is drawn on the canvas and the values are typed into the form as Part 3 shows.

The form set for this example, as Part 3's steps 3 to 7 set it:

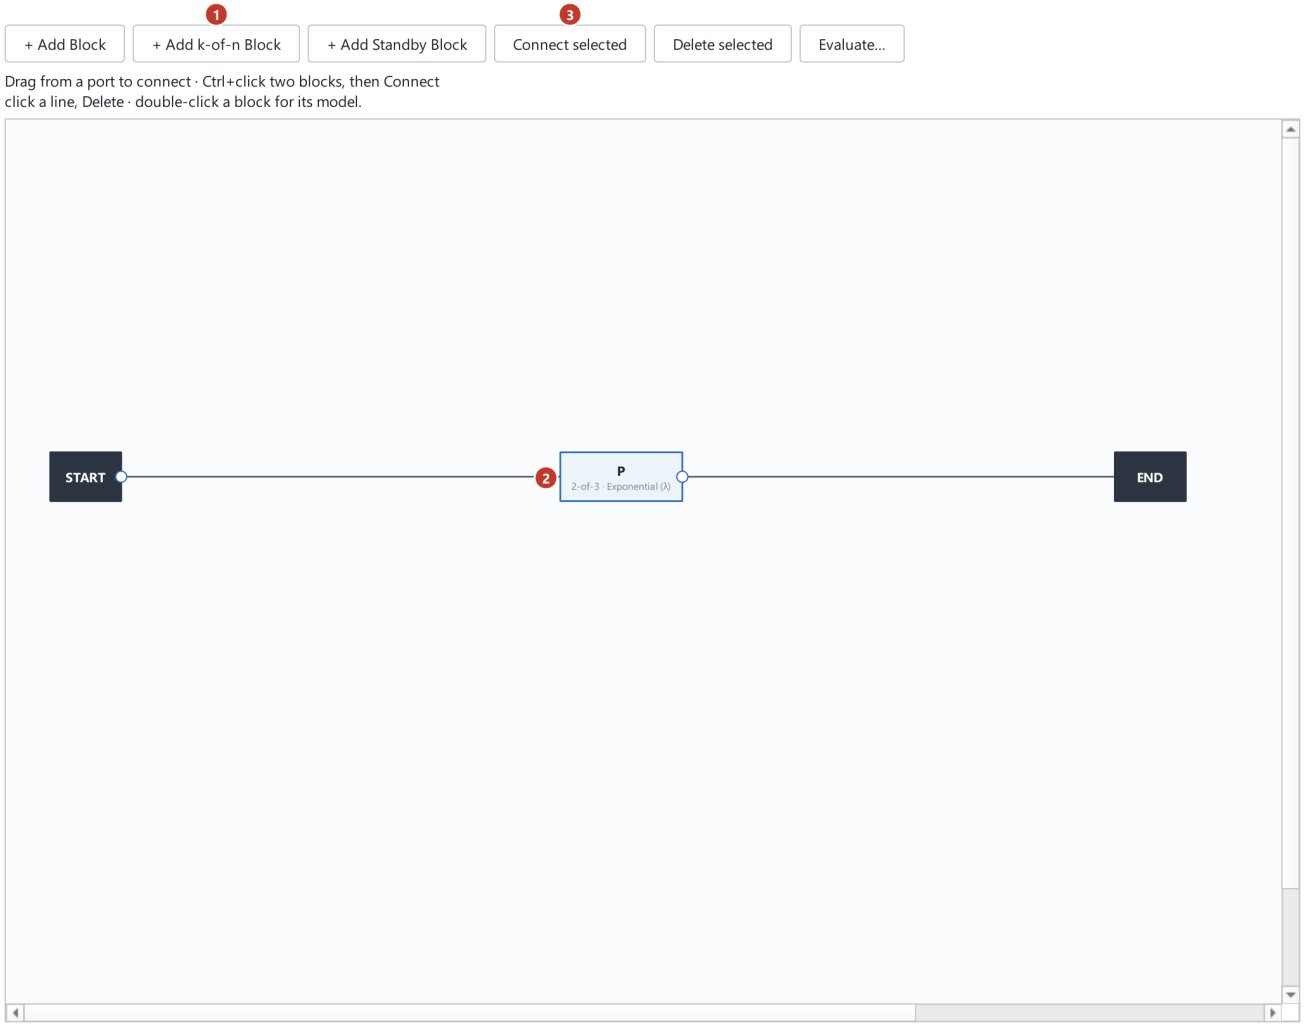


The example's values in the form

- ① and ② Simulate RBD on the canvas; row P's Repair model None (non-repairable).
- ③ to ⑤ Mission time T 30, Replications N 2000, Confidence % 95.

Part 3 — How to use it, step by step

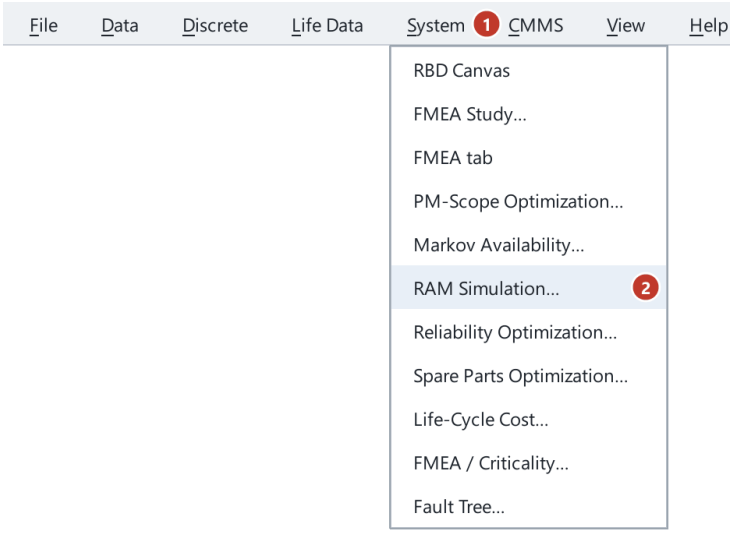
1. **Draw the station on the RBD canvas.** On the main tab RBD Canvas, click + Add k-of-n Block, name it P, give it 3 units with 2 needed and the life Exponential (λ) 0.02, and connect it between the start and the end with Connect selected.



Step 1 — the station on the RBD canvas

- ① + Add k-of-n Block: the button that adds the block.
- ② Block P, 2 of 3 units needed, connected between START and END.
- ③ Connect selected: joins the selected blocks.

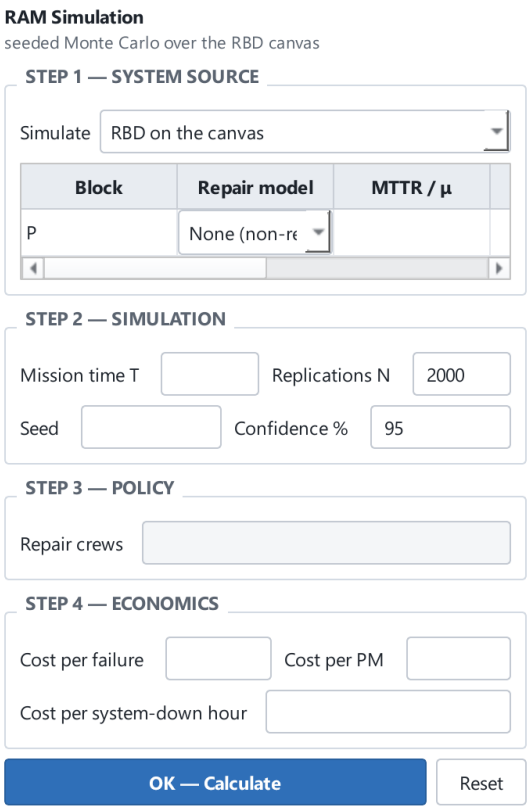
2. **Choose System ▸ RAM Simulation…** in the menu bar. The form RAM Simulation opens in the left panel with one policy row per block of the canvas.



Step 2 — the menu path

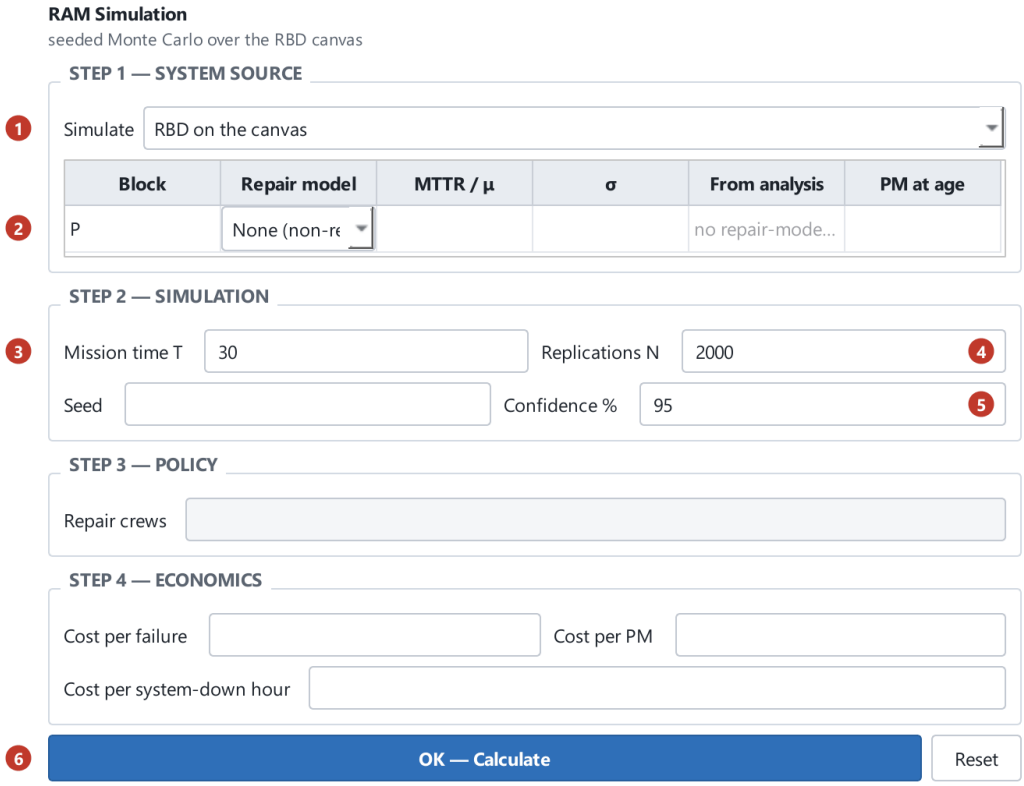
- ① System: the menu.
- ② RAM Simulation…: the entry to click.

The form as it opens, before any box is set:



Step 2 — the form as it opens

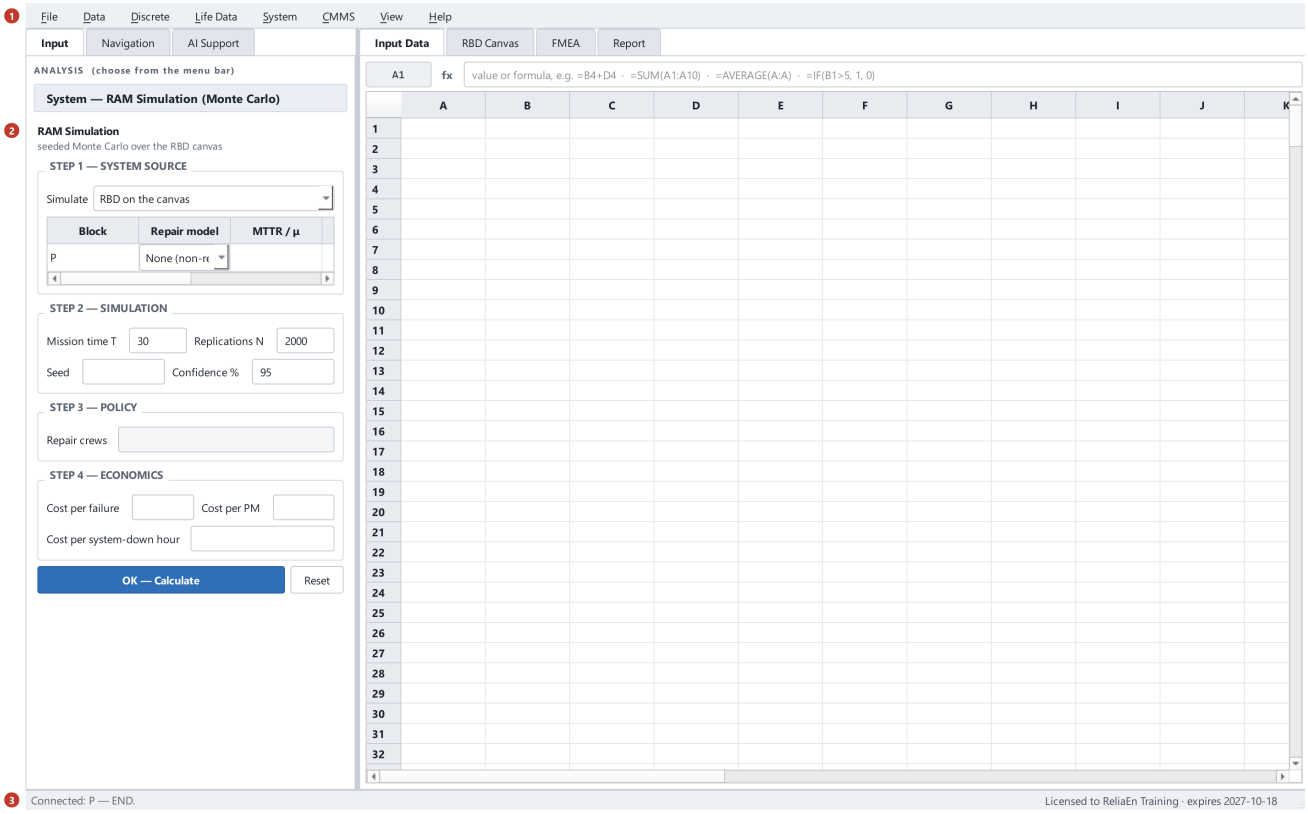
Steps 3 to 7 set the form's boxes, marked ① to ⑤ in the picture; ⑥ is the button of step 8:



Steps 3 to 7 — the form set for this example

3. ① **In Simulate, keep RBD on the canvas** - the diagram drawn on the canvas now, the only source in version 1.0.0.
4. ② **Check that row P's Repair model reads None (non-repairable).** PM at age stays blank: corrective only.
5. ③ **In Mission time T, type 30.**
6. ④ **In Replications N, keep 2000.** Seed stays blank: the house seed 20260825.
7. ⑤ **In Confidence %, keep 95.** Repair crews and the three costs stay blank.

8. **Click OK — Calculate** at the foot of the form (② in the picture below). ReliaEn simulates the 2,000 missions and shows the report.



Step 8 — the window just before OK

- ① The menu bar: File, Data, Discrete, Life Data, System, CMMS, View, Help.
- ② The form, set as in steps 3 to 7; OK — Calculate at its foot.
- ③ The status bar: Connected: P — END. (the last canvas action) and Licensed to ReliaEn Training · expires 2027-10-18.

Part 4 — The report, card by card

	INPUT DATA	
1	System	3 unit(s): P unit 1, P unit 2, P unit 3
2	Mission time T	30
3	Replications N	2000
4	Seed	20260825
5	Confidence	95%
6	Crews	unlimited
7	Policy	corrective only
8	Repair — P	None (non-repairable)

Input data

- ① System 3 unit(s): P unit 1, P unit 2, P unit 3: the units simulated, in the order drawn.
- ② Mission time T 30.
- ③ Replications N 2000.
- ④ Seed 20260825: the same seed gives the same report.
- ⑤ Confidence 95%.
- ⑥ Crews unlimited: every failed unit would be repaired at once.
- ⑦ Policy corrective only.
- ⑧ Repair — P None (non-repairable).

	RESULTS	
1	Availability — mean (SE)	0.823103 (SE 0.005918)
2	Availability — empirical 95% band (percentile)	[0.155349, 1.000000]
3	Mission reliability R̂(T)	0.5750 [0.5530, 0.5968] (Clopper-Pearson)
4	System failures — mean	0.4250
5	System downtime — mean (h)	5.3069
6	MTBF_sys (up-time ÷ system failures)	58.1014

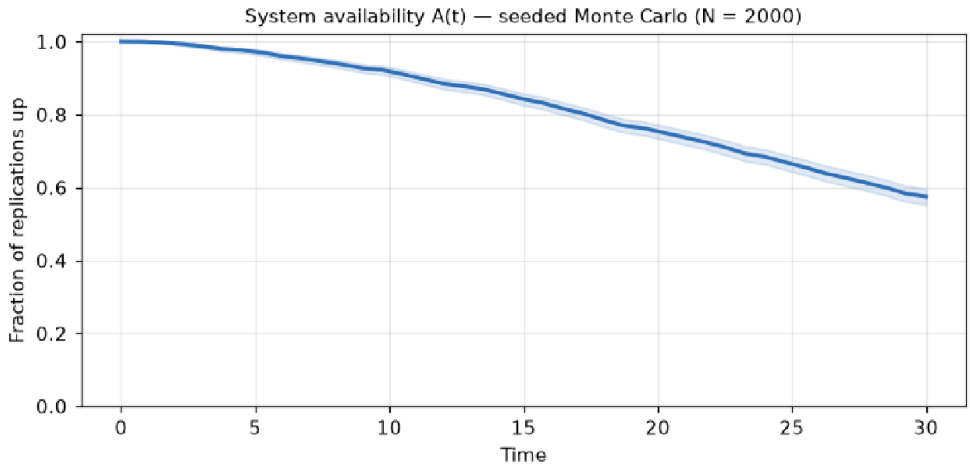
Results

- ① Availability — mean (SE) 0.823103 (SE 0.005918): the fraction of the mission the station is up, on average.
- ② Availability — empirical 95% band (percentile) [0.155349, 1.000000]: the middle 95% of the simulated missions.
- ③ Mission reliability R̂(T) 0.5750 [0.5530, 0.5968] (Clopper–Pearson): the fraction of missions with no system failure, with its bounds.
- ④ System failures — mean 0.4250 per mission.
- ⑤ System downtime — mean (h) 5.3069 hours per mission.
- ⑥ MTBF_sys (up-time ÷ system failures) 58.1014: the up-time per system failure.

	INTERPRETATION
1	Every number on this report reproduces from the seed above: keyed random streams per (replication, block) mean the same inputs always give the same digits. The identity card audits this run against the verified analytic engines wherever an analytic twin exists.

Interpretation

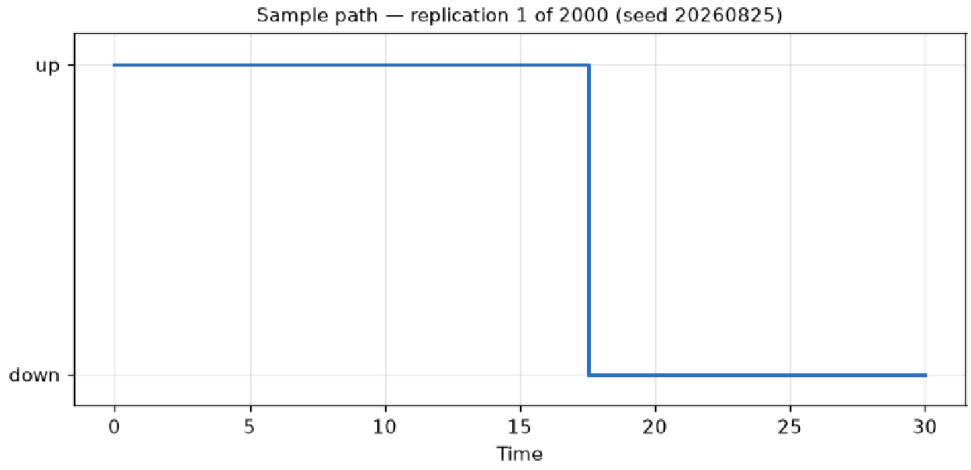
- ① Every number on this report reproduces from the seed above: keyed random streams per (replication, block) mean the same inputs always give the same digits. The identity card audits this run against the verified analytic engines wherever an analytic twin exists.



Interactive Plot

Availability against time

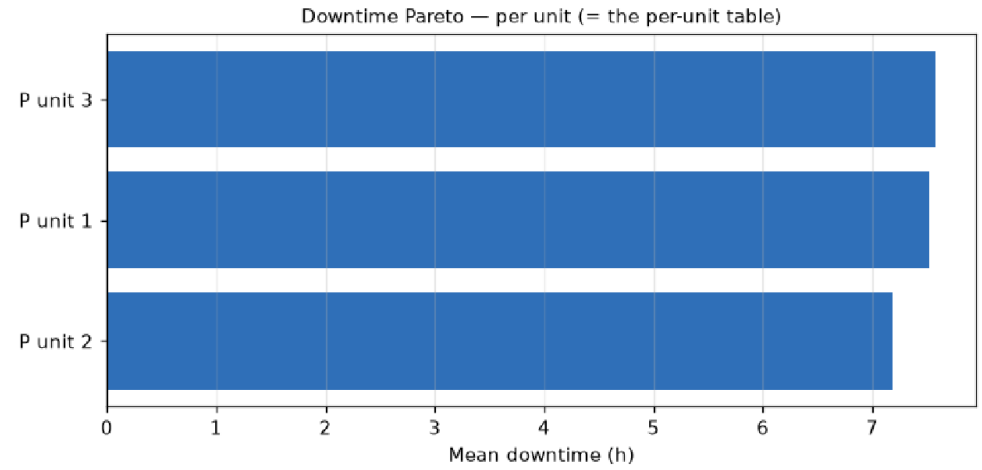
- The fraction of replications up at each time t from 0 to 30, with its Clopper–Pearson band; at t = 30 hours it reads 0.5750.



Interactive Plot

Sample path of replication 1

- The system up or down against time in the first replication; each step is a row of the sample-path card.



Interactive Plot

Downtime by unit

- One bar per unit, its mean downtime per mission - the rows of the per-unit card. The tallest: P unit 3, 7.5744 hours.

BUILT-IN IDENTITY CHECKS — THE SIMULATION AUDITED AGAINST THE VERIFIED ANALYTIC ENGINES

1 Each tolerance is 3× the identity's own standard error (derivations in ram_oracle.md); I5 and I6 are exact. Not applicable here: I2, I2a, I3a — closed forms cover a single exponential unit with exponential repair only; the 2-parallel presets are pytest-pinned (K-3); I3b — the renewal twin M(T) is pytest-pinned (tests/test_ram_sim.py), not an in-engine arm; I5 — economics off (no cost typed).

2 Id	3 Check	4 Simulated	5 Analytic	6 Tolerance	7 Verdict
8 I1	mission reliability vs analytic RBD	0.575000	0.572985	0.033182	agree

Built-in identity checks — the simulation audited against the verified analytic engines

- ① Each tolerance is 3× the identity's own standard error (derivations in ram_oracle.md); I5 and I6 are exact. Not applicable here: I2, I2a, I3a — closed forms cover a single exponential unit with exponential repair only; the 2-parallel presets are pytest-pinned (K-3); I3b — the renewal twin M(T) is pytest-pinned (tests/test_ram_sim.py), not an in-engine arm; I5 — economics off (no cost typed).
- ② to ⑦ The columns: Id, Check, Simulated, Analytic, Tolerance, Verdict.
- ⑧ I1 · mission reliability vs analytic RBD · 0.575000 · 0.572985 · 0.033182 · agree.

PER-UNIT RESULTS

1 The downtime Pareto's bars ARE these rows — same numbers, stated (no image-only state).

2 Unit	3 Failures (mean)	4 PM (mean)	5 Downtime (mean h)	6 System-down share
7 P unit 1	0.4590	0.0000	7.5156	0.3290
8 P unit 2	0.4370	0.0000	7.1818	0.3213
9 P unit 3	0.4600	0.0000	7.5744	0.3497

Per-unit results

- ① The downtime Pareto's bars ARE these rows — same numbers, stated (no image-only state).
- ② to ⑥ The columns: Unit, Failures (mean), PM (mean), Downtime (mean h), System-down share.
- ⑦ P unit 1 · 0.4590 · 0.0000 · 7.5156 · 0.3290.
- ⑧ P unit 2 · 0.4370 · 0.0000 · 7.1818 · 0.3213.
- ⑨ P unit 3 · 0.4600 · 0.0000 · 7.5744 · 0.3497.

A(T) GRID — THE AVAILABILITY PLOT'S OWN VALUES (ALL 41 POINTS)

① t	② Fraction up	③ lower	④ upper
0	1.0000	0.9982	1.0000
0.75	0.9995	0.9972	1.0000
1.5	0.9975	0.9942	0.9992
2.25	0.9930	0.9883	0.9962
3	0.9875	0.9816	0.9919
3.75	0.9800	0.9729	0.9857
4.5	0.9765	0.9689	0.9827
5.25	0.9700	0.9616	0.9770
6	0.9600	0.9505	0.9682
6.75	0.9535	0.9433	0.9623
7.5	0.9455	0.9346	0.9550
8.25	0.9370	0.9254	0.9473
9	0.9270	0.9147	0.9380
9.75	0.9220	0.9094	0.9334
10.5	0.9105	0.8971	0.9227
11.25	0.8975	0.8834	0.9105
12	0.8845	0.8697	0.8982
12.75	0.8775	0.8623	0.8916
13.5	0.8695	0.8539	0.8840
14.25	0.8560	0.8398	0.8711
15	0.8420	0.8253	0.8577
15.75	0.8310	0.8138	0.8472
16.5	0.8155	0.7978	0.8323
17.25	0.8020	0.7838	0.8193
18	0.7845	0.7658	0.8023
18.75	0.7690	0.7499	0.7873
19.5	0.7620	0.7427	0.7805
20.25	0.7495	0.7299	0.7684
21	0.7370	0.7171	0.7562
21.75	0.7245	0.7044	0.7440
22.5	0.7100	0.6896	0.7298
23.25	0.6930	0.6723	0.7132
24	0.6840	0.6631	0.7043
24.75	0.6690	0.6479	0.6896
25.5	0.6545	0.6332	0.6754
26.25	0.6385	0.6170	0.6596
27	0.6260	0.6044	0.6473
27.75	0.6130	0.5913	0.6344
28.5	0.5995	0.5776	0.6211
29.25	0.5830	0.5610	0.6047
30	0.5750	0.5530	0.5968

A(t) grid — the availability plot's own values (all 41 points)

- ① to ④ The columns: `t`, `Fraction up`, `lower`, `upper` - the availability plot as numbers, 41 points from 0 to 30 hours.

SAMPLE PATH — REPLICATION 1 OF 2000 (SEED 20260825)

① The sample-path plot draws exactly these transitions.

Transition time	System state
② 0	up
③ 17.5297	down

Sample path — replication 1 of 2000 (seed 20260825)

- ① `The sample-path plot draws exactly these transitions.`
- ② 0 · up: the station starts up.
- ③ 17.5297 · down: in replication 1 the second pump is lost at 17.5297 hours.

Part 5 — How to interpret the results

Can the run be trusted? Read the identity check first. With no repair and no PM the exact reliability of a 2-of-3 group exists: $3p^2 - 2p^3$ with p = e^(−0.02 × 30), that is 0.572985. The simulation gives 0.575000, within the tolerance of 0.033182, and the check reads agree. The run reproduces the exact answer where the two overlap - so the same diagram can be trusted when repairs, crews or PM are added.

How often the station keeps two pumps. In a fraction 0.5750 of the 30-hour runs the station never loses its second pump; with 95% confidence the true value lies between 0.5530 and 0.5968. The availability plot ends at the same 0.5750 at t = 30 hours.

What one run looks like. On average the system is up 0.823103 of the run and down 5.3069 hours, with 0.4250 system failures. The band [0.155349, 1.000000] is the real planning number: a bad run can leave the station short of pumps for most of the 30 hours.

Which pump to improve first. The one with the largest System-down share in the per-unit card: P unit 3 (0.3497). The three pumps are identical, so the shares are nearly equal; their differences are the simulation's own spread, not a weak pump.

What to do next. Add the repair model of the pumps (an exponential MTTR or a lognormal repair-time fit) and, if crews are limited, their number; with costs typed, the ECONOMICS card prices the mission. Raise Replications N before a decision: the bands and tolerances narrow as N grows. Keep the seed with the report.