

Buffer management: priorities and recurring causes

Use a declared buffer-priority policy to choose an eligible next action, expose blocked high-risk work and distinguish immediate recovery from recurring-cause improvement.



The red status needs an action, not another color

Fictional planning team: three jobs have locally defined protective-time statuses. J1 has high consumption but lacks a required drawing approval. J2 has increasing consumption and is ready; J3 has lower consumption and is ready. Coordinator Rhea must apply the local rule and investigate repeated approval delays.

Your role and the reference condition

Planning coordinator with the work-center and approval owners.

Normal: Protection status informs priorities and response without bypassing readiness or quality requirements.

Gap: The team considers starting blocked J1 because its color is strongest, or enlarging every buffer to make the red statuses disappear.

Work within these conditions

1. The case supplies qualitative zones only; no universal percentages or sizing algorithm are prescribed.
2. Local policy: work on the highest-priority eligible job and escalate blocked high-consumption work in parallel.

Fictional teaching case. Supplied observations support the exercise; proposed actions are not verified customer results.

See how the method works

A buffer only helps management when its consumption is visible and tied to a response. Define the buffer type, units and operating policy before using colored zones. Track which work is consuming protection, apply the agreed priority rule and record what caused the disruption. Repeated risk should trigger investigation and review of the system, not a permanent emergency mode. The diagram's zones are illustrative rather than prescribed thresholds. Production, project and replenishment buffers have different meanings, so their measures must not be mixed. Use the history to improve reliability and decide whether the protection or release policy needs revision.

Lower consumption

Monitor under agreed policy

Increasing consumption

Check emerging risk

High consumption

Act by defined priority rule

Agree local priority and escalation rules. No universal percentage thresholds.

Method exhibit: Measure protection, Prioritize response, Record disruption, Improve reliability.

Follow the evidence and decisions

Inspect the supplied case inputs

Job / rule	Supplied evidence
J1	High consumption; required drawing approval missing
J2	Increasing consumption; ready and eligible
J3	Lower consumption; ready and eligible
Recurring log	Drawing approval appears in 3 of the last 5 high-consumption episodes
Priority policy	Highest eligible status first; blocked high status escalated
Zone thresholds	Local definitions; numerical thresholds not supplied

01 / Verify the status and its clock

Rhea checks the data source, commitment boundary and agreed status calculation before acting on the display. She does not substitute elapsed job age for buffer consumption.

Why: A status is useful only if its clock and policy are understood. Different applications may use different protective mechanisms and measurements.

Evidence / check: Local status basis confirmed; no numerical consumption percentage is invented.

02 / Separate risk from readiness

J1 has the highest reported risk status but cannot legally enter the declared process without the required approval. Rhea records the missing decision and its owner.

Why: Priority does not grant technical or quality authority. A blocked urgent job requires a response, not a pretend start.

Evidence / check: J1 remains blocked with approval owner and next response time.

03 / Choose the eligible next work

Under the supplied policy, the work center takes J2 before J3 while the owner escalates J1's approval. Any later reprioritization follows the agreed safe interruption rule.

Why: The team can protect system progress while addressing the most exposed blocked job. A high-risk blockage need not imply idling all usable capacity.

Evidence / check: J2 selected; J1 escalated in parallel; J3 waits under normal policy.

04 / Separate recovery from learning

Record the cause and recovery for J1, then examine the five-episode log. Three drawing-approval episodes justify investigating that dependency, not declaring its root cause proven.

Why: Firefighting the current job and improving the recurring system condition are different tasks. A count prioritizes inquiry, not causal certainty.

Evidence / check: Approval dependency selected for investigation, with examples and timestamps retained.

05 / Review protection through its owner

After representative evidence, the owner reviews reliability, release assumptions and buffer adequacy. Do not enlarge protection solely to remove a color from the screen.

Why: A parameter change can conceal delays while increasing lead time or inventory. It needs a system-outcome prediction and follow-up.

Evidence / check: Policy/parameter review names evidence, decision and subsequent delivery check.

Completed priority and learning record

Item	Immediate decision	Follow-up
J1 high, blocked	Escalate approval; no invalid start	Record missing decision and response
J2 increasing, ready	Next eligible work	Monitor commitment protection
J3 lower, ready	Remain queued	Review when status changes
3 of 5 approval episodes	Investigate recurring dependency	Do not infer root cause yet
Buffer enlargement	Not automatic	Owner reviews reliability and assumptions

J1 becomes ready while J2 is underway

The approval arrives after J2 has started.

Decision

Reassess priority using the agreed interruption policy, setup costs and current process state; do not assume a mid-task stop is harmless.

Reasoning

Protection status informs coordination but does not replace safe execution or the declared dispatch rules.

What would change the decision?

The next dispatch decision records readiness change and the applicable interruption rule.

Practise with a changed case

Two high-consumption jobs with different readiness

Separate fictional queue: K1 is high consumption but missing accepted material, K2 is high consumption and ready, K3 is increasing consumption and ready. The same highest-eligible policy applies. Recent recurring causes are missing material, missing material, drawing change and missing material.

Job / log	Status
K1	High; material not accepted
K2	High; ready
K3	Increasing; ready
Recent causes	Material, material, drawing, material

Your task

1. Select the next eligible job and the parallel response.
2. Choose a recurring issue to investigate without claiming a proven cause.
3. State what cannot be calculated from the supplied qualitative zones.

Use the next page to record your reasoning before reading the answer.

Your practice worksheet

Two high-consumption jobs with different readiness

Status and readiness

Eligible dispatch

Blocked-job response owner

Recurring evidence

Missing numeric/policy information

Complete your reasoning before turning to the answer. Use the separate learner workbook for group practice.

Check your reasoning and transfer it

1. K2 is the highest-status eligible job and should be selected under the stated policy. K1 requires material acceptance response; its status does not authorize use of unaccepted material.
2. Three of four listed episodes involve material, making that dependency a reasonable investigation focus. Determine why readiness failed.
3. No buffer-consumption percentages, universal red thresholds or optimal buffer sizes can be calculated from these inputs.

Suggested completed record

Action	Evidence	Decision
Dispatch	K2	High and eligible
Escalate	K1	Material acceptance owner
Investigate	3 of 4 material episodes	Dependency, not proven root cause
Calculate	No numeric buffer data	Do not invent thresholds

Plausible wrong turns

1. High priority does not create eligibility.
2. Repeated red status does not automatically prove the buffer should be enlarged.

Evidence of a sound answer

1. Apply the supplied policy consistently.
2. Keep recovery and recurring-cause inquiry distinct.
3. Respect unavailable numerical information.

Take the learning back to work

Commitment	Agreed follow-through
Owner	Planning/system owner with readiness and execution counterparts
Record	Status/readiness history, dispatch exceptions and recurring-cause log
Review	During execution and at the agreed buffer-policy review
Verification evidence	Reliable commitments, timely blocked-work response and reduced recurring readiness failures
If unresolved	Review release, reliability and protection assumptions explicitly; do not hide the signal through unexamined parameter changes.

Help someone else apply the method

Prepare

1. Job cards with separate status and readiness fields
2. Recurring-cause log; answer kept separate

Minutes	Activity and coaching prompt
4	Explain the status Which clock and rule make the display meaningful?
8	Dispatch the original jobs How can the team help J1 while working J2?
8	Use the K-job variant What does three material episodes establish, and what not?
5	Review protection What outcome would justify a parameter change?

Debrief

1. Reject invented universal zone percentages.
2. Ask who owns the missing approval or material decision.

Individual or remote practice

Write a dispatch decision and a separate improvement question before opening the answer.

Connect the method and check its boundaries

Before application

1. Chosen DBR/CCPM policy
2. Reliable status data

Outside this lesson

1. Color is not a universal percentage policy
2. Expediting everything defeats prioritization

TOCICO: Buffer Management

TOCICO: Buffer Management

Buffer management and recurring red-zone instability; indexed conference text

Method reference; original OPEX scenario and diagram are synthetic teaching content, not source case results.

[Open this lesson and its visual aids](#)