ROGER BOWN
OPERATIONS · SUPPLY CHAIN · ERP

OPERATIONS IMPROVEMENT GUIDE

Demand Driven MRP

Protecting flow instead of chasing the forecast

The only thing to maximise is customer satisfaction. Everything else is a constraint. DDMRP places inventory where it protects flow — and replenishes on what customers actually consume.

THE FIVE CORE STEPS

- **01** Position the buffers
- **02** Size the buffers
- **03** Monitor daily
- **04** Replenish on consumption
- **05** Adjust over time

Practical leadership. Clear priorities. Lasting progress.

PROCESSES

Your planning system and business processes

“Your system is perfectly designed to give you the results you are getting.” — W. E. Deming

Planning system	Business planning	RRP
	Sales & operations planning	RRP
Operations	Master production scheduling	RCCP
	Material requirements planning	CRP
	Production activity control	I/O

RRP resource requirements planning · RCCP rough-cut capacity planning · CRP capacity requirements planning · I/O input–output control

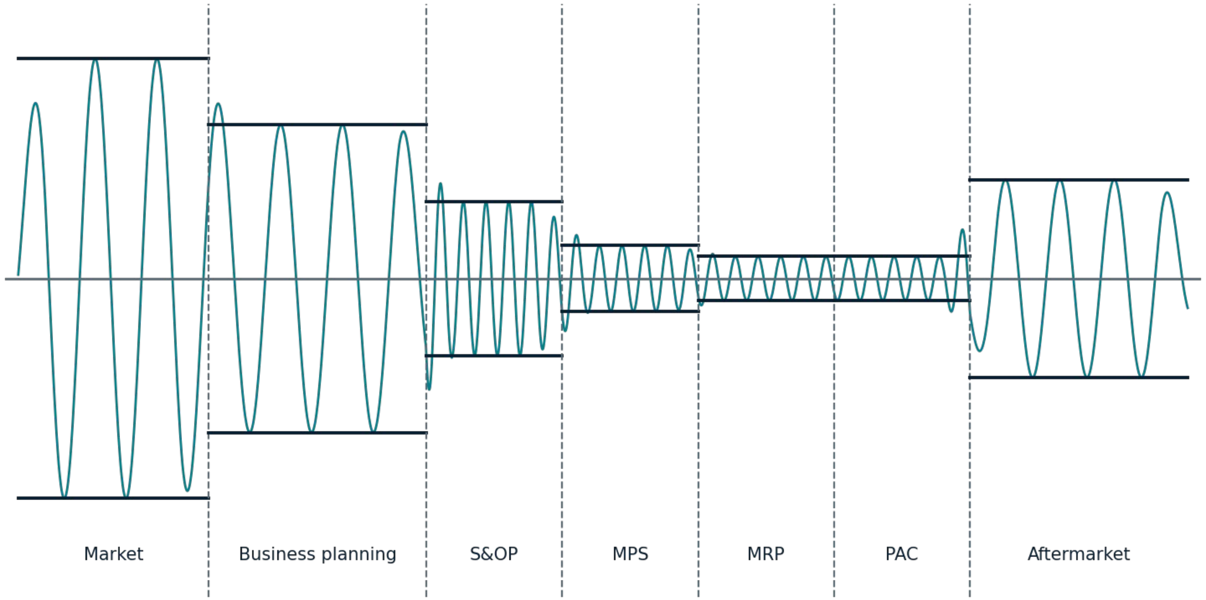

Each planning level narrows the variability the market throws at you. Where a level is missing or ignored, that variability reaches the shop floor and the customer.

S&OP – Sales & Operations planning · MPS - Master Production Scheduling · MRP - Materials requirements planning · PAC Production activity control

When the planning system fails, procurement, scheduling and fulfilment become a daily fire-fight: high inventory and poor service (out-of-stocks, availability, OTIF) — and the cost shows up in EBIT, OTIF and stock turns.

PROCESSES

The value drivers DDMRP moves

Forecast quality and inventory performance are directly linked — improve one and you move the other

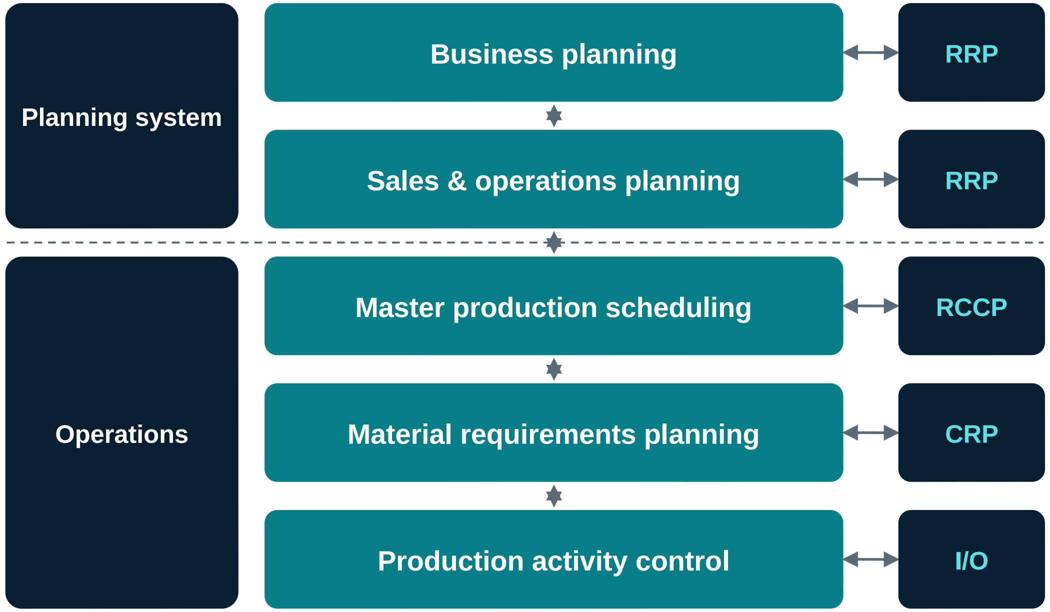

RRP resource requirements planning · RCCP rough-cut capacity planning · CRP capacity requirements planning · I/O input–output control

WHAT IT DOES FOR YOUR BUSINESS

Improve capital investment · Reduce inventories · Protect service

Key metrics

FORECAST QUALITY ← *direct relationship* → INVENTORY PERFORMANCE

Forecast quality

- Demand forecast error
- Mean absolute percentage error (MAPE)
- Forecast bias

Inventory performance

- Inventory turnover
- Inventory ageing
- Inventory days
- Holding costs
- Stock-outs
- OTIF

Better demand signals and better-placed stock raise turns and service at the same time — that is the DDMRP promise.

What Demand Driven MRP means for your business

PROCESSES

A planning method that protects flow instead of trying to forecast perfectly

DDMRP builds strategically placed inventory buffers that absorb variability and let you respond quickly to real demand. It grew out of Goldratt's Theory of Constraints, combined with Lean and MRP logic.

THE SIMPLE IDEA

Instead of *"predict exactly what customers will buy months in advance"* DDMRP says **"place smart buffers at key points and adjust them on actual demand."**

What companies typically see when it is implemented well

INVENTORY	SERVICE LEVEL	LEAD TIME	EXPEDITING
20 – 40%	**5 – 15%**	**30 – 60%**	**Dramatically**
lower	higher	shorter	less

Typical ranges from published DDMRP implementations — your baseline diagnostic sets the figures that apply to you.

PROCESSES

The five core steps

What your team actually does, day to day

01

Choose where to hold inventory

- Not everywhere — only at strategic decoupling points
- These break long supply chains into shorter, more manageable segments

02

Size the buffers

- Three zones per buffer:
- Red — safety
- Yellow — working stock
- Green — replenishment

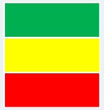

03

Monitor daily

- Instead of reviewing forecasts weekly or monthly, watch how much stock remains in each buffer and how fast it is being consumed

04

Replenish on consumption

- When inventory drops into yellow or red, replenishment is triggered
- Based on what was actually used — not on the forecast

05

Adjust over time

- Buffers are resized when demand patterns shift, lead times change or variability changes

Positioning and sizing are done once and reviewed; monitoring and replenishing are the daily routine.

Flow logic — traditional MRP versus DDMRP

Where the demand signal enters, and what it triggers

PROCESSES

Traditional MRP

Customer forecast
↓
Master production schedule
↓
MRP explosion (BOM + lead times)
↓
Planned orders
↓
Suppliers / production

Small forecast changes at the start cascade through the entire supply chain.

DDMRP

Customer orders
↓
Inventory buffer (decoupling point)
↓
Replenishment signal
↓
Planned orders
↓
Production / supplier

Buffers at key locations decouple the stages and respond to actual demand, preventing cascading disruption.

The difference is the trigger: a forecast that may be wrong, or consumption that has already happened.

How the two approaches compare

Key features and characteristics side by side

PROCESSES

Feature	Traditional MRP	DDMRP
Planning driver	Forecast	**Actual consumption**
System type	Push	**Pull**
Reaction to variability	Frequent rescheduling	**Buffers absorb it**
Inventory location	Spread across many stages	**Strategic buffers**
Stability	Low — schedule nervousness	**High — stable replenishment**
Lead times	Often long	**Shorter**
Planning complexity	High	**Visual buffer management**
What is critical	Forecast accuracy	**Flow protection**

Traditional MRP optimises the plan. DDMRP protects the flow.

PROCESSES

DDMRP in one picture

Five components in three groups — position, protect, pull

Model / re-model the environment			Plan	Execute
POSITION	PROTECT		PULL	
1	2	3	4	5
Strategic decoupling	**Buffer profiles and levels**	**Dynamic adjustments**	**Demand-driven planning**	**Visible, collaborative execution**
Where the buffers go	How big each zone is	How they change over time	When and how much to order	Buffer status everyone can see

What DDMRP typically delivers

- Lower overall inventory
- Higher service levels
- Shorter lead times
- Fewer expediting crises
- More stable production schedules

Green

Yellow

Red

The buffer everyone works to

Traditional MRP optimises plans.

DDMRP protects flow.

PROCESSES

In pictures — where the signal enters and where stock sits

Traditional MRP above, DDMRP below

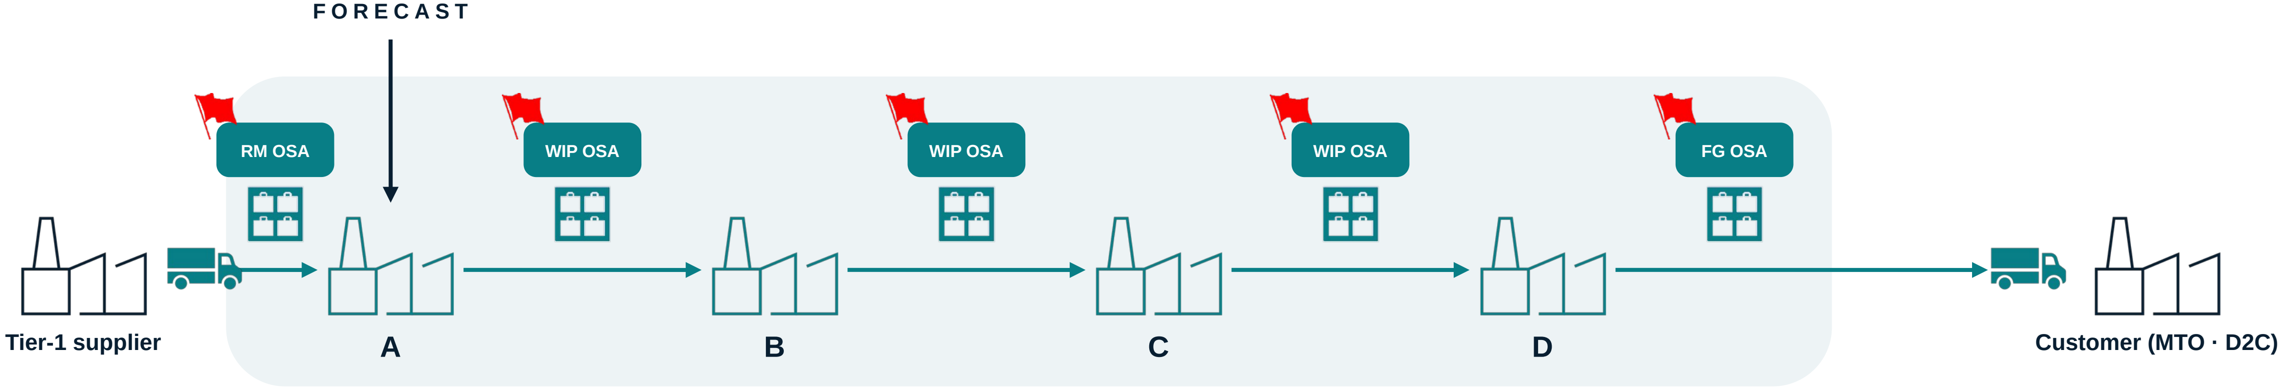

Traditional MRP — everything tightly connected: a small change at the start ripples everywhere.

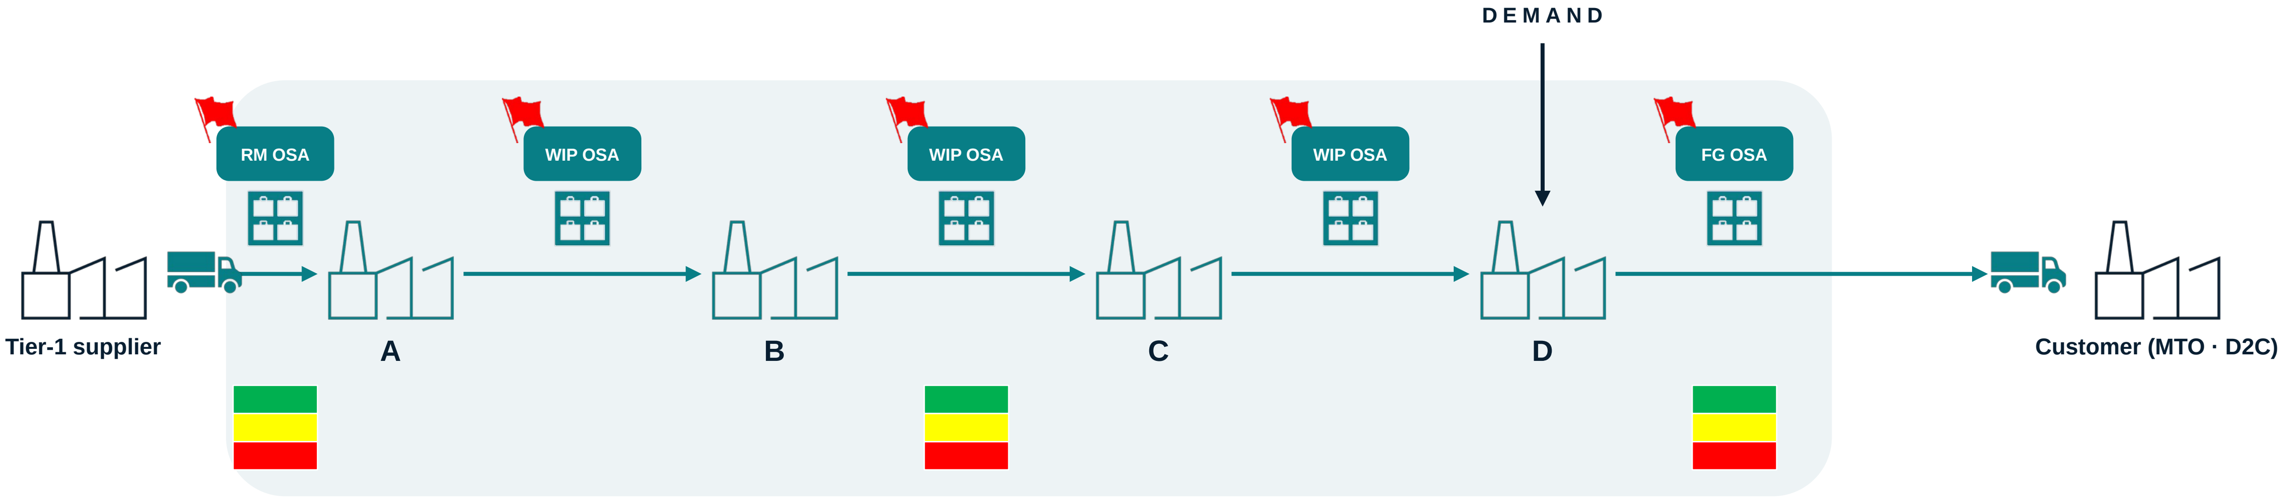

DDMRP — buffers at the raw-material, constraint and finished-goods points decouple the chain; replenishment follows real demand.

DDMRP manages where inventory should sit in the supply chain; Drum–Buffer–Rope controls how production flows through the factory constraint.

Control signals and control points

PROCESSES

CONTROL SIGNALS — demand pulls, right to left

Supplier orders ← Factory production ← Replenishment signal ← Buffer consumption ← Customer demand

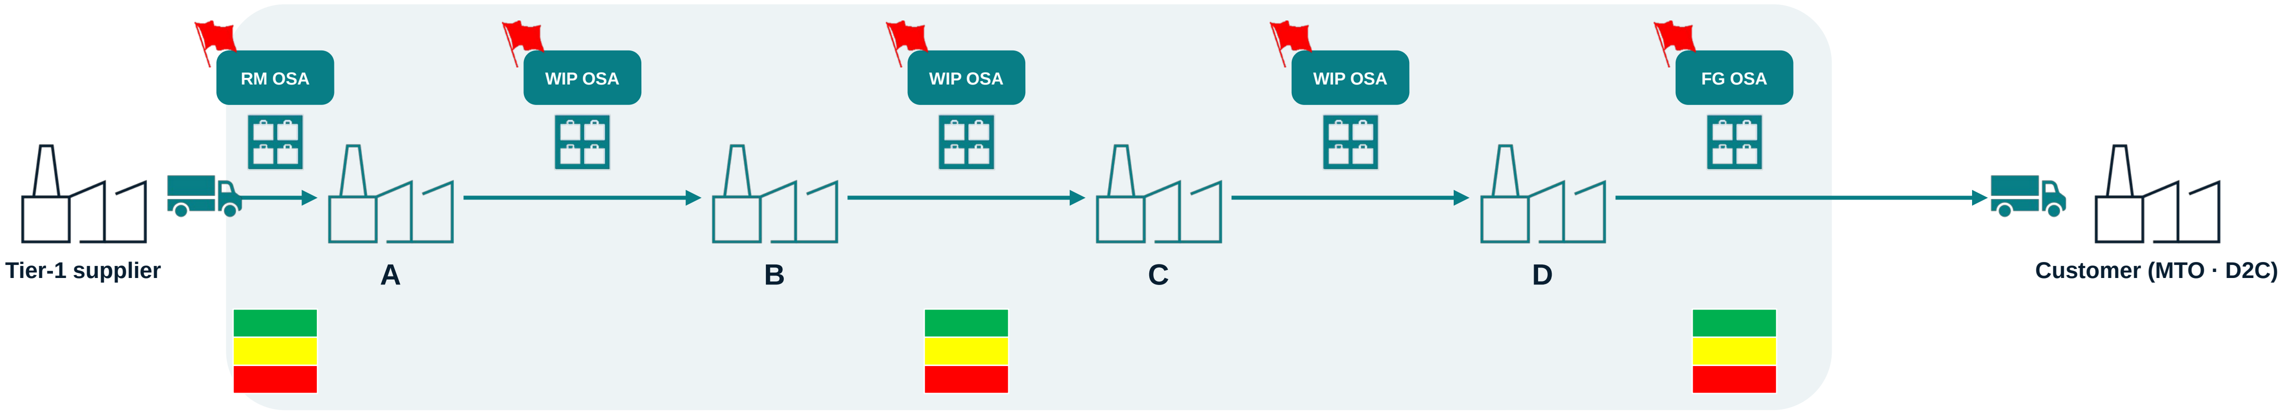

Supplier layer — raw-material buffers

Purpose Protect against supplier lead-time variability and stabilise inbound availability.

Placement Long-lead components · imported parts · high-variability materials

Factory layer — Drum, Buffer, Rope

Purpose Maximise throughput, never starve the constraint, limit excess WIP.

Placement Immediately before the system constraint

Distribution layer — finished-goods buffers

Purpose Absorb demand spikes and shorten the customer lead time.

Placement Regional DCs · high-demand markets · service-critical products

CONTROL POINTS — the three buffer layers

Inside the factory — Drum, Buffer, Rope

How production is scheduled around the constraint

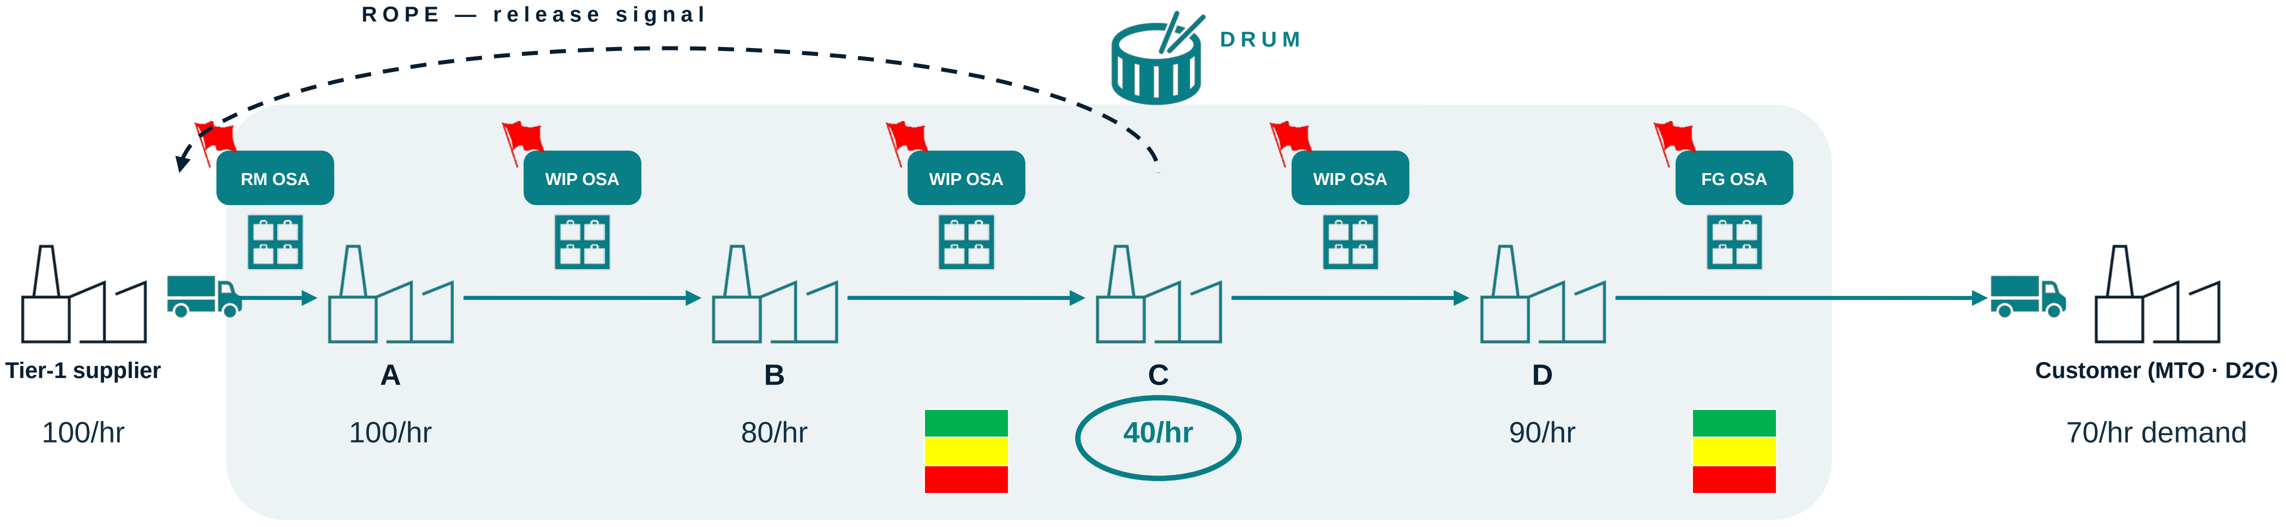

ROPE — release control

- A signal from the drum back to the release point
- Controls how much material enters the system
- If the drum runs at 40/hr, release at 40/hr

BUFFER — protection

- A time buffer, not just inventory, placed before the constraint
- Protects the drum from upstream variability, downtime and supplier delays

DRUM — the pace-setter

- The system constraint sets the beat and never waits
- All upstream activity is subordinated to it
- Every lost minute is lost revenue

Inbound & production scheduling | **Production execution & quality** | **Warehousing & outbound**

Supply planning — material requirements planning and inventory control

In pictures — the buffer and how it moves

Three zones, three planning levels, and buffers that flex with demand

PROCESSES

Anatomy of a buffer

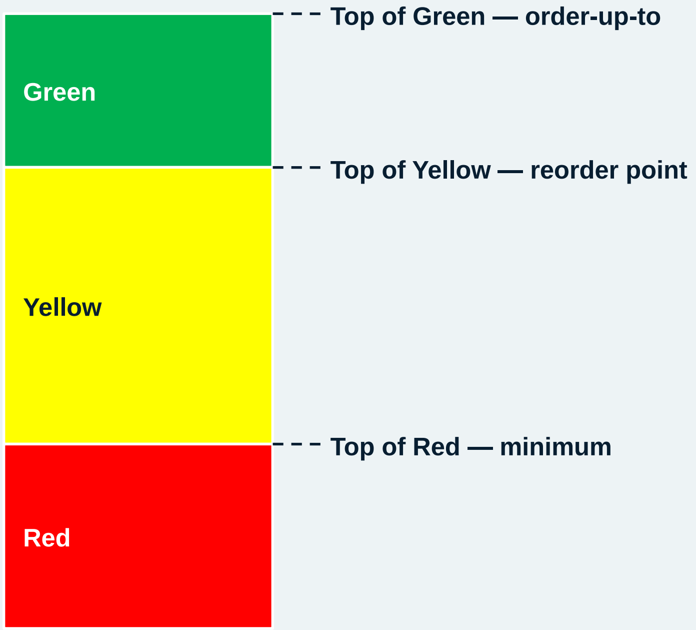

Green: replenishment signal · Yellow: working stock · Red: safety protection

Buffers adjust as demand changes

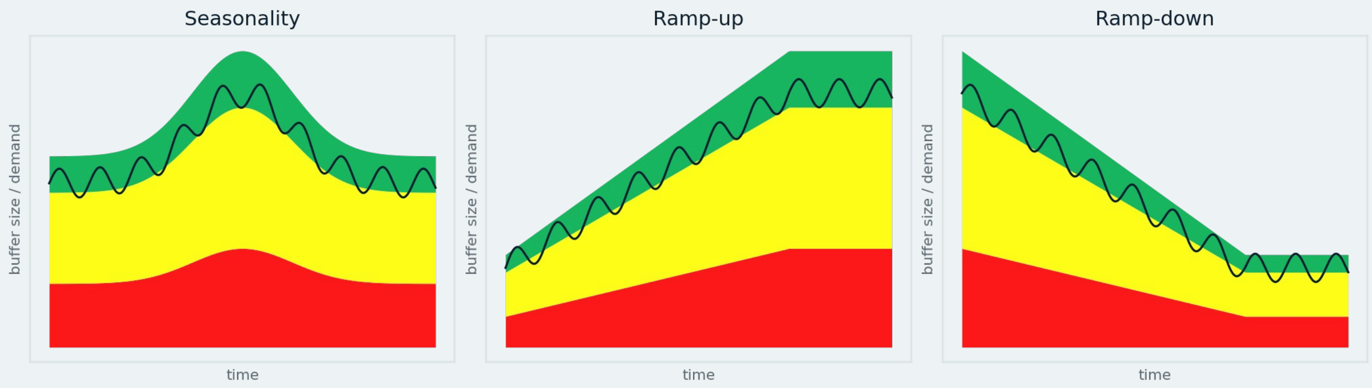

Zone sizes are recalculated from average daily usage, so the buffer grows into a season or ramp-up and shrinks out of it — without anyone re-forecasting.

Everyone reads the same three colours: green means leave it, yellow means an order is due, red means act now.

The systems that support the architecture

Manufacturers typically implement the model through their ERP and planning systems

PROCESSES

Layer	Typical system capability
Demand sensing	Advanced planning / demand-sensing tools
DDMRP buffers	Inventory optimisation modules
Factory execution	MES or production scheduling
Constraint scheduling	DBR-style finite scheduling
Supplier collaboration	Supplier portals / EDI

THE COMBINED ARCHITECTURE GIVES YOU

- A stable supply chain
- Protected constraints
- Predictable throughput
- Responsive replenishment

Most ERPs already hold the data DDMRP needs — usage, lead times, on-hand and on-order. The change is in the logic, not the system.

Sizing the buffers — the inputs you need

Per item, at each buffer position

ADU
Average daily usage

Units per day, from actual consumption history

DLT
Decoupled lead time

Days to replenish this buffer — not the end-to-end supply chain lead time

LTF
Lead-time factor

Dimensionless, typically 0.2–0.8; higher for long or unstable lead times

VF
Variability factor

Dimensionless, typically 0.2–1.0; higher for volatile demand or supply

MOQ / pack
Order constraints

Optional — minimum order quantity and pack size

OC
Order cycle

Optional but recommended — how often you review and replenish

Vendors and implementations differ slightly in formulas and tuning factors. The common logic: build zones from usage and lead time, then adjust for variability.

Sizing the buffers — base level and red zone

Step 1: the baseline. Step 2: the protection

PROCESSES

STEP 1 — BASE BUFFER LEVEL

A simple baseline: how much you will consume during the replenishment lead time.

$$\text{Base} = \text{ADU} \times \text{DLT}$$

Everything else is built from this number — get ADU and DLT right first.

STEP 2 — RED ZONE (PROTECTION)

Often modelled as two parts:

1. Safety from lead time

$$\text{Red(LT)} = \text{Base} \times \text{LTF}$$

2. Safety from variability

$$\text{Red(VAR)} = \text{Base} \times \text{VF}$$

$$\text{Red} = \text{Red(LT)} + \text{Red(VAR)}$$

Red grows when lead time is long or uncertain, or when demand or supply is volatile.

Red is the part of the buffer you never plan to use — it is what keeps the customer served when something goes wrong.

Sizing the buffers — yellow and green zones

Working stock and the replenishment signal

PROCESSES

YELLOW ZONE (WORKING STOCK)

The cover for the replenishment lead time.

$$\text{Yellow} = \text{Base} = \text{ADU} \times \text{DLT}$$

Yellow is the core operating inventory that bridges one replenishment to the next.

GREEN ZONE (REORDER / ORDER-UP-TO)

Reflects order cycle, MOQs and the replenishment batch size you want. A simple formula:

$$\text{Green} = \text{Yellow} \times \text{GTF}$$

GTF (green time factor) is often 0.2–1.0 depending on ordering economics.

If you use an explicit order cycle (OC) in days:

$$\text{Green} = \text{ADU} \times \text{OC}$$

then enforce Green $\geq$ MOQ (or $\geq$ pack size)

Green is the order quantity: when stock falls into yellow you order enough to get back to the top of green.

From zones to decisions

Step 3: planning levels. Step 4: the replenishment quantity

PROCESSES

STEP 3 — PLANNING LEVELS

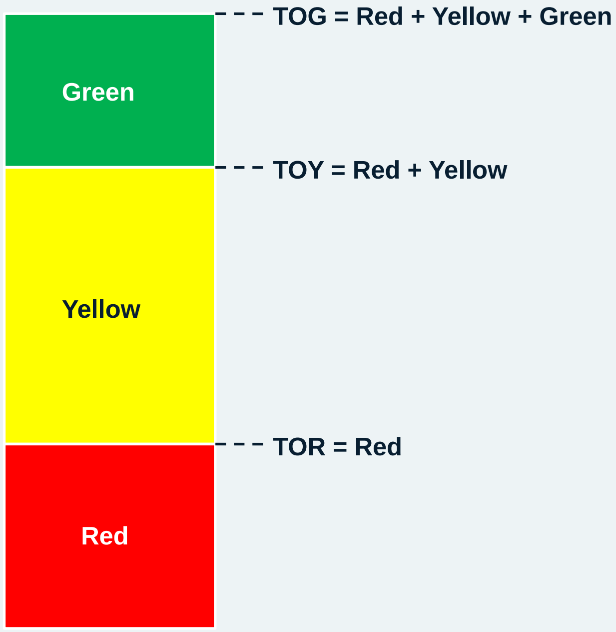

STEP 4 — REPLENISHMENT QUANTITY (SIMPLE RULE)

First compute the net flow position — what you have, what is coming, and what is already spoken for:

$$\text{NFP} = \text{On-hand} + \text{On-order} - \text{Qualified demand}$$

Then order up to the top of green:

$$Q = \text{TOG} - \text{NFP}$$

Round or adjust to MOQ, pack size and capacity constraints. The three thresholds — TOR, TOY, TOG — are what buffer status and every replenishment decision are read against.

Where NFP sits against the three levels is the whole daily conversation: green, yellow or red.

Worked example

One item, one buffer position — from inputs to order-up-to level

PROCESSES

ASSUME

ADU	100 units/day
DLT	5 days
LTF	0.30
VF	0.40
OC	2 days
MOQ	150

CALCULATION

1. Base (yellow)	100 × 5 = 500
2. Red	Red(LT) = 500 × 0.30 = 150 Red(VAR) = 500 × 0.40 = 200 Red = 350
3. Green (using OC)	100 × 2 = 200 · MOQ check: 200 ≥ 150 ✓
Planning levels	TOR = 350 TOY = 350 + 500 = 850 TOG = 350 + 500 + 200 = 1,050

RESULT

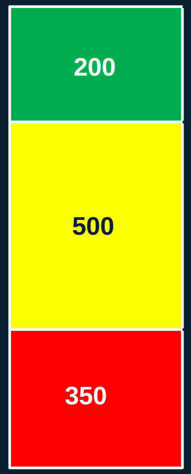

Order up to 1,050 units, and read buffer status by where NFP sits against 350 / 850 / 1,050.

The same arithmetic, run for every buffered item each day, is the whole planning engine.

Choosing LTF, VF and OC — practical starting points

Heuristics that get you running; tune from service level and buffer-penetration history

PROCESSES

LTF — lead-time factor

increases with

- Long or unstable replenishment lead time
- Low supplier reliability

VF — variability factor

increases with

- High demand volatility
- Yield issues or quality fallout
- Transport variability

OC — order cycle

comes from

- How often you want to replenish (daily, twice weekly, weekly)
- Capacity constraints and ordering costs
- MOQs and pack sizes

A SANE STARTING POINT MANY TEAMS USE

LTF **0.25 – 0.40** VF **0.25 – 0.60** OC **1 – 5 days**

depends on cadence and MOQs — then tune on service level and buffer-penetration history

Start conservative, measure penetration into red, and tighten the factors as the data comes in.

The bullwhip effect — what we are trying to reduce

A standard measure is variance amplification

PROCESSES

BULLWHIP RATIO

$$BW \equiv \frac{\text{Var (Orders)}}{\text{Var (Demand)}}$$

Greater than 1 means your ordering rule is amplifying demand noise upstream.

What it means for you

Bullwhip happens when the ordering rule turns small demand fluctuations into larger upstream order fluctuations. Your suppliers see a far more volatile signal than your customers ever generated — and you pay for it in stock, expediting and capacity.

In classical order-up-to (OUT) analysis — the logic behind most MRP-style replenishment — bullwhip is driven mainly by:

- Forecast updating — moving averages and exponential smoothing that chase the latest data
- Lead time — physical and information delay between demand and supply

DDMRP and DBR each attack one of these drivers — the next slides show how.

Why traditional MRP amplifies variance

The intuition behind the maths, without the derivation

PROCESSES

A SIMPLIFIED ORDER-UP-TO RULE

$$O(t) = \hat{D}(t, L) + (S - IP(t))$$

$\hat{D}(t, L)$ forecast of demand over the lead time L
S order-up-to target
IP(t) inventory position

Every order is a fresh forecast plus a correction back to target — so every forecast wobble becomes an order wobble.

A — Forecast nervousness

$$F(t) = F(t-1) + \alpha \cdot (D(t-1) - F(t-1))$$

A higher α puts more weight on the latest data: forecasts react faster and orders become more volatile.

B — Lead time multiplies forecast error

Because you forecast demand over L periods, the variance of the lead-time forecast error grows with lead time (and with how demand is correlated). That raises Var(O) — and therefore BW.

Bottom line: longer lead times plus frequent forecast updates mean larger order-variance amplification.

How DDMRP reduces bullwhip

It changes the structure of the ordering equation

PROCESSES

DDMRP introduces decoupling buffers and replaces the forecast term with actual consumption. Two effects follow.

A — It shortens the effective lead time

$$L \rightarrow DLT$$

Instead of planning with end-to-end lead time, DDMRP plans with the decoupled lead time — the replenishment time to one buffer position. The L that sits inside lead-time forecasting and error propagation gets smaller, so variance amplification falls.

B — It replaces the forecast with replenish-to-buffer

$$Q(t) = TOG - NFP(t)$$

NFP(t) = on-hand + on-order − qualified demand. This behaves like a base-stock pull rule driven by consumption, removing the explicit forecast-update term that drives OUT bullwhip. Forecasts may still size the buffers strategically, but daily execution is consumption-driven.

Shorter L and no daily forecast term: the two biggest bullwhip drivers are both reduced.

Buffers as a filter — and DBR inside the factory

Two further mechanisms that dampen variability

PROCESSES

C — Buffers create a deadband (a low-pass filter)

- Replenishment is triggered by buffer penetration, so small demand noise is absorbed inside the zones and does not cause proportional order swings
- Mathematically this adds state buffering and reduces the sensitivity $\partial O/\partial D$ at high frequencies — short-term noise — lowering Var(O)

D — Drum, Buffer, Rope attacks internal variability

Rope — a rate limiter on release

$$\text{Release}(t) \approx \min(\text{available WIP}, \text{drum rate})$$

Even if upstream tries to surge, the rope prevents excess WIP and chaotic starts and stops — internal order variability that would otherwise propagate back to suppliers.

Buffer — protects the constraint

A time/WIP buffer before the drum reduces starvation, so completions are smoother — less expediting and less volatile replenishment.

DDMRP calms the signal between sites; DBR calms the flow inside the plant.

The whole argument in one inequality

Why DDMRP with DBR produces less bullwhip than MRP with order-up-to

PROCESSES

$$\text{BW (DDMRP + DBR)} \quad < \quad \text{BW (MRP / OUT)}$$

If bullwhip in an MRP / order-up-to environment increases with lead time L and forecast responsiveness α, then moving to DDMRP with DBR lowers it because it:

01 Reduces L to DLT

Decoupling buffers shorten the lead time that drives error propagation

02 Lowers the effective α

Daily execution runs on consumption, not high-frequency forecast updates

03 Caps release variability

The rope and drum regularise what the factory starts and when

Less amplification means less stock, less expediting and steadier suppliers — the value drivers from the start of this guide.

THE FIRST MANAGEMENT DECISION

Where should your decoupling points be — and what is the inventory and service at stake?

A short diagnostic answers that with evidence: one product family followed from supplier to customer, the candidate buffer positions identified, and the stock, lead-time and service effect of buffering them quantified.

WHAT HAPPENS NEXT

01 Map the flow and candidate decoupling points

02 Size buffers for a pilot family

03 Run replenish-on-consumption for one cycle

04 Measure turns, OTIF and expediting; then scale