

Pumping station PS-04 — Operating principles

Trainee document · crude transfer to terminal

Contents

How to read this document	2
1. The feed vessel T-100 and the mass balance	2
1.1 Why level is the first thing to look at	2
1.2 The mass balance: what comes in minus what goes out	2
1.3 The three level settings, and why they sit where they do	2
1.4 Settings of this chapter	3
2. The pumps P-101 A/B: curve, duty and standby, NPSH	3
2.1 Why two pumps	3
2.2 The operating point: two curves that meet	3
2.3 Motor current	3
2.4 Pump changeover while running: three phases, not a switch	4
2.5 NPSH: why level and flow decide together	4
2.6 Settings of this chapter	4
3. Pressure control PIC-205	5
3.1 Why control pressure and not flow	5
3.2 The measurement chain, and why the lag matters	5
3.3 The three terms, and what each one does	5
3.4 Anti-windup, and why limiting the output is not enough	5
3.5 Automatic, manual, and bumpless transfer	5
3.6 Parameters of this chapter	6
4. The safety instrumented system: 2oo3 voting and SDV-301	6
4.1 Why voting and not a single transmitter	6
4.2 The drift fault, and how to read it	6
4.3 What the trip does	6
4.4 First-out: which cause came first	6
4.5 ESD and reset	7
4.6 Settings of this chapter	7
5. ISA-18.2 alarms and settings editable on line	7
5.1 Why alarm management, and not a list	7
5.2 The three states	7
5.3 The three priorities	7
5.4 Settings change on line, but under a rule	7
5.5 Settings of this chapter	8
6. Machine protections: vibration and temperature	8
6.1 Why vibration is measured	8
6.2 The zones and the settings	8
6.3 Bearing temperature	8
6.4 The injectable faults and what to watch	8
6.5 Settings of this chapter	9

7. The start and stop sequence	9
8. What the model simplifies	9

How to read this document

The PS-04 simulator reproduces a transfer station: a feed vessel, two centrifugal pumps in parallel with one in service and one on standby, a control valve, a safety shutdown valve and the line to the terminal.

Every number the panel displays comes out of a model solved ten times a second. There are no fake animations: close the valve and the pressure really rises; drain the vessel and the pump really loses suction. This document explains **why** each system exists and **how** it reacts, so that what you see on the panel can be predicted rather than endured.

Instrument tags (PIT-205, LSSL-102, SDV-301...) are never translated in any language: they are identifiers, not words. Each chapter ends with its settings in a table; the last section of the document states what the model simplifies, because knowing where the simulator stops is part of the training.

1. The feed vessel T-100 and the mass balance

1.1 Why level is the first thing to look at

A centrifugal pump does not suck: it receives. It is the head of liquid above the suction nozzle that pushes the crude into the impeller. If the level falls, so does the pressure available at suction, and below a certain point the liquid starts to vaporise right at the impeller inlet — the cavitation of chapter 2. The level of T-100 is therefore not an inventory figure: it is the condition that makes pumping possible at all.

That is why level appears three times on the panel — on the mimic, in the process values and in the settings table — and it is the only parameter that appears both as a start permissive and as a trip cause.

1.2 The mass balance: what comes in minus what goes out

The level is not set anywhere: it is the result of a balance.

$$dL/dt = (\text{feed} - \text{transferred flow}) / 3600 / \text{area}$$

With an area of 2.0 m², a difference of 100 m³/h between feed and transfer moves the level by about 0.014 m/s, that is almost a metre per minute. It is a rate the trainee needs to have in mind: an imbalance that looks small on the board empties the vessel in a few minutes.

The panel writes the instantaneous balance under the feed command — *even, filling, draining* — and when the vessel is draining it adds the time remaining to the very low level trip. That is the information that lets you decide whether there is time to correct or whether it is better to cut the transfer at once.

1.3 The three level settings, and why they sit where they do

LAH-102 (high, 10.5 m) warns that the vessel is filling beyond normal: at 11.9 m the model shuts off the field feed to avoid overflow.

LAL-102 (low, 4.00 m) is both an alarm and a **permissive**: below this setting starting is inhibited. It does not stop a pump already running — that would be a trip, not a permissive — but it prevents starting one in conditions that would lead straight to cavitation.

LSSL-102 (very low, 3.20 m) trips the pumps. Its position is not arbitrary: at nominal flow the cavitation point (where available NPSH equals required NPSH) falls around 2.9 m of level. The trip setting sits **above** that point, not below. A level setting placed below the cavitation point would be a useless protection: the machine would be damaged by vibration before the level ever got low enough to stop it — which is exactly the situation one of the training scenarios presents to the trainee.

1.4 Settings of this chapter

Tag	Description	Starting value	Action
LAH-102	High level	10.5 m	Alarm; at 11.9 m feed shut off
LAL-102	Low level	4.00 m	Alarm + start inhibited
LSLL-102	Very low level	3.20 m	Pumps tripped

All three are editable on line from the settings panel (chapter 5).

2. The pumps P-101 A/B: curve, duty and standby, NPSH

2.1 Why two pumps

The transfer cannot stop for maintenance nor for a failure. The two machines are identical and interchangeable: one is **in service**, the other **on standby**. The selection on the panel is not a switch that changes over instantly but the start of a sequence: chapter 2.4 explains why.

2.2 The operating point: two curves that meet

A centrifugal pump has no flow rate of its own. It has a **curve**: the more flow passes, the less head it can deliver.

pump head: $\Delta P = 48 \cdot n^2 - 1.8 \cdot 10^{-4} \cdot Q^2$ [bar, m³/h]

where n is the relative speed commanded by the drive (0 to 1). The first term says head grows with the square of speed; the second that it falls with the square of flow.

On the other side there is the system, which offers a resistance:

system curve: $Q = 46 \cdot kv \cdot \sqrt{(Pd - 8.0)}$

with kv the opening of control valve PCV-201 and 8.0 bar the pressure of the line to the terminal. The **operating point** is where the two curves meet, and the simulator finds it at every scan by bisection on the header pressure.

Two facts follow, visible at once on the panel, that surprise anyone used to thinking of a pump as a flow generator:

- **closing PCV-201 makes pressure rise and flow fall**, not the other way round;
- **increasing speed raises both**, but not in proportion.

If the valve closes completely the pump goes into *deadhead*: no flow, all the head available, pressure at full scale. It is the manoeuvre used to prove the overpressure chain of chapter 4.

2.3 Motor current

Motor current follows speed to the power 2.7 and grows with flow:

$I = 145 \cdot (0.12 + 0.88 \cdot n^{2.7}) \cdot (0.72 + 0.34 \cdot Q/280)$ [A]

It is used to recognise an abnormal condition before an alarm arrives: a pump running fast while drawing little is delivering little flow, and the reason is to be looked for downstream (valve closed) or upstream (suction in trouble).

2.4 Pump changeover while running: three phases, not a switch

Changing the duty pump **with the machine stopped** is immediate: it only changes which of the two will start. **While running** it is not, and it is the most instructive manoeuvre on the panel.

The problem is the non-return valve on each pump's discharge. The incoming machine, as long as its head is below the pressure already present in the header, cannot open it: it spins without delivering anything. If the outgoing pump were stopped before the incoming one had taken the load, the transfer would be interrupted and the pressure would collapse.

The implemented sequence avoids this gap in three phases:

1. **PRE-LUBE** (3 s) — the incoming machine prepares, the outgoing one keeps sustaining the transfer on its own.
2. **STARTING** — the incoming machine ramps up in speed until it **matches the header pressure**. The value is not fixed but computed: $n = \sqrt{((P_d - \text{suction pressure}) / 48)}$, with a small margin above, and in any case not less than 35 %. Only then does its non-return valve open.
3. **CHANGEOVER** — the load passes to the incoming machine, the controller realigns bumplessly (chapter 3.5) and the outgoing one ramps down to a stop.

The result, measured on the simulator: throughout the changeover the discharge pressure stays between 30.8 and 32.9 bar against a set-point of 32.0, and the flow never has a gap. An earlier version that ramped the incoming pump to a fixed 35 % instead produced a 4.5 bar excursion and a flow collapse to 127 m³/h: the non-return valve stayed shut because the head was not enough to open it.

If a trip arrives during the changeover, or a stop is commanded, the sequence is aborted.

2.5 NPSH: why level and flow decide together

Required NPSH is how much head the pump demands in order not to cavitate, and it **grows with the square of the flow**:

$$\text{NPSHr} = 2.0 + 5.5 \cdot (Q/300)^2 \quad [\text{m}]$$

Available NPSH is what the plant offers, and depends on the vessel level:

$$\text{NPSHa} = \text{level} + 1.6 - 0.9 \quad [\text{m}]$$

When NPSHa falls below NPSHr the liquid vaporises at the impeller inlet: the bubbles collapse shortly after, and the panel shows it as a step in vibration proportional to the deficit. It is not an alarm that trips: it is a phenomenon you watch grow.

The operational consequence is less obvious than it seems: **cavitation can be fought by reducing flow**, not only by raising the level. If the level cannot be raised — because the field feed is what it is — the only way is to ask for less flow, and that is what one of the training scenarios requires the trainee to understand.

2.6 Settings of this chapter

Tag	Description	Starting value	Action
IAH-220	High motor current	155 A	Alarm only
YA-101	Insufficient NPSH, cavitation	$\text{NPSHa} < \text{NPSHr}$	High priority alarm
YA-104	Mechanical seal, leakage	fault injected	High priority alarm

3. Pressure control PIC-205

3.1 Why control pressure and not flow

The line to the terminal accepts whatever is sent to it, but demands a minimum pressure to reach its destination and does not tolerate an excessive one. The transfer is therefore controlled on **discharge pressure**, letting flow be the consequence.

The controlling element is not the valve: it is the **pump variable speed drive**. PCV-201 remains a manual command of the operator, and from the controller's point of view it is a **disturbance**: if the trainee closes it, the controller must slow the pump down so as not to exceed the set-point. This is the difference that separates a control loop from a direct command, and the panel lets you try it both ways.

3.2 The measurement chain, and why the lag matters

Between what happens at the discharge and what the controller reads there are three lags:

Element	Time constant
Discharge header inertia	1.2 s
Transport delay to the transmitter	0.6 s
Transmitter damping	1.0 s

It is the lags that make instability possible: without them you could raise the gain as much as you like without anything happening, and tuning the controller would teach nothing. With these values, maximum gain produces a 10.4 bar peak-to-peak oscillation, which derivative action at 0.09 s brings back to 3.0 bar.

3.3 The three terms, and what each one does

output = $K_p \cdot \text{error} + K_i \cdot \int \text{error} \cdot dt - K_d \cdot (\text{derivative of the measurement})$

- **Kp · proportional** (0.012 at start) — responds at once, in proportion to the error. On its own it always leaves a permanent offset: to hold the valve open an error is needed that produces that opening.
- **Ki · integral** (0.030 s^{-1}) — accumulates the error over time and cancels it. It is the term that actually brings pressure onto set-point. Too fast, it overshoots and hunts.
- **Kd · derivative** (0 at start) — anticipates, opposing the rate of change.

The derivative acts **on the measurement, not on the error**. This is a deliberate choice: differentiating the error, every time the operator moves the set-point the derivative would see a step and kick the output — the so-called *derivative kick*. Differentiating the measurement alone, a set-point move produces no kick at all, and the derivative keeps doing its job on disturbances. The panel states this in the note under the sliders.

3.4 Anti-windup, and why limiting the output is not enough

When the controller output is saturated — speed already at maximum, or at minimum — the error still exists and the integral keeps accumulating, even though it achieves nothing. On coming out of saturation that accumulation has to be unwound: the machine stays off set-point for as long as it takes to discharge it.

The remedy implemented is **back-calculation**: when the output saturates, the integral is recomputed backwards from the value actually applied, instead of being left to run. The difference is measurable: an earlier version of the simulator integrated even with Ki at zero, making the gains effectively irrelevant — a defect found by testing, not by reading the code.

3.5 Automatic, manual, and bumpless transfer

In **manual** the operator commands the speed directly and the controller looks on. On the **manual → automatic** transition the integral is preloaded to the current speed: the controller restarts from where the machine is instead of from zero, and there is no step. The same realignment happens at the end of a pump changeover while running.

3.6 Parameters of this chapter

Parameter	Start	Range	Effect
Kp proportional gain	0.012	0 to 0.060	Promptness and amplitude of the response
Ki integral action	0.030 s ⁻¹	0 to 0.120	Speed at which offset is cancelled
Kd derivative action	0.000 s	0 to 0.200	Damping of oscillations
Pressure set-point	32.0 bar	14 to 44	Controller target

The *Factory tuning* button returns the three gains to their starting values.

4. The safety instrumented system: 2oo3 voting and SDV-301

4.1 Why voting and not a single transmitter

A protection has two ways of being wrong: **failing to act** when needed, and **acting** when not needed. A single transmitter exposes both: if it fails low, the protection does not trip; if it drifts high, it stops the plant for no reason.

2 out of 3 voting solves both with three independent transmitters, PT-205 A/B/C: the trip acts when **at least two** exceed the setting. One faulty channel, on its own, neither stops anything nor prevents anything. It is the architecture that IEC 61511 and ISA 84 prescribe for safety instrumented functions of this level.

The panel shows the three values separately, each with its own bar and setting marker: it is not decoration, it is how you recognise **which** channel is drifting.

4.2 The drift fault, and how to read it

The *PT-205B drift* fault slowly raises the reading of channel B alone. What you see: B's bar separates from the other two, and when the deviation exceeds 3 bar the **YD-205** suspect channel alarm appears. The plant keeps running: that is exactly the intended behaviour, and anyone without the three separate readings would have no way of understanding what is happening.

4.3 What the trip does

When the voting acts — or when any of the other causes acts — the sequencer moves to the **TRIPPED** step: the pumps stop on a fast ramp and the panel signals the state. In the specific case of overpressure or emergency shutdown, **SDV-301** also closes, the shutdown valve towards the terminal: the line is isolated, not just deprived of head.

4.4 First-out: which cause came first

When a plant stops, the causes pile up within seconds: the pump stops, so pressure collapses, so other alarms arrive. Knowing **which came first** is the difference between understanding the failure and guessing at it.

The alarm list marks the first critical cause with the **First-out** label. The marking stays until a reset is performed.

4.5 ESD and reset

The **emergency shutdown** (HS-001) is a direct operator command: it stops the pumps and closes SDV-301 with no conditions. **Reset trips** rearms the logic, but only if the cause has cleared: it is not a button that cancels the problem, it is the one that rearms the protection after the problem has been solved.

4.6 Settings of this chapter

Tag	Description	Starting value	Voting	Action
PAH-205	High discharge pressure	41.0 bar	—	Alarm
PSHH-205	Very high discharge pressure	46.0 bar	2 out of 3	Trip + SDV-301 closed
HS-001	Emergency shutdown	command	—	Trip + SDV-301 closed
YD-205	Deviation between transmitters	3 bar	—	Suspect channel alarm

5. ISA-18.2 alarms and settings editable on line

5.1 Why alarm management, and not a list

A plant that signals everything signals nothing. The ISA-18.2 standard exists because of this problem: distinguish alarms by **priority**, keep track of their **state**, and make sure the operator sees first what matters.

5.2 The three states

State	Meaning
ACT	Active: the condition is present and the alarm has not been acknowledged
ACK	Acknowledged: the operator has taken it on, the condition is still present
RTN	Returned: the condition has cleared, but the alarm has not been acknowledged yet

The distinction between ACK and RTN is the one most often forgotten. A **returned** alarm stays in the list precisely so that the operator can notice that something happened while looking elsewhere: it disappears only on acknowledgement. For the same reason the banner at the top of the panel follows the worst priority among the **active** alarms, not among all of them: a returned alarm must not keep the console red, or people stop looking at it.

5.3 The three priorities

Critical (red) are the trips and the emergency; **High** (yellow) the conditions that require action but do not stop the plant — cavitation, high vibration, high temperature, high current, seal; **Low** (blue) the status indications — levels, deviation between transmitters.

5.4 Settings change on line, but under a rule

All the panel settings can be changed while running. This is not a convenience: it is how a protection chain is proof-tested in the field, by lowering the setting until it acts.

The rule the panel enforces by itself is the **alarm → trip hierarchy**: the trip setting always stays beyond the alarm one. If the trainee moves one of the two past the other, the second moves accordingly and the note under the table says so. The reason is that a trip acting before its own alarm takes away from the operator the only window of time available to intervene.

A second criterion is not automatic and remains the responsibility of whoever sets the values: **a setting outside the range the machine can reach is not a protection**. A pressure setting placed above the maximum head of the pump (about 50 bar at shut-off) can never act, and one of the training scenarios is built exactly on this.

5.5 Settings of this chapter

All those already seen, plus motor current. The panel highlights the row in alarm in yellow and the one in trip in red, so the table is also a diagnosis.

6. Machine protections: vibration and temperature

6.1 Why vibration is measured

Vibration is the symptom that precedes nearly every mechanical failure of a pump: unbalance, misalignment, bearing degradation, cavitation. It is also the only one that shows up **before** the damage, whereas bearing temperature arrives when the damage is already under way.

In the simulator the vibration of the machine in service is, in healthy conditions:

$$v = 0.85 + 4.4 \cdot n^2 \quad [\text{mm/s}]$$

that is, it grows with the square of speed. Abnormal contributions add to this: cavitation in proportion to the NPSH deficit (up to +6.5 mm/s), a degraded seal (+3.2 mm/s), a failing bearing (+5.4 mm/s).

6.2 The zones and the settings

The settings follow ISO 10816 criteria for machines of this size: **VAH-208 at 7.10 mm/s** is entry into zone C — operation allowed only for limited periods — and **VAHH-208 at 11.0 mm/s** is the trip.

6.3 Bearing temperature

$$T = 34 + 44 \cdot n^2 + 2.6 \cdot (\text{vibration above } 5 \text{ mm/s}) \quad [^{\circ}\text{C}]$$

plus 42 °C if the bearing fault is present. Temperature is **slow**: its time constant is about 38 seconds, against half a second for vibration. Anyone waiting for temperature to notice a failing bearing notices it far later than someone watching vibration — which is why both measurements exist.

TAH-215 at 85 °C alarms, **TAHH-215 at 95 °C** trips.

6.4 The injectable faults and what to watch

Fault	What it does	What shows first
Cavitation	Cuts available NPSH by 7.5 m	Available NPSH below required, vibration rising
Mechanical seal	+3.2 mm/s of vibration	YA-104 alarm, vibration towards zone C
Bearing	+5.4 mm/s and +42 °C	Vibration, then temperature much later
PT-205B drift	One channel out of three rises	B's bar separates, YD-205 alarm

Fault	What it does	What shows first
Feed loss	Vessel feed to zero	Balance draining, then LAL and LSLL

6.5 Settings of this chapter

Tag	Description	Starting value	Action
VAH-208	High vibration, zone C	7.10 mm/s	Alarm
VAHH-208	Very high vibration	11.0 mm/s	Machine tripped
TAH-215	High bearing temperature	85 °C	Alarm
TAHH-215	Very high bearing temperature	95 °C	Machine tripped

7. The start and stop sequence

The sequencer has seven steps, shown in the header:

Step	What happens
STOPPED	Machine stopped, ready
PERMISSIVES	Check: level above LAL-102, no trip active
PRE-LUBE	3 s of preparation
STARTING	Speed ramp up to 35 %
RUNNING	The controller takes over (in automatic) or the operator does (in manual)
RAMP DOWN	Controlled ramp down to a stop
TRIPPED	Fast stop on action of a protection

If the permissives are not satisfied the sequence returns to STOPPED and the panel writes the reason for the refusal under the feed command: that is where the explanation is to be found, not in the alarm list.

The time scale is ×3: one second of clock is worth three seconds of process. The times shown on the panels are process times; the exam session stopwatch, by contrast, runs in real time.

8. What the model simplifies

Knowing where the simulator stops is part of the training. The declared simplifications:

- **The two pumps are identical.** No difference in wear, tuning or efficiency between P-101 A and P-101 B.
- **The 2 out of 3 voting works on three readings derived from the same pressure**, with independent noise and, when the fault is injected, a drift applied to channel B alone. They are not three physically independent measurement chains with their own calibration errors.
- **The line to the terminal is a constant pressure** (8.0 bar). There are no line transients, water hammer or back-pressure changes from the terminal.
- **Cavitation is represented through its effects** — loss of NPSH, vibration, deterioration of the operating point — but not as a two-phase phenomenon with progressive impeller erosion.
- **Crude temperature is constant.** No viscosity effect on the pump curve nor on line losses.

- **The vessel has no stratification and no sediment:** one level, one fluid.
- **The seal fault produces no loss of flow:** it raises vibration and generates its own alarm, but the leakage does not appear in the mass balance.

None of these simplifications touches what the panel is there to train: reading the balance of a vessel, tuning a control loop, recognising cavitation, understanding 2 out of 3 voting and rebuilding the hierarchy between alarms and trips.