

What a Level Switch Can — and Cannot — Tell You About a Leak

Reservoir inventory, containment accumulation,
and fail-safe interlocking in liquid-cooled systems.

Compac Engineering Inc. · Paradise, California

Issued 15 October 2026

SUMMARY

A level switch does not detect leaks. It detects one of two *consequences* of a leak, in two different places, and the two are not interchangeable:

- **Loop inventory sensing** — a falling level in the reservoir of a closed loop. Catches loss from anywhere in the circuit. Tells you nothing about where.
- **Containment accumulation sensing** — fluid arriving in a basin, drip pan, or sump where there should be none. Tells you roughly where. Only reports fluid that has already escaped.

The distinction matters because the two topologies have different sensitivity floors, different failure modes, and opposite wiring conventions if you want the installation to be fail-safe. Specifying one when the application needs the other is the most common way a correctly-rated switch ends up in a system that still doesn't notice the leak.

This paper covers how to choose between them, how to calculate the smallest loss each can resolve, how to keep thermal expansion from masking a real leak or faking one, and how to wire either so that a broken wire reports a fault rather than hiding one.

Scope. Closed-loop liquid systems with a vented or low-pressure reservoir — coolant distribution units, chiller and glycol loops, machine-tool coolant systems, process recirculation skids, and deionized-water polishing loops. The arithmetic applies to any closed liquid circuit with a reservoir.

1 The question behind the question

“Can I use a level switch for leak detection?” is almost always asked when a designer has a closed loop, a pump that must not run dry, and a requirement that

says the machine shall detect loss of coolant. One component looks like it could satisfy all three.

It can satisfy two of them well, and the third only within limits that are worth knowing before the drawing is released.

A point level switch is a **threshold device**. It reports one bit: the level is above the actuation point, or it is below it. That bit is extremely reliable — a hermetically sealed reed switch has no contact with the liquid, no mechanical wear surface in the signal path, and a rated life of 100 million operations for the

Standex/OKI ORD229 element used across Series 10 and Series 15. What it is not is a measurement. It will not tell you a rate, a trend, or a volume.

Any leak-detection claim built on it is really a claim about **how much fluid must be lost before the bit flips** — and that number is set by the tank geometry and the setpoint, not by the switch. Once the question is framed that way, it becomes answerable with arithmetic.

2-3 The two topologies

TOPOLOGY A

Loop inventory

The switch sits in the reservoir. As fluid leaves the circuit the reservoir level falls, because in a closed loop the reservoir is the only place the inventory can come from.

TOPOLOGY B

Containment accumulation

The switch sits in the basin, drip pan, or sump under the equipment. Normally dry. Actuation means fluid is somewhere it should not be.

A — what loop inventory catches, and what it misses

Catches: loss from anywhere — a weeping fitting three metres downstream, a failed quick-disconnect, a cracked cold plate, a seeping pump seal. The reservoir doesn't care where the fluid went.

- It **cannot locate** the leak. One bit, no direction.
- It **cannot distinguish** a leak from legitimate inventory change — thermal expansion, air purging out of the loop after service, or a filter change that retained fluid. §4.2 deals with this.
- It **cannot see a loss smaller than its resolution**, which may be larger than you expect.

The second job is usually the one that justifies it. A reservoir level switch is almost always doing double duty, and the second duty is **dry-run protection for the pump**. A centrifugal pump run dry loses its prime, loses its cooling, and can destroy a seal face in seconds. The same low-level bit that is weak evidence of a slow leak is excellent protection against that — because dry-run is exactly the case where the level has fallen a *lot*, fast.

Specify the reservoir switch for dry-run protection first. Treat its leak-detection value as the bonus it genuinely is, and size it per §4 so you know what that bonus is worth.

Placement. The sensing point belongs in the reservoir **upstream of the pump inlet**, at or above the level at which the inlet begins to draw air. Not in the discharge line, not after a filter, and not where the pump's own suction draws the surface down locally.

B — what containment accumulation catches, and what it misses

Catches: fluid that has physically escaped, in the zone the pan covers. Because a pan is local, a pan switch *does* localise — one pan per rack, per CDU, per manifold,

and the alarm tells you which.

- It cannot see a leak that lands **outside the pan footprint**, runs down a cable tray, or wicks into insulation.
- It cannot see a leak that **evaporates as fast as it arrives**. A slow weep onto a warm surface may never accumulate.
- It reports **only after escape**. There is no early warning; by definition the fluid is already out.

Sensitivity is a geometry problem, and it is the opposite problem from the reservoir. A pan is wide and shallow, so a small volume spreads thin. The switch must actuate at the lowest practical height above the pan floor, which rules out long-stem float assemblies and favours a **side-mount switch through the pan wall**. Series 15 is **4.45 in** overall with a **0.70 in** float, installing outside-in through **1/2 NPT** or inside-out through a **1/4 in** straight thread with jam nut and O-ring — and needs no headroom above the pan at all.

A sump, or a deliberately-formed low corner in the pan, improves sensitivity substantially, because it converts a wide shallow volume into a narrow deep one. If the enclosure design is still open, a 2–3 in sump under the sensing point is the cheapest sensitivity improvement available.

4 Calculating what you can actually resolve

4.1 Level change per unit volume

For a reservoir of constant cross-section, a lost volume V produces a level drop:

$$\Delta h = V / A$$

where A is the free surface area. For a cylindrical reservoir of diameter D , $A = \pi D^2/4$. Working in US gallons ($1 \text{ gal} = 231 \text{ in}^3$):

RESERVOIR DIA.	SURFACE AREA	LEVEL CHANGE PER GALLON
8 in	50.3 in ²	4.60 in
12 in	113.1 in ²	2.04 in
16 in	201.1 in ²	1.15 in
18 in	254.5 in ²	0.91 in
24 in	452.4 in ²	0.51 in
36 in	1,017.9 in ²	0.23 in

The design consequence: **a wide, shallow reservoir is a poor leak sensor**. At 36 in diameter a full gallon moves the surface less than a quarter inch. The switch's own actuation band — the level difference between actuate and reset, published per model — must be comfortably smaller than the Δh you intend to catch, and on a wide reservoir it may not be.

TAKE IT FROM THE DATASHEET

Read the actuation band from the datasheet for the model you are specifying. Do not assume it; it varies with float geometry and mounting.

4.2 Thermal expansion — the error that masks real leaks and invents false ones

Coolant volume changes with temperature. A loop that is full at operating temperature is *not* full at ambient, and the difference is routinely larger than the leak you were hoping to detect. This is the single most common reason an inventory-based leak threshold either never trips, or trips every night when the load drops.

Method. Do not guess a coefficient — take the volumetric expansion figure from the coolant manufacturer's data sheet for your specific fluid and concentration, then:

1. Compute the loop's total volume change across the full operating temperature band, from coldest cold-start to hottest steady state:

$$\Delta V_{\text{thermal}} = V_{\text{loop}} \times (\text{expansion over the band})$$

2. Convert to level change using the reservoir area:

$$\Delta h_{\text{thermal}} = \Delta V_{\text{thermal}} / A$$

3. Place the low-level trip **below the coldest normal operating level**, by more than $\Delta h_{\text{thermal}}$, or accept nuisance trips on cooldown.
4. Compare your target leak volume against $\Delta V_{\text{thermal}}$ **as volumes**.

Step 4 decides whether the approach works at all, and it has a property worth understanding: **the leak-to-thermal ratio is independent of reservoir geometry**. Both scale with $1/A$, so a bigger reservoir does not improve your thermal margin — it only makes the level change harder to resolve.

Reservoir geometry affects **switch resolution**. It does not affect **thermal masking**.

4.3 Worked example

A CDU with a 25 gallon loop inventory and a 16 in diameter vented reservoir. From the table, **1.15 in** per gallon.

QUANTITY	VOLUME	LEVEL	VERDICT
Target leak to detect	0.500 gal	0.57 in	Actuation band must sit well inside this
Thermal swing @ 0.35 %	0.0875 gal	0.10 in	Trip point goes at least this far below coldest level
Margin, 0.5 gal target	5.7 : 1		Comfortable — distinguishable from thermal breathing
Margin, 0.1 gal target	1.1 : 1		Leak is the same size as the noise

The 0.35 % figure above is a stand-in for whatever your coolant data sheet yields; the method is what transfers. Change only the target to **0.1 gallon** and the ratio against the same thermal swing falls to 1.1 : 1. No reservoir geometry and no better switch fixes that, because the ambiguity is volumetric.

This is the paper's central conclusion. Inventory sensing resolves large, fast losses. Containment sensing resolves small, slow ones, because a pan has no thermal baseline to subtract — dry is dry.

The two topologies are **complementary, not alternative**. A system that needs to catch both a burst and a weep needs both — and that is a two-switch answer, not a better single switch.

5 Fail-safe wiring — and why both topologies land in the same place

5.1 Define the terms, because convention will mislead you

“N.O.” and “N.C.” are ambiguous for float switches: they describe the contact in the *de-actuated* position, and whether de-actuated means float-up or float-down depends on mounting. Series 15 makes this explicit by setting contact form **by**

installation orientation — one SKU per material, with an arrow on the body indicating direction.

So specify behaviour, not label:

- **Closes on rise** — contact closes as the float lifts.
- **Opens on rise** — contact opens as the float lifts.

5.2 The permissive principle

A protective circuit should be **closed in the healthy state**, so that the fault condition and the circuit's own failure modes produce the same result. A broken wire, a corroded splice, an unplugged connector, and a switch that has failed open must all read as *fault*.

INSTALLATION	HEALTHY STATE	REQUIRED BEHAVIOUR	BROKEN WIRE READS AS
Reservoir dry-run / inventory	Level adequate, float up	Closes on rise	Low level → pump stops ✓
Containment pan leak alarm	Pan dry, float down	Opens on rise	Leak → alarm ✓

Both end up **closed when the system is healthy and open on fault** — despite being opposite orientations and opposite-sounding specifications. Wire them in series into a single healthy-state permissive string and any one fault, or any break in the string, drops the whole chain.

THE INTUITIVE WIRING IS THE WRONG WIRING

Closing the circuit when liquid is present is the obvious scheme for a pan switch, and it is not fail-safe. A severed cable in that scheme reports “no leak” forever — and a pan switch is precisely the device most likely to have its cable damaged, because it lives on the floor of the enclosure where the service technician kneels.

5.3 What fail-safe wiring still does not cover

A contact that has failed **shorted**, or a wire pinched into a short across the contact, presents as healthy in either scheme. Detecting that requires a supervised circuit — end-of-line resistor, series supervision, or a periodic forced test — and that is a control-system design task, not something a mechanical switch provides. State the limitation in the FMEA rather than discovering it in the field.

6 Interlock, not telemetry

There is a strong case for a hardwired switch in a system that already has instrumented monitoring, and it is not redundancy for its own sake.

A level switch wired in series with the pump contactor coil is a **physical permissive**. It functions with the building management system offline, the network partitioned, the controller mid-firmware-update, or the monitoring software in a state nobody anticipated. There is no firmware in the path, no calibration to drift, no configuration to be lost in a replacement, and no addressing to be wrong. It is the layer that still works when the layer above it does not.

That is a different job from telemetry, and it should be specified as a different job. Analog level transmitters, flow meters, conductivity and pH monitoring, and differential pressure measurement all answer *“how is the loop doing?”* A switch answers *“may this pump run?”* Systems that conflate the two usually end up with good dashboards and no interlock.

6.1 Driving the load — check the contact rating against the coil

The interlock only works if the contact can carry what you put across it, and **contactor and solenoid coils are inductive loads**. Published switch ratings are resistive. This is where level-switch interlocks most often get specified wrong.

SERIES	SWITCH ELEMENT	RESISTIVE RATING	INDUCTIVE CAPABILITY
Series 10	Standex/OKI ORD229 reed	70 VA · 0.5 A AC / 0.7 A DC	Requires external arc suppression
Series 15	Standex/OKI ORD229 reed	70 VA / 50 W · 0.5 A AC / 0.7 A DC 300 VAC / 350 VDC max	Requires external arc suppression
Series 20	Yaskawa Bestact™ R25	240 V – make 10 A / break 0.5 A 115 VDC break 0.3 A	120 VA @ 240 VAC rated

Three practical consequences:

1. **Series 20 can drive a typical contactor coil directly**, within its 120 VA inductive rating. Its Bestact R25 element is also the one with a published *make* rating — 10 A at 240 V — which matters for coil inrush.
2. **Series 10 and Series 15 reed elements need help** with an inductive coil. Add external arc suppression, or isolate the coil behind a **Series 50 solid-state relay** and let the reed switch only the relay input. An unprotected reed switching a coil will weld or erode its contacts long before the 100-million-operation rated life.
3. **Mechanical life differs from contact life.** Series 20 is rated 100 K operations AC / 300 K DC. For a reservoir switch that cycles on every makeup event, check the duty cycle against that figure; for a containment pan switch that should actuate approximately never, it is irrelevant.

Arc-suppression circuit selection, including nine worked cases for DC and AC inductive loads and an RC value selection table, is covered in **TN-120** —

Protecting Relay and Switch Contacts.

7 Selecting hardware: pressure is the real constraint

For water-glycol service, chemical compatibility is usually *not* what limits the choice, and designers frequently optimise the wrong variable.

7.1 Temperature limits, and which one governs at each end

THE UPPER LIMIT — THREE NUMBERS, THREE QUESTIONS

A float switch in a glycol loop has at least three temperature figures attached to it, and they are routinely read as competing answers to one question. They are not. They answer three different questions:

- **What does the polymer tolerate?** — the plastic's own temperature rating, in isolation.
- **What does the polymer tolerate *in this medium*?** — the chemical compatibility rating, which is medium-specific and can be higher or lower than the generic figure.
- **Does the assembled switch still seal, and still switch?** — the manufacturer's assembly rating. **This is the one that governs function.**

What the 90 °C ERECTA SWITCH rating actually is. It is not a conservative restatement of a plastic's limit. It is an assembly rating, and it accounts for what no material datasheet can: **part fits, each plastic's coefficient of thermal expansion, the effect of that differential expansion on leakage, and switch operation.**

A float level switch is an assembly of mating injection-moulded parts whose clearances change with temperature — and at different rates, because the parts are not all the same polymer. A reed switch depends on float-to-switch geometry. A seal depends on its fit. A polymer's temperature rating says nothing about any of that. This is also why the figure is a flat 90 °C across Series 10, 15, and 20 rather than varying by material: it rates the assembly, not the plastic in it.

Compac publishes it as the limit of its own warranty of function. Above 90 °C the decision belongs to the OEM engineer, who knows the duty cycle, the mounting, and the consequence of a leak. That is a statement about where responsibility sits — not a claim that the part fails at 91 °C.

The design rule is simply to take the lower of the two applicable figures — compatibility in your medium, and the assembly rating — every time. Which of the two binds depends on the material, and for one of the three it is not the one a designer would assume.

MATERIAL	A-RATED IN GLYCOL TO	ASSEMBLY MAX	GOVERNING LIMIT
Polypropylene (PP)	225 °F (107 °C)	90 °C (194 °F)	Assembly — 90 °C
Acetal (AC)	140 °F (60 °C)	90 °C (194 °F)	Material — 60 °C
PVDF Kynar (KR)	275 °F (135 °C)	90 °C (194 °F)	Assembly — 90 °C

For PP and Kynar the assembly binds — 31 °F below polypropylene's rating and 81 °F below Kynar's. **Selecting Kynar to gain temperature headroom in a glycol loop buys nothing**, because the assembly caps out first. Kynar is the correct

specification when the *chemistry* demands it — an aggressive inhibitor package, an oxidising biocide, or a high-purity requirement — not when the temperature does.

ACETAL IS THE EXCEPTION, AND THE TRAP

In glycol, Acetal is A-rated only to **60 °C (140 °F) — 30 °C below the switch assembly's own 90 °C rating**. It is the one material where the plastic binds first, and the one case where reading “90 °C max” off a ratings table gives a part rated 54 °F lower than the number in front of you.

Acetal remains the correct choice for hydrocarbons, where PP is incompatible. In a glycol loop it is usable but narrow, and rarely the right pick for a circuit that may run warm or stagnate hot.

Polypropylene is the standard specification for water-glycol — lowest cost, and 17 °C of margin over the assembly limit it never reaches.

For context, the 90 °C ceiling sits far above normal liquid-cooling supply temperatures; ASHRAE's facility-water classes for liquid cooling run from roughly 17 °C to 45 °C supply. Confirm against your own loop's worst-case excursion, including stagnation after a pump failure, rather than against the steady-state design point — and note that a 60 °C Acetal limit leaves far less room for that excursion than a 90 °C one.

THE LOWER LIMIT — THREE MECHANISMS, AND THE MEDIUM SETS IT

The minimum is **above the freezing point of the medium**. That is a materially different statement from “above 0 °C”, and the difference matters most in exactly the application this paper is about.

Three separate things go wrong at and below the freeze point. They are worth separating because their consequences differ:

1. **The switch stops working.** A float cannot float in a frozen or slushy liquid, and a piston-type flow switch cannot move. A *functional* failure with no damage —

and a silent one, because a switch frozen in its healthy state reports healthy.

2. **The plastic becomes brittle.** Impact and shock resistance fall with temperature. This is a property of the polymer and **does not move with the coolant chemistry.**
3. **The medium expands as it freezes, and can break the body.** The destructive mechanism, and the only one of the three that is permanent.

Glycol is in the loop precisely to move mechanisms 1 and 3 out of the way. A water-glycol mix at an appropriate concentration has a depressed freeze point, so a glycol loop can legitimately operate below 0 °C — as outdoor dry coolers and cold-climate chiller loops routinely do. Take the freeze point from the coolant data sheet at *your* concentration rather than a generic figure, and remember that concentration drifts with makeup water over service life. Mechanism 2 does not move with it, and no amount of glycol addresses it.

Where this bites: the stagnant branch. A circulating loop resists freezing; an isolated one does not. A valved-off bypass, a dead leg, a drained-but-not-dry branch, or a reservoir on an unheated outdoor skid can freeze locally while the main loop is perfectly warm. Those are also precisely the places a protective switch gets installed — a flow switch in a bypass is in a low-flow branch by design.

DESIGNING THE STAGNANT BRANCH OUT — A TEMPERATURE SWITCH THAT ESTABLISHES FLOW

That hazard has a direct answer, and it is worth designing in rather than accepting. Sense the liquid temperature in the branch and **open a valve to establish flow before it freezes.** The primary mechanism is heat transport — flow continuously replaces cooling fluid with warmer fluid from the active part of the loop — and moving liquid also resists ice nucleation.

A Series 11 bimetallic switch (the **11-800** element) installed in a **pipe tee** does this directly:

- **1/8" NPT**, 0.5" body and 0.5" hex — fits a standard tee without a custom boss.
- Medium is **gas or liquid**, so it senses the fluid itself rather than pipe-wall temperature.
- 16 factory setpoints across 0–100 °C, codes A–P.
- **5 A at 120 VAC inductive** — enough to drive a solenoid valve coil *directly*, which the reed-based series cannot do without arc suppression or a Series 50 relay (\$6.1).

±5 °C ACCURACY IS THE NUMBER THAT GOVERNS THIS DESIGN

Series 11 accuracy is **±5 °C typical**. Choose a setpoint at least the accuracy band above the medium's freeze point, or the switch may not act until the fluid is already freezing — which is the one outcome the whole arrangement exists to prevent.

And **the setpoint is a part number, not an adjustment**. Each of the 16 codes is a distinct SKU with no field trim, so the margin calculation has to be right before the order goes in.

Creep-action gives a **slow response**, which is the correct behaviour here: a pipe's thermal mass makes freezing a slow event, and a slow switch will not chatter around the setpoint.

Wiring it fail-safe is \$5.2 applied to a valve instead of a pump. The protective action is *valve open*, and failing to act is destructive — so specify a valve that **fails open** (de-energised open) and wire the switch to energise it shut only while the fluid is warm, i.e. contact closed on rising temperature:

STATE	CONTACT	VALVE COIL	RESULT
Warm — normal	closed	energised	Valve shut, no bleed ✓
Cold — approaching freeze	opens	de-energises	Valve opens, flow established ✓
Broken wire / failed-open switch	open	de-energises	Valve opens, flow established ✓

Every failure mode drives to the protective action, and the circuit is closed in the normal state. Note that this is the *opposite* valve convention from the one a designer reaches for first — a normally-closed valve opened by the switch would leave a severed cable with the branch unprotected and no indication.

The constraint: Series 11 is rated 50 PSI at 20 °C, so confirm the branch pressure before specifying it in line. A freeze-protection bleed on a low-pressure or vented branch is comfortably inside that; a pressurised main is not.

A CONTAINMENT PAN IS THE WORST CASE OF THE THREE

What collects in a pan is whatever leaked, at whatever concentration it happened to be, in a space that may not be conditioned. If that fluid **freezes with the float down, the pan switch stays in its healthy state and will not alarm** — the same class of lying failure as the stuck-shut check valve in §7.4, and it defeats the §5.2 fail-safe wiring entirely. Freezing with the float up alarms, which is the safe direction.

Treat a pan in an unconditioned space as requiring heat, relocation inside the conditioned envelope, or an explicit acceptance of the limitation in the FMEA. Do not assume the pan inherits the loop's freeze protection: it holds escaped fluid, not loop fluid, and a dilute leak freezes at a higher temperature than the loop it came from.

7.2 Specific gravity: check it, then stop worrying about it

Every float switch has a minimum specific gravity below which the float will not lift. The ERECTA SWITCH minimum is **0.80**. Water-glycol mixtures sit at or slightly above 1.0, so the floor is satisfied with wide margin. Take the figure from the coolant data sheet and confirm it; this is a check, not a constraint, for aqueous coolants. It becomes a real constraint for light hydrocarbons and some dielectric fluids.

7.3 Pressure: this is what actually decides the series

SERIES	MAX PRESSURE @ 20 °C	DERATING	SUITS
Series 5	200 PSI	Derates to zero @ 90 °C	Inline flow sensing; weak link is the switch-to-fitting seal
Series 10	100 PSI	Derates to zero @ 90 °C	Pressurised reservoirs and in-circuit vessels
Series 15	50 PSI	Derates to zero @ 90 °C	Side-mount through a tank or pan wall
Series 20	20 PSI	Derates to zero @ 90 °C	Vented / near-atmospheric reservoirs, 2 in NPT top mount

What a published pressure rating is — and the number it is not. These figures come from empirical testing, and each has a named weak link. For Series 5 the weak link is not the body: it is **the seal between the flow switch and the accessory fittings**. That seal has been tested successfully at **500 PSI**. The published rating is **200 PSI**.

The gap between those two numbers is the most useful thing to understand about a thermoplastic pressure rating. **A proof test tells you about burst. It tells you nothing about service life.** Held under sustained, unrelieved stress a thermoplastic *creeps* — it flows slowly and permanently — so a joint that holds 500 PSI on a bench can relax and weep after months at a fraction of it. The published figure is a stress the joint can hold indefinitely, not the pressure at which it fails.

Which is also why **a short hydrostatic test on your own assembly does not validate a higher operating pressure**. It validates the wrong failure mode.

READ THE DERATING COLUMN CAREFULLY — AND KNOW WHY IT DERATES

Every series in the table derates to zero at 90 °C — Series 5, 10, 15 and 20 alike. There is no ERECTA SWITCH series that holds its rated pressure at its rated temperature. That is the same creep mechanism with the other variable moved: **creep rate rises with stress and with temperature**. The same seal at the same pressure flows faster when hot, so the derating curve is not a cliff at which something breaks — it is the point beyond which no sustained pressure can be warranted over service life.

So the pressure and temperature ratings are not independent, and a specification that reads “50 PSI, 90 °C” — taken from two separate rows of one table — describes an operating point the part does not have.

Most CDU and chiller reservoirs are vented and near atmospheric, which puts all three in range and makes mounting geometry the deciding factor. A pressurised

vessel narrows it to Series 10.

Series 20 additionally carries **UL Class 1, Division 2, Groups A–D**, which is not typically required in a data hall but is relevant for industrial and process installations of the same hardware. Series 15 is UL File **E70063**.

7.4 Flow verification — and the four places to put the switch

Flow switches are routinely specified alongside level switches for the same protective function, and the selection logic is less obvious than it looks, because **a flow switch is a restriction in the line**. That single fact governs every architecture below.

Series 5 covers **0.008 to 1.00 GPM** of actuation setpoint — five standard points from 0.10 to 1.00 GPM, with the low-flow set actuating at 0.042 GPM with spring, or 0.008 GPM springless and vertically mounted. A corrosion-proof titanium spring, and a built-in clean-out port for in-line service.

The cost in pressure, for a **5–20** angle-body switch in water at 20 °C:

SETPOINT	ΔP AT ACTUATION	ΔP AT 3 GPM
0.10 GPM	0.5 PSI	30 PSI
0.25 GPM	0.6 PSI	30 PSI
0.50 GPM	0.75 PSI	25 PSI
0.75 GPM	1.2 PSI	22 PSI
1.00 GPM	1.4 PSI	20 PSI

Characteristic performance $\pm 10\%$, catalogue G520-4. The 5-21 straight body is lower throughout — 0.22 to 0.7 PSI at actuation.

Read the two columns together. At the actuation threshold the switch costs **well under 1.5 PSI**. Push 3 GPM through the same switch and it costs **20 to 30 PSI** — between **14 and 60 times** the actuation figure, depending on setpoint, and worst

for the most sensitive switches. This is why 3 GPM is the recommended inline maximum, and it is the reason the three architectures after the first one exist. The limit is not the switch's ability to sense. It is what the switch does to the line while sensing.

ARCHITECTURE A

Inline

Lines at or below 3 GPM. Simplest, cheapest, lowest parts count — and at 0.008 GPM it resolves flows most inline meters cannot see at all. The default for makeup lines, filter and polishing bypasses, deionised-water loops, seal-flush lines, and single cold-plate branches.

ARCHITECTURE B

Per-branch — one switch per served path

For a large system the instinct is to find a switch that handles the main flow. Invert it: put a cheap switch at the **inflow of each cooling path** instead. Branch flows sit in Series 5's native range, so no bypass is needed, and the result is categorically better information.

This is the same argument as §3, moved from level to flow. A single main-line meter tells you the loop is flowing. A per-branch switch tells you **which branch stopped** — and a blocked cold plate, a kinked hose, or a mis-seated quick-disconnect is a per-branch failure, not a loop failure. **Localisation comes from placement, not from the instrument.** What makes it practical is unit cost: this architecture is only affordable if the switch is cheap enough to fit dozens of them.

ARCHITECTURE C

Fixed-constriction bypass

Divert a small percentage of system flow through an unobstructed, close-coupled bypass placed **across a constriction** in the main line. The differential pressure across the constriction drives the bypass. With a 0.1 GPM setpoint switch, under 1 pint per minute at roughly 3 ounces of pressure actuates it — about **1 % of the flow in a 10 GPM system**. The constriction may be deliberate plumbing, or may already exist as a filter or a reduced section; as little as 7 % is sufficient.

This is the method in current ERECTA SWITCH literature, and it works well in the middle and upper part of the range. **What it cannot do is sense very low flows.** An orifice differential falls with the square of flow, so as main-line flow drops the bypass is starved precisely when you most want a reading.

ARCHITECTURE D

Bypass with a spring-opposed check valve

Replace the fixed constriction with a **variable** one, and the low-flow blind spot disappears. A spring-opposed check valve sits in the main line with the switch in the bypass around it. At low flow the valve is shut and *all* flow is forced through the bypass, so the switch can sense right down to its own setpoint. It is the only architecture here that delivers high low-flow sensitivity *and* low high-flow pressure drop.

For Architecture C, the constriction diameter sets the main-line trip point. For a 5-21 set at 0.1 GPM, water, constriction in 1/2" Sch 40 pipe:

CONSTRICTION DIA. (IN)	0.187	0.250	0.312	0.375	0.437	0.500	0.562	0.602
Main-line GPM, pull-in	1.05	1.75	2.55	3.5	4.75	6.0	7.5	—
Main-line GPM, drop-out	1.0	1.7	2.5	3.3	4.5	5.75	7.25	—

Hysteresis is tight — a few percent across the range — which is what makes the arrangement usable as a threshold rather than merely as a flow/no-flow indication. **The limitation is that the constriction is a fixed penalty at every flow rate.** Size it to trip at 1 GPM and you carry that restriction at 40 GPM too.

How Architecture D behaves across the range

- **At low flow the valve is shut.** Spring preload exceeds the bypass pressure drop, the main line is blocked, and effectively **all** flow routes through the bypass. The switch sees the full main-line flow — so the sensing threshold is the switch's own setpoint, not some multiple of it.
- **The bypass itself is the sensing restriction.** As bypass flow rises, the pressure drop across it — switch plus plumbing — rises too. Because the bypass and the check valve sit in parallel, that same differential appears across the valve. Nothing else has to be added to the main line to generate it.
- **When the bypass differential exceeds the spring preload, the valve begins to open** — against spring resistance, so it *divides* flow rather than dumping it. Some flow continues through the bypass at every higher rate, and that is what keeps the switch made.
- **At high flow the bulk passes the open valve,** so main-line pressure drop stays low instead of growing with the square of flow through a fixed orifice. The spring's characteristic sustains the bypass differential as flow climbs.

What this buys, in the paper's own numbers. From the Architecture C table, the most aggressive constriction tested — 0.187" in 1/2" Sch 40 pipe — still needs **1.05 GPM** in the main line to trip a switch set at 0.1 GPM. Because an orifice differential falls with the square of flow, no constriction recovers the bottom of the range; the published table simply stops there.

With the check valve shut, the main-line sensing threshold *is* the switch setpoint — **0.1 GPM**. Same pipe, same switch, same setpoint: a better than **tenfold improvement in low-flow sensing**. Against the springless 0.008 GPM configuration the gap is wider still.

Three design points, and one of them is a genuine hazard rather than a nuisance:

1. **Size the bypass to deliver at least the setpoint** at the lowest main-line flow that must be proven. The valve's crack pressure sets where the handover happens; below it, bypass flow and main-line flow are the same thing.
2. **A blocked bypass reads as no-flow**, and a **valve stuck open** starves the bypass at low flow and also reads as no-flow. Both are fail-safe in the sense of §5.2 — they alarm.
3. **A valve stuck shut is not fail-safe**. All flow routes through the bypass, the switch reports flow present, and that reading is true — but the served load is now receiving only bypass flow, which may be a small fraction of what it needs.

THE ONE FAILURE MODE THAT LIES

In the stuck-shut case the switch is answering “is there flow in the bypass”, and that has stopped being a proxy for “is the load being served”. Put it in the FMEA. Where the consequence is unacceptable, prove flow *at the load* with Architecture B rather than proving it at the main.

Selection order for flow: B if the served paths are individually addressable and localisation has value — which, in a multi-server or multi-zone cooling system, it almost always does. **A** if there is one path at or below 3 GPM. **D** where one line

must be proven across a wide flow range, and especially where low flow must be detected on a line that also carries high flow. **C** where a constriction already exists in the plumbing for other reasons and the bottom of the range does not need proving.

8 What this approach cannot do

Stated plainly, because a protection scheme is only as good as the honesty of its stated limits. A level switch cannot:

- **Locate a leak within a zone.** Pan granularity is the resolution. Finer localisation means more pans and more switches, or a different technology.
- **Measure rate or trend.** A point switch produces one bit. Rate-of-loss monitoring requires a continuous transmitter.
- **Detect loss below the resolution floor** derived in §4.
- **Detect vapour-phase loss.** In immersion and two-phase systems, fluid can leave as vapour with no liquid ever arriving in a pan.
- **Monitor coolant chemistry.** Conductivity, pH, inhibitor depletion, and particulate loading are all invisible to a level switch and all matter to loop life.
- **Catch a shorted contact or a shorted cable,** per §5.3.
- **Alarm once the pan contents have frozen with the float down,** per §7.1 — a switch frozen in its healthy state reports healthy.

Where the requirement is continuous telemetry, chemistry monitoring, or sub-zone localisation, a level switch is the wrong instrument and should be specified alongside the right one rather than instead of it.

Selection summary

REQUIREMENT	TOPOLOGY	CANDIDATE	BEHAVIOUR	NOTE
Pump dry-run protection	Reservoir	Series 10 (stem) or Series 20 (2 in top mount)	Closes on rise	Sense upstream of pump inlet
Large / fast loss detection	Reservoir	Series 10 / 15 / 20	Closes on rise	Size per §4; usually the same switch as dry-run
Small / slow leak detection	Containment pan	Series 15 side-mount	Opens on rise	Lowest practical actuation height; add a sump
Pressurised vessel	Reservoir	Series 10 (100 PSI)	Either	Highest of the level series — but every series derates to zero at 90 °C
Direct contactor coil switching	Either	Series 20 (120 VA inductive)	Either	Reed series need suppression or a Series 50 relay
Flow proof, single path ≤ 3 GPM	Flow — inline (A)	Series 5 (0.008–1.00 GPM)	SPST N.O., closes on rising flow	ΔP under 1.5 PSI at actuation
Flow proof, which branch stopped	Flow — per-branch (B)	Series 5, one per served path	SPST N.O., closes on rising flow	Localises; needs low unit cost to be practical
Flow proof, wide-range main	Flow — bypass (C / D)	Series 5 in bypass	SPST N.O., closes on rising flow	D (spring-opposed check valve) senses to setpoint and keeps ΔP low; C is blind at low flow

REQUIREMENT	TOPOLOGY	CANDIDATE	BEHAVIOUR	NOTE
Glycol loop, may run warm	Material	PP (90 °C) — not Acetal	—	Acetal in glycol is A-rated only to 60 °C
Loop temperature limit	Temperature	Series 11 (0–100 °C, 16 setpoints)	COR / OOR	6 A @ 120 VAC resistive, 1/8 in NPT, 50 PSI
Freeze protection, stagnant branch	Temperature in a pipe tee	Series 11 / 11-800	Closed on rising temp	Drives a fail-open valve directly at 5 A inductive; setpoint $\geq \pm 5$ °C above freeze point

The three decisions that matter most, in order

1 Which consequence are you detecting

—
 falling
 inventory,
 or
 arriving
 fluid?
 Get
 this
 wrong
 and
 nothing
 downstream
 recovers
 it.

2 What is the smallest loss that must trip it

, *volumes*
 and

does
§4
say
that
is
resolvable?
Compare
for *level change*
thermal
masking;
compare
against
the
actuation
band
for
switch
resolution.

3 Is the healthy state a closed circuit?

If a
severed
cable
reads
as
"fine",
the
installation
is
not
a
protective
device.

Related library documents

TN-108 — Thermal Expansion · bears directly on the §7.1 assembly rating

TN-111 — Understanding The Flow Switch Bypass · the §7.4 architectures

TN-120 — Protecting Relay and Switch Contacts (arc suppression, nine circuit cases, RC selection)

Why Reed Switches Need Contact Protection — white paper

Series 10 — Product Family Overview and Ratings

Series 15 — Application Notes

Series 20 — Operating Principle and Wiring Diagrams

Design Engineer's Guide — Series 50 Solid State Relay Selection

Design Engineer's Guide — Series 11 Bimetallic Temperature Switch Selection

Compac Engineering Inc. has manufactured ERECTA SWITCH liquid level, flow, and temperature switches in Paradise, California since 1964. No minimum order quantities. Free evaluation samples. Engineering support by phone at 1-800-828-9044.

Ratings in this document are published ERECTA SWITCH specifications current at the date of issue. Confirm against the datasheet for the specific model and material before releasing a design. Coolant properties — specific gravity, volumetric expansion, and material compatibility for a particular inhibitor package — must be taken from the fluid manufacturer's data sheet.