


Why the BMS is part of the Quasar family

Quasar is an AI harness: how EIS data from 10 HiBlade batteries becomes assessed, independently checked advance warnings.

Safiery Pty Ltd · 30 September 2026



1. Summary

The Quasar BMS in every HiBlade pack is part of the Quasar family because it is the first gate of the Quasar harness. The BMS makes the EIS measurement, decides whether the measurement is valid, and acts locally on anything imminent. Quasar ONE then carries the validated results to the Quasar Twin, where the harness reads every cell group in every pack against its history and against Safiery's signatures.

Quasar is not an AI agent. It is the harness an agent works inside: the measurements, the validity rules, four sets of learned signatures, the history, the gates and the record. The AI model is a replaceable tenant. Safiery's knowledge of how HiBlade cells age and fail lives in the harness, not in the model.

For a 10-pack installation this means 140 cell groups read together every day. Each is compared with its own baseline, its siblings and the other packs, at matched temperature and state of charge. Any warning is checked by a second, independent agent before it is issued, and protection never waits for either of them.

2. Quasar is an AI harness, not an AI agent

An AI agent is a large language model given tools and a task. On its own it knows general electrochemistry. It does not know how the specific cell in a Heartbeat compression carrier ages at 35 °C and 75 % SOC, it cannot tell a clean sweep from a disturbed one, and it natively has no limits on what it can do.

Think of the "AI harness" as a harness on a wild horse (AI Model). To control this wild horse, it is a smart multi-layer harness that keeps the LLM in the direction and with the limited authority the software designers' permit. The skill is in the design and implementation of the harness, not the raw LLM model, its genius belongs to others.

Quasar™ family consists of Quasar Glass, Quasar ONE, Quasar Master and Quasar BMS

Safier has learnt, in the use of agents, the value of a second agent verifying alerts and warnings without knowledge of the logic the first agent used. With Quasar this is implemented in Quasar Twin.

LLMs are released at a blistering pace. The Quasar AI harness is agnostic to the model. A better model next month can be accessed and used easily after a short validation process. Quasar is not beholden to the strength of proprietary LLM models. Open source models with weightings specific to this application are the most likely end game for high volume use.

The structure, the signature learnings, the rules, the context and the gates are the proprietary knowhow to Safiery, going forward. The structure is described here. The signature data, thresholds, weightings and instruction sets remain Safiery's proprietary knowhow.


The harness supplies everything the model lacks:

- **Measurements.** Validated EIS parameters for each 3P cell group, with SOC %, temperatures and node timestamps.
- **Knowledge.** Four signature sets loaded by Safiery (section 4): failure modes, condition variants, ageing and compression.
- **Rules.** The scan validity logic, and the persistence rules that decide what may be reported.
- **Context.** Each group's signed factory baseline and full history in the Quasar Twin.
- **Gates.** A second, independent agent checks any warning before it is issued. Local protection in the BMS never waits for either agent.
- **Record.** Every assessment is written back to the twin as a data node, so it can be reviewed, audited and replayed.

The harness also chooses how the model is instructed. A routine daily sweep, a failed validity check, a local warning raised by the BMS and a change in the rate of drift each call for a different instruction set, and the harness selects it from the event.

Because the model is a tenant, it can be replaced or upgraded without changing the safety case. The signatures, the validity rules and the gates are the durable asset.

Why the BMS belongs in the family. A harness is only as good as its first measurement. The BMS is where that measurement is made and first judged, so it is the edge of the harness, not a data source bolted on to it. The same logic runs on the battery whether or not the rest of Quasar is connected.



3. The measurement chain, end to end




Figure 1. The chain from sweep to assessment. Steps 1 to 4 run in the battery and need no network. Steps 5 to 9 run when Quasar ONE is connected.

3.1 At the battery: an isolated EIS function

- The EIS function is electrically isolated from the rest of the BMS. A fault in the EIS path cannot compromise cell monitoring, balancing or protection.
- On initiation, the BMS injects a sine-wave current of about 5 A on the 48 V pack and steps its frequency from 0.1 Hz to 976 Hz. The voltage response of each 3P cell group is a few millivolts, held under 10 mV so the cell stays in its linear range.
- The NXP BMA7418 measures each of the 14 cell groups with 10 $\mu\Omega$ resolution, using synchronous (lock-in) detection at the excitation frequency. One pack scan takes about 55 seconds.
- The sweep runs in the quiet window: the pack is rested from the bank and the whole-battery shunt reads under 6 mA, so the 5 A sine is the only AC current in the pack. In a bank, packs are scanned in turn while the others carry the load.

3.2 Is the scan clean? The validity check in the BMS

The BMS decides whether each sweep is valid before anything uses it. A sweep is accepted only when all four tests pass:

- The pack current stayed under the 6 mA shunt threshold for the whole sweep.
- The real and imaginary parts of the impedance are consistent with each other under the Kramers-Kronig relations.
- The equivalent-circuit fit (R_0 , the two arcs and the Warburg tail) has small residuals.
- The result repeats within the group's known noise floor.

The Kramers-Kronig test. For a system that is linear, causal, stable and unchanged during the measurement, the real and imaginary parts of the impedance are linked: either can be calculated from the other. In practice the BMS fits the sweep with a chain of simple RC elements (the Lin-KK method), which can only fit data that obeys the relations, and checks the residuals at each frequency. If they are large, the cell drifted during the sweep, the excitation was too large, or interference got into the measurement.

If the sweep fails, the BMS runs it again. Only valid sweeps are used for local warnings or sent to the twin. Where failures cluster is useful in itself: repeated low-frequency failures on one pack usually mean its quiet window is not truly quiet, not that its cells are changing.

3.3 Immediate warnings, decided at the battery

Some failures cannot wait for the cloud. The BMS compares each valid sweep with the group's signed baseline and its recent sweeps, and acts on four imminent patterns without any network connection.

Warning	EIS pattern at the battery	Confirming signal	Why it cannot wait	Local action
Lithium plating	R_{ct} falls suddenly and the arc 2 apex frequency rises, at matched temperature	A cold or high-rate charge in the charge log before the sweep	Plated lithium can grow into dendrites and an internal short; further charging makes it worse	Restrict charge current and block charging below the temperature limit; Protect if it persists
Soft short or dendrite	R_{ct} down with erratic dips in R_0 on one group	Balancing demand and self-discharge rising on that group	An internal short can progress to a hard short and heating	Stop charging the pack, rest it from the bank, raise an alert
Open cell or failed weld	A step of about +50 % in group resistance in a single sweep	A tab thermistor rise on that group	The two remaining cells each carry 50 % more current, and a failing joint heats	Derate the pack by at least a third or rest it from the bank; raise an alert
Rapid gas generation (pore blockage)	R_{ct} and the Warburg tail rising in steps, far faster than the ageing curve	Recent high voltage and temperature; swelling, if instrumented	Gas generation is the early stage of swelling and venting	Stop charging, rest the pack, raise an alert

An alarming result triggers a confirmation sweep shortly afterwards, as in the current BMS logic. For plating, the charge restriction can be applied while the confirmation runs, because the cost of waiting is far higher than the cost of a restricted charge.

These warnings work with or without Quasar ONE. A HiBlade pack on its own is protected; Quasar adds the installation-wide view and the slower diagnoses.

3.4 From the battery to the twin: Quasar ONE, NATS and JetStream


When a Quasar ONE module is connected, it publishes the result of every valid sweep for every cell group in every pack:

- the fitted values per cell group: R_0 , R_{SEI} , R_{ct} , the Warburg slope σ , the arc apex frequencies and the CPE exponents;
- the pack's SOC % and the group temperatures at the time of the sweep;
- the validity result, and any local warning the BMS raised.

The values go to each battery's digital twin in the Quasar Twin over NATS with JetStream. **Every value is timestamped at the node where it was measured, not when it arrives.** That gives three things:

- exact ordering across 140 groups and 10 packs, whatever the network latency;
- a temperature reading lined up exactly with the sweep it is used to correct;
- values measured while the link was down keep their true time when they are forwarded later.

vessel SN0847 : 10 HiBlade packs : 14 cell groups per pack : battery 01. group 07 expanded : values are illustrative



A number on its own means nothing to an agent. "488" could be a resistance, a voltage or a count, from any battery on any boat. The semantic model gives every value a single, permanent place in one tree: vessel, system, bank, battery, cell group, facet and leaf. The address is the meaning:

That one line says which boat, which bank, which battery, which cell group, which measurement and which unit. Nothing has to be guessed, looked up in a mapping table or inferred by a model. This matters for six reasons.

2. Assessments sit beside the data they judge. The harness writes its findings back into the same tree, as assessment leaves on the cell group they concern. The status, likely modes, warning period and validation action for group_07 live next to group_07's impedance values and temperatures. Anyone reviewing a warning, whether an owner, an engineer or an auditor, finds the evidence and the judgement in one place.

Question	Subject query
How does group_07 compare with its 13 siblings?	vessel.SN0847.electrical.house_bank.battery_01.*.eis.r_ct_uohm
How does battery_01 compare with the other 9 packs?	vessel.SN0847.electrical.house_bank.*.pack.temperature_c
Which cell groups are on watch across the fleet?	vessel.*.electrical.house_bank.*.assessment.status

· Safiery Pty Ltd · Page 5

4. Time is exact and history is kept. Each value is timestamped at the node where it was measured, and JetStream keeps the history of every subject. The harness can replay any cell group from its signed baseline to today. A battery that was offline, like battery_10 in the figure, rejoins with its readings in their true place in time rather than in the order they arrived.

5. Authority is scoped by the subject. Credentials in Quasar are granted per subject, which makes the semantic model part of the safety case. The harness can read eis, signals and bms leaves, and can write only assessment leaves. It has no path to write a BMS setting or a protection threshold. The checking agent reads the same subjects and the proposed assessment, but never the first agent's reasoning.

6. The model is the stable interface. LLMs will keep changing; the semantic model does not. A new model is given the same subjects, the same signatures and the same rules, so replacing it does not mean re-integrating any battery, boat or display. The Quasar Twin, the Quasar ONE app and the agents all read the same leaves, so an owner, a technician and an AI agent are always looking at the same facts.

In short, the signatures tell the harness what to look for, and the semantic model tells it exactly where to look. Without it, fleet-level EIS reading would be a custom integration for every installation. With it, adding a boat means adding one trunk to the tree.

3.6 In the Quasar Twin: an assessment node for each cell group

The harness runs the AI agent over the new data, with each group's history as the reference, and writes the result back to the twin as a data node attached to that cell group.

Assessment field	What it contains
Status	Normal ageing, watch, or warning.
Likely modes	A ranked shortlist of the most likely failure modes, each with its evidence.
Commentary	Plain language: what moved, by how much, against which reference, and which corroborating signal supports it.
Warning period	The likely time to failure, as remaining margin divided by the fitted rate, with the basis of the estimate.
Recommended validation action	The step that would confirm or rule out the leading mode, where it can be done: for example an EIS sweep twice a day instead of once, a sweep at a set SOC, or inspection of a named joint.
Evidence	Links to the sweeps, the baseline and the signatures matched.
Independent check	The result of the second agent's review, whenever a warning is proposed.

Validation actions close the loop back to the battery. Moving one pack from one sweep a day to two doubles its data, from about 365 to about 730 sweeps a year per cell group, and shortens the time to a confident call on that pack alone.

3.7 A second agent before any warning

When the first agent proposes a warning, a second, independent agent reviews it. It sees the data and the proposed assessment, but not the first agent's reasoning. If they disagree, the finding is held at "watch" and validation action is requested instead. This reduces the chance of a false alarm reaching an owner, and of one model's blind spot passing unchallenged.

4. What Safieri loads into the harness

Safieri loads the harness with four sets of learned signatures, together with the scan validity logic. The failure-mode set names the mechanism; the other three define the reference it is measured against.

Set	What it holds	What it lets the harness do
Validity logic	Quiet-window rule, Kramers-Kronig test, fit-residual limits, repeatability, and the repeat-sweep rule	Trust only clean sweeps, and recognise a pack whose measurement conditions are poor
1. Failure-mode signatures	The fifteen modes, each as a pattern in level, rate and shape plus one corroborating non-EIS signal	Name the likely mechanism behind a change, and say what would tell the candidates apart
2. Condition variants	How each parameter and each signature looks at several SOC % levels and temperatures	Compare at matched conditions, so a warm day does not make a group look younger
3. Ageing signatures	Expected trajectories of normal calendar and cycle ageing, calibrated against the cell maker's cycle-life data	Separate normal ageing from a fault, and estimate the margin left
4. Compression signatures	How the impedance, especially R_0 and the contact arc, responds to stack compression pressure in the Heartbeat carrier	Recognise compression loss, separate reversible contact effects from delamination, and flag a carrier problem

A failure signature is a change against a reference. Sets 2, 3 and 4 give that reference at the right temperature, state of charge, age and compression, which is what makes set 1 usable in the field. The sets improve as field data accumulates across installations, and they stay with Safieri whichever model reads them.

5. Why reading the whole installation works

Each BMS sees 14 groups and its own data. The harness sees all 140 groups, all 10 packs and every non-EIS signal beside them, which turns the three hardest parts of EIS diagnosis (matched conditions, sibling comparison and single-cell dilution) into things it can measure.

5.1 The installation in numbers

Quantity	Per pack	Installation (10 packs)
3P cell groups	14	140
Group sweeps per day (one sweep a day)	14	140
Group sweeps per year	about 5,100	about 51,000
Fitted EIS values per year (about 8 per sweep)	about 41,000	about 410,000
Scan time, packs scanned in turn	about 55 s	about 9 to 10 min

Because the whole bank is scanned within about ten minutes, all 140 groups are measured under nearly the same ambient conditions, which is what makes comparison across packs fair.

Imagine a thousand boats with ten batteries being assessed, that is 410 million fitted values to assess a year. At that scale no team of engineers can read the results. Agents do the reading, and people receive only the checked findings.

5.2 A signature needs two kinds of evidence

Each signature is an EIS pattern confirmed by a signal that is not EIS. The harness holds both for every group, with the history needed to see when each started to move.

Mode	EIS pattern	Corroborating signal	What the harness correlates
Lithium plating	R_{ct} drops, apex frequency rises	Cold or high-rate charge in the charge log	The drop against the time and temperature of the preceding charge
Weld or busbar	R_0 rigid shift, steps or scatter	One joint runs warm	Which group moved against which thermistor rose; one group or the pack terminal
Soft short	R_{ct} down, erratic R_0 dips	Balancing demand rising on one group	Balancing effort per group over weeks against the impedance trend

5.3 Matched conditions (signature set 2)

Temperature moves R_{ct} by more than the single-cell faults the system is looking for. If R_{ct} roughly halves for every 10 °C of warming, the correction is $R_{ct}(T_{ref}) = R_{ct}(T) \times 2^{((T - T_{ref})/10)}$.

Temperature error	Apparent R_{ct} change	Compared with one cell's R_{ct} doubling (+20 % at group level)
1 °C	about 7 %	a third of the signal
3 °C	about 19 %	the same size as the signal
5 °C	about 29 %	larger than the signal

The halving rule is an approximation; the real slope depends on chemistry, SOC and age. Signature set 2 gives the harness the measured variants, and 140 groups at slightly different temperatures and SOC let it refine the correction for each installation as the cells age. Around 51,000 sweeps a year also give a measured noise floor for every parameter, so "significant" is a number, not a guess.

5.4 Three levels of comparison

The BMS can compare a group with its 13 siblings. The harness adds two more levels: each pack against the other nine, and the whole installation against the expected ageing curve (signature set 3). Where the drift appears points to the kind of cause.

Drift pattern	Compared against	Most likely points to	Example modes
One group moves, 13 siblings agree	The other groups in its pack	A local fault in that group, a joint or one cell	Weld or busbar, plating on one cell, soft short, delamination, pore blockage
All 14 groups in one pack move, the other 9 packs agree	The other packs in the bank	Something specific to that pack: position, duty, terminal or compression	Pack in a hot spot, a pack carrying more current, terminal or cable, compression loss, dryout
All 140 groups move together	The expected ageing curve and the cell maker's data	An installation-wide cause: climate, charge regime or normal ageing	Calendar ageing, SEI growth, charge held at high SOC

Each level cancels the common cause above it, in the same way common-mode rejection works in instrumentation, applied twice. With 140 groups and 10 pack medians the harness can use robust statistics, such as the median and median absolute deviation, and flag an outlier with a known false-alarm rate.

5.5 Single cells seen through 3P dilution

The measurement is per 3P group, so a change in one cell is diluted by its two parallel siblings. The dilution follows fixed ratios, which the harness uses as fingerprints.

Single-cell event	Group value (three cells of resistance R)	Change at group level	Size at the illustrative baseline
One cell's R_{ct} doubles	0.40 R instead of 0.33 R	+20 % on R_{ct}	about +90 $\mu\Omega$ on 0.45 m Ω , 9 × the 10 $\mu\Omega$ floor
One cell goes open	0.50 R instead of 0.33 R	+50 %, as a step	about +225 $\mu\Omega$ on an R_0 of 0.45 m Ω , over 20 × the floor
Plating on one cell	lower R_{ct}	R_{ct} falls	the opposite direction to every ageing mode

The baseline figures are the illustrative ones from the reading guide, not measured HiBlade values. A 20 % rise could be one cell doubling or all three rising by 20 %; the first tends to be sudden, local and visible in the tab thermistor or balancing demand, the second gradual and shared with the pack. Single-cell events sit only 9 to 20 times above the floor, so the matched-condition correction in 5.3 is what keeps them visible.

5.6 Level, rate and shape (signature set 3)

A threshold alarm asks one question: is this number too high? The signatures ask three: how far, how fast, and in what pattern. The harness fits each group's history to the growth shapes the modes produce and reports a ranked shortlist.

Growth shape over time	Typical modes
Rises with the square root of time	SEI growth
Steady linear creep, $\mu\Omega$ per week	Current-collector corrosion
Accelerating	Transition-metal dissolution
Sudden steps	Weld failure, pore blockage from gas, open cell
Rises and partly recovers with temperature cycles	Loss of stack compression
Steps that follow hard cycling rather than calendar time	NCM microcracking

A finding must persist over several sweeps and survive the sibling comparison before it is reported. Plating is the exception, because the cost of waiting is lopsided.

5.7 Compression (signature set 4)

Loss of stack compression is one of the fastest modes, with 2 to 8 weeks of warning. It raises R_0 , adds a contact arc at high frequency, and partly recovers through temperature cycles. Delamination looks similar but never recovers. Set 4 lets the harness tell the two apart, and when all 14 groups of one pack move together in the compression pattern, it points to the carrier rather than to the cells.

6. Safety Protection Boundaries

The agents are advisory. Protection stays in the Quasar BMS and its hardware, and every claim stays as firm as the data allows.

- **Not in the safety loop.** Overvoltage, overtemperature, short-circuit protection, the local warnings in 3.3 and the Protect response all run in the Quasar and its hardware, whether or not Quasar ONE or the cloud is connected.
- **Only valid sweeps.** Nothing that fails the validity check is trended or assessed.
- **Not what EIS cannot see.** A high-resistance soft short shows first in self-discharge. A weld that fails suddenly is flagged as a step, not predicted. A thermal event in progress belongs to temperature, voltage and pressure monitoring.

Quasar™ family consists of Quasar Glass, Quasar ONE, Quasar Master and Quasar BMS

- **Warning periods are estimates.** The per-mode windows come from published ageing rates and the detection floor, not yet from HiBlade field ageing data, and are reported with that framing.
- **Findings need persistence and a second check.** Apart from plating, nothing becomes a warning on a single sweep or on one agent's view.

7. Subscriptions: you are hiring agents to check State of health daily

Quasar ONE is a low cost US\$500 on boat device complete with its own modem and eSIM. In its first year, the subscription service is included so the user gets first hand experience of its value.

The second and subsequent years are a paid subscription. For the “Battery Detailed State of Health” assessment and notification, the subscription price is a function of the number of batteries.

There are 3 benefits to a subscription:

1. For a continuous subscription, the warranty period is increased from 3 years to 5 years.
2. The visibility of continuous State of Health Measurement by cell group, by battery, visible in the Quasar Twin.
3. At resale, a persistent record of the state of health of the batteries and an estimate of the ageing life to 80% capacity (same benchmark as most lithium batteries)

If the subscription lapses, the BMS validity check, local warnings and Protect response continue unchanged.

8. IP Notice

There are 8 patents filed on the Quasar BMS, and the compression carrier of the HiBlade battery.

There are 5 patents filed on the Digital twin creation and onboarding process plus the AI harness structure.

Copyright Safiery Pty Ltd 2026