ADDITIVE ASSURANCE

WHITE PAPER

PROOF & TRUST IN METAL AM

Validation and Compliance without Slowing Production

QUALITY ASSURANCE SERIES

1. EXECUTIVE SUMMARY

Metal AM is not short of measurement; it's short of proof. Proof that survives scrutiny, scales across machines and sites, and doesn't collapse into endless post-build inspection or a "monitoring dashboard" nobody can sign off.

The first phase of any serious quality assurance journey is therefore not about more data. It's about trust — the ability for Quality, Operations, Certification/Regulatory, and IT/Security to agree early: what evidence is required, how it will be generated, how it will be retained, and how it will be defended.

THIS WHITE PAPER ADDRESSES:

- Why many in-situ initiatives stall at the governance boundary.
- What "first value" should look like in the first month after installation.
- Why pilots should be framed as evidence programmes that move at the pace of risk acceptance.
- How "golden build" similarity thinking provides proof of control and drift detection.

WRITTEN FOR: TEAMS SPECIFYING A SOLUTION TO VALIDATE AM BUILDS — ESPECIALLY THOSE WHO NEED RESULTS THAT ARE OPERATIONALLY USABLE, GOVERNANCE-READY, AND COMPLIANCE-CREDIBLE.

2. THE UNCOMFORTABLE TRUTH

Most QA tools don't fail technically — they fail organisationally. In regulated manufacturing, projects rarely die because a sensor can't see something. They die because nobody translated "monitoring" into governance. That governance boundary — between Quality and IT/Security — is where good intentions go to stall.

The questions that block progress are rarely "Does it detect anomalies?" They're the unglamorous ones:

- Where does it run — cloud, on-prem, air-gapped?
- Who controls the data? Who can access it? Who can export it?
- What is retained, for how long, and at what level of detail?
- What does the report look like when an auditor asks for evidence?

> **When these questions aren't answered early, "monitoring" stays a science project. When they are answered, in-situ evidence becomes something a Quality leader can actually approve.**

> **IN REGULATED MANUFACTURING, PROJECTS RARELY DIE BECAUSE A SENSOR CAN'T SEE SOMETHING. THEY DIE BECAUSE NOBODY TRANSLATED 'MONITORING' INTO GOVERNANCE.**

3. EVIDENCE, NOT HEATMAPS

A mature quality function doesn't buy dashboards; it buys decisions it can defend. There is a key distinction at the heart of everything: **monitoring versus assurance**.

Monitoring shows you information. Assurance produces decision-ready outputs — clear reporting, traceability, and evidence packaging that can be reviewed quickly, documented consistently, and defended later.

That seemingly subtle difference changes everything. It widens who can realistically consume the outputs (not just PhDs), it reframes how a pilot should be run (as an evidence programme rather than a technology trial), it shifts the deployment conversation to the issues that actually decide adoption (governance first), and it redefines success in practical terms (behaviour change and decision-making discipline — not simply prettier plots).

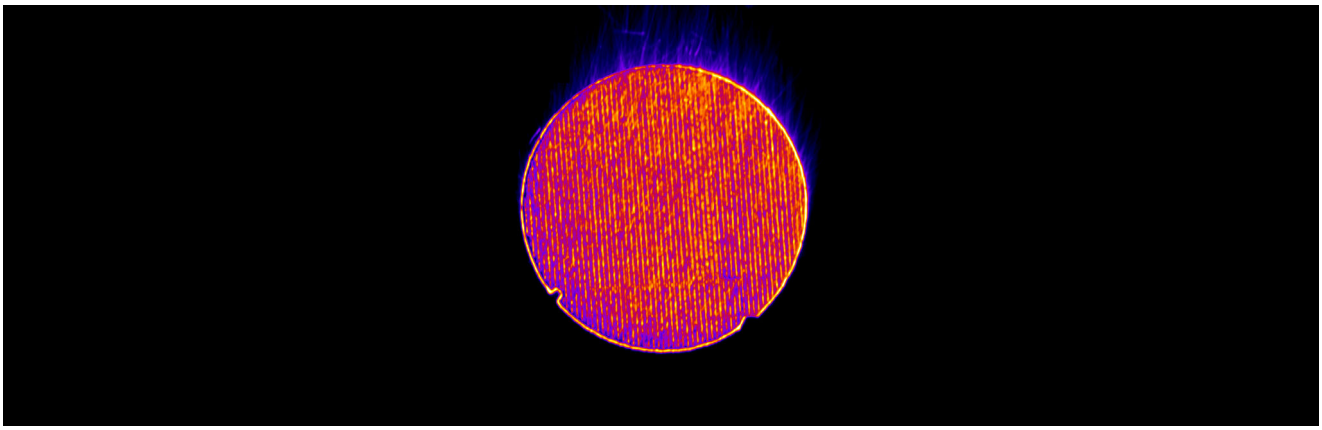

Layer 569 melt-pool emissions colourmap captured by AMiRIS. Raw sensor data alone is not evidence — it requires structured interpretation to become decision-ready output.

In practice, decision-ready evidence has three characteristics:

COMPACT	STRUCTURED	CONTEXTUAL
Compact enough to be used. A Quality reviewer cannot audit "all the pixels."	Structured enough to be stored and retrieved — not a folder full of screenshots.	Tied to thresholds so "proceed / hold / audit" isn't an opinion.

The most advanced analytics won't help if the evidence doesn't land in a form a regulated organisation can sign.

4. DEPLOYMENT MODELS DEMYSTIFIED

A surprising number of in-situ projects stall because deployment is treated as an "IT detail" to fix later. Deployment is a buying criterion — not because one model is always better, but because each model shifts **who owns risk**.

CLOUD DEPLOYMENT
Cloud is often the shortest path to evidence when policy allows it. It reduces internal IT workload and can make scaling and updates simpler. The real question is not "Is cloud secure?" — it becomes: "Can your policies, contracts, customer requirements, and audit posture support the required data handling?"

ON-PREMISES
On-prem means the system runs inside your network and storage boundary, under your controls. It's often essential where data residency, supplier access constraints, or customer requirements are non-negotiable. The trade-off is ownership: patching, infrastructure planning, and governance become your responsibility.

AIR-GAPPED
Air-gapped is simply on-prem with isolation — no dependence on internet connectivity. It removes an entire category of external connectivity risk and can simplify conversations with customers who mandate isolation.

AMiRIS sensor unit installed at the SLM machine's build chamber window.

THE 4 GOVERNANCE DECISIONS

Whatever model you choose, success is determined by four governance decisions that IT/Security and Quality can sign together:

1. **DATA VOLUME PLANNING:** Generates ~30 MB per cm³ of build volume. Needs planning, not surprise.
2. **RETENTION POLICY:** Define retention by audit need, not habit. What must be retained vs archived.
3. **ACCESS CONTROL:** Who can view, export, administer — minimise and audit access.
4. **REPORTING OUTPUTS:** Quality needs decision-ready reporting and traceability, not "all the pixels."

5. WHAT "WORKING" LOOKS LIKE

Most organisations measure early success with the wrong metric: "Is it installed?" or "Does it show anomalies?" The more operational lens is: **does it change day-to-day decision-making within 30 days?** If not, it's usually an implementation problem, not a technology problem.

You don't need to perfect the model in week one. What you do need to do is to perfect the rules of engagement.

WEEK 1
CLEAR ACTIONS

The first week should not be spent chasing perfect analytics. Create shared rules: who reviews what, an action ladder (inform → review → hold), and what "good" looks like operationally.

WEEK 2
OUTPUTS REPLACE THEATRE

Outputs should be short, clear, defensible summaries that reduce cognitive load. A consistent review cadence, simple disposition language, and a repeatable "what changed?" narrative.

WEEK 3
GOVERNANCE UNBLOCKED

Access control, retention policy, and storage planning are locked down. "IT says no" stops being a default objection because sizing and controls are already done.

WEEK 4
REPEATABLE ROUTINE

First value is visible as reduced friction. Fewer late surprises, faster internal alignment, less time wasted interpreting anomalies. Economics become clear.

6. PILOTS AS EVIDENCE PROGRAMMES

Regulated teams have a healthy allergy to fixed "transform in 90 days" promises — for good reason. We reframe the pilot as an evidence programme: first evidence in around 30 days, then proceed at the pace your risk acceptance, customer requirements, and certification reality demand. That isn't failure. That's appropriate governance.

This reframing matters because it stops you promising what you can't control (approval cycles) and focuses you on what you *can* control: producing a repeatable, decision-ready evidence pack.

PHASE 1 (0–30 DAYS)

DEPLOYMENT & FIRST EVIDENCE

The first month should feel operational, not experimental. The outcome is alignment around three questions: What decision are we trying to enable? What does "good evidence" look like here? What is the minimum viable report?

PHASE 2 (TIMELINE SET BY RISK)

VALIDATION DEPTH

Validation is a spectrum, determined by what you are trying to prove, and to whom. Expand validation only as far as your risk model requires, and keep the programme grounded in representative builds and agreed acceptance rules.

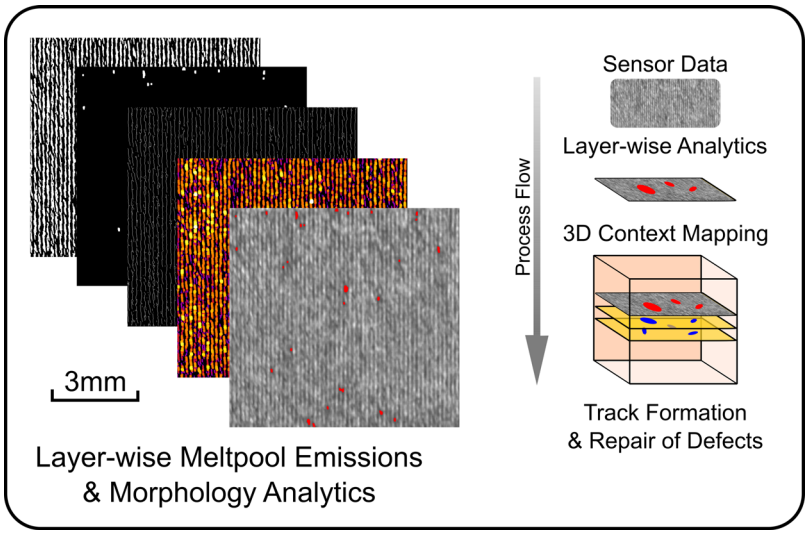

AMiRIS layer-wise meltpool emissions and morphology analytics — from raw per-layer sensor data to 3D context mapping and track formation analysis.

7. GOLDEN BUILD METHODOLOGY

Drift is rarely dramatic. It arrives as "slightly off." Then one day it's scrap, or a quality investigation with "containment" and "re-qualification" in the subject line.

Most organisations attempt to manage drift reactively. By the time you've found the issue, you've already made it. **Golden build methodology** is the grown-up approach: define a "known good" reference build and continuously compare new builds against it so drift is detected early — while the cost of acting is still low.

A golden build isn't a theoretical perfect part. It's a build you trust because it has been validated. The mistake is treating it as a trophy. The point is to use it as a control reference.

SIMILARITY BEATS DASHBOARDS

Dashboards make people argue. Similarity makes people decide. Similarity answers three questions that are genuinely useful:

- **Is this part behaving like the reference?**
- **Where is it different?**
- **Is the difference large enough to matter?**

DRIFT ACROSS CONTEXTS

Drift is often across time and context: machine A vs machine B, laser-to-laser variation, site-to-site differences, before/after maintenance, powder lot changes.

> **AMiRIS tracks across machines and sites with real-time alerts and 2D/3D visualisation to demonstrate process stability over time.**

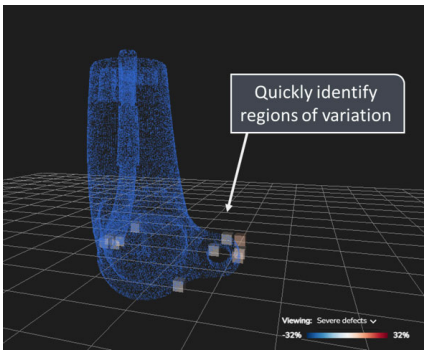

AMiRIS 3D defect viewer — quickly identify regions of variation across the full build volume.

8. CONCLUSION

If you're specifying a solution to validate AM builds, don't start by asking what the system can detect. Start by asking what your organisation needs to approve.

A. EVIDENCE & REPORTING

- What does a decision-ready report look like?
- Can it support disposition quickly and consistently?
- Can it be packaged as true traceability?

B. GOVERNANCE

- Where does the system run?
- Who controls the data, access, and export?
- What is your defensible retention policy?

C. OPERATIONAL ADOPTION

- What changes in the first 30 days?
- Who reviews what? What triggers escalation?
- How do you force decision discipline?

D. PROOF OF CONTROL

- How will you detect drift across time and sites?
- Do you have a "golden build" strategy?
- Does similarity thinking reduce noise?

> **If you're trying to scale metal AM, don't just ask "Can we detect defects?" Ask the more grown-up question: "Can we prove this process is under control without drowning in raw data?"**

Metal AM will not scale on inspection alone. The path to repeatable production is proof that can be approved — evidence that is decision-ready, governance-ready, and defensible under audit.

The central shift is moving from "monitoring" to assurance: compact, structured, contextual outputs that Quality can sign, Operations can act on, and IT/Security can deploy with clear controls.

When deployment is treated as a buying criterion, and when "working" is measured by behaviour change in the first month, teams can reach first evidence quickly and validate at the pace their risk model demands.

THE PRAGMATIC NEXT STEP

A JOINT ALIGNMENT SESSION (IT/SECURITY AND QUALITY) USING A TIGHT DEPLOYMENT AND GOVERNANCE CHECKLIST TO REMOVE FRICTION AND ACCELERATE CONFIDENT ADOPTION.