AURUS INSTITUTE FOR RESOURCE DEVELOPMENT

07

WHITE PAPER · NUMBER SEVEN

Data You Can Bank

Drilling, sampling and assay QA/QC as the foundation of resource confidence: from the raw sample to a grade a lender will underwrite

TECHNICAL PAPER · EXPLORE

July 2026 · ~50-min read

Aurus Technical Committee

AURUS INSTITUTE FOR RESOURCE DEVELOPMENT

Evidence before assertion.

MINING · INFRASTRUCTURE · ENGINEERING · ENVIRONMENT

Contents

The grade is a measurement, and a measurement without a control is an assertion

Every mineral resource estimate rests on a chain of physical steps: a sample is cut from the ground, reduced in stages, digested and read on an instrument, and the number that comes back is written into a database and, eventually, into a public report a bank or a stock exchange will rely on. Quality assurance and quality control, QA/QC, is the discipline that makes each link in that chain measurable, so that the reported grade carries a stated accuracy and precision rather than a hope. This paper treats QA/QC not as a laboratory formality but as the financial control that decides whether a resource is bankable.

It is written for two readers whose questions differ. A project **owner or exploration manager** needs to know what a defensible program actually inserts, measures and discloses, and where such programs most often fail. A **lender, investor or Qualified Person reviewing the work of others** needs a precise account of what the reporting codes require and a vocabulary for telling a controlled dataset from an uncontrolled one. The paper serves both by reading the guidelines and codes directly, and then setting out the execution discipline that turns them into a program that holds.

The obligation is explicit, not optional. The CIM Mineral Exploration Best Practice Guidelines require that a quality-assurance program be systematic, that it apply to each drilling and sampling campaign and to all types of analytical data, and that it cover the full range of values measured and not just the high or unusual results. The program must include external blanks, certified reference materials and duplicate samples, plus regular check sampling by a third-party laboratory, with blanks and reference materials inserted frequently enough to give statistical confidence in the results. Complete disclosure of the QA/QC results, the pass and fail criteria, and the actions taken on failures is recommended at every stage.

The reporting codes make the same demand as a disclosure test. The JORC Code, 2012 Edition, requires a public report to address, on an "if not, why not" basis, the "nature of quality control procedures adopted (eg standards, blanks, duplicates, external laboratory checks) and whether acceptable levels of accuracy (ie lack of bias) and precision have been established". In the CRIRSCO family of codes (JORC 2012, SAMREC and PERC) and the CIM Definition Standards that NI 43-101 incorporates, a resource is only as strong as the data that can be audited back to the raw record. A Qualified or Competent Person signs against that audit trail, not against the number alone.

> QA/QC · quality assurance is the set of planned, systematic activities that build confidence a process will meet its data-quality objectives; quality control is the measurement and testing, the reference materials, blanks and duplicates, that check whether it did. Assurance designs the control; control detects the failure. The two words are not interchangeable, and a program that has one without the other is incomplete.

What a controlled program delivers: five findings

- **The error is set before the laboratory opens the bag.** The largest single source of uncertainty in the whole chain is usually the initial sampling step, not the assay. A program that pours its attention into the laboratory and neglects drill-sample recovery and sub-sampling is controlling the wrong end (Chapter 2).
- **Three insertions answer three different questions.** Certified reference materials test accuracy and bias, blanks test contamination, and duplicates test precision. Each is a distinct question, and a scheme that inserts one heavily and the others sparingly leaves a real risk unmeasured (Chapters 3 to 5).

- **Failures are defined in advance or they are argued after the fact.** Control limits, warning at two standard deviations and action at three, and a written batch-failure rule are set before the data arrives, so a failing batch triggers a re-assay rather than a negotiation (Chapter 5).
- **Most errors enter through the keyboard, not the crusher.** Transcription, unit and merge errors in the database are a common and under-controlled failure, which is why the guidelines prefer a secured relational database over spreadsheets and require a data-verification program (Chapter 6).
- **The chain of custody is part of the assay.** A number is only defensible if the sample that produced it can be traced, secured and shown to be the sample it claims to be, from the core tray to the ISO-accredited laboratory (Chapters 2 and 6).

Who should read this paper

For **owners and exploration managers**, Chapters 1 through 6 set out what a defensible program inserts, measures and stores, and Chapters 7 and 8 give a disclosure test and a method to hold a program against. For **lenders, investors and reviewing Qualified Persons**, the same chapters provide a precise account of what the codes require and a diagnostic for separating a controlled dataset from an uncontrolled one. The paper cites the primary instruments directly: the CIM Mineral Exploration Best Practice Guidelines, the JORC Code 2012 Table 1, the recommended-practice review published in Geochemistry: Exploration, Environment, Analysis, and ISO/IEC 17025.

How to use this paper

Read **front to back** for the full method, from why QA/QC is a financial control to the closing standard. Read **Chapter 8 first** for the instrument, a template protocol with its KPIs and non-conformance workflow, then trace each element back to the guideline or code that earns it. Or read **by exhibit**: the exhibits and the five stat tiles, each with its source line, carry the paper's framework content, and the exhibit index in the back matter maps every one to its references. The signature exhibit, the control-chart triptych in Chapter 5 that reads precision, accuracy and contamination side by side, is the paper's centrepiece and, like the closing protocol, is a framework instrument built on illustrative values rather than any specific project's data.

The through-line, stated plainly

The argument reduces to one proposition: a grade is a measurement, and a measurement is only worth what its controls can prove. Two projects can report the same number, and one can be bankable while the other is not, because one carries a disclosed and passing QA/QC record and the other carries an assertion. The reporting codes exist to make that difference visible, and a disciplined team builds the control in from the first drill hole rather than reconstructing it under diligence, when the reference materials were never inserted and the duplicates were never taken.

Everything downstream follows from taking that seriously. The insertion scheme has to be designed before the program starts, because a certified reference material cannot be inserted retroactively into a batch that has already been read. The chain of custody has to be recorded as the samples move, because it cannot be reconstructed once they are gone. And the database has to be verified as it is built, because a transcription error found at resource estimation is far more expensive than one caught at data entry. The chapters that follow take each of these in turn; this summary is the map, not the territory.

A final orientation. This is a technical paper, not a laboratory manual and not a certification of any specific dataset; it reads the guidelines and codes that govern the work and sets out how a disciplined team executes against them. Its claims are framework claims, cited to published texts. A small number of numeric conventions, insertion rates and control limits among them, are widely used industry practice rather than a single code's mandated figure; the paper flags each of those as established practice and says where the binding requirement is only qualitative, because a reader relying on this for diligence is entitled to know exactly how firm each claim is.

AT A GLANCE

Five figures that frame a defensible program

INSERTION

1 in 20

COMMON INSERTION RATE FOR EACH QC TYPE (ABOUT 5%), SO QC MATERIAL RUNS NEAR 15 TO 20 PERCENT OF A BATCH

Recommended practice (Abzalov, 2011); qualitative basis CIM, Mineral Exploration Best Practice Guidelines, 2018

ACTION LIMIT

±3 SD

STANDARD DEVIATIONS FROM A REFERENCE MATERIAL'S CERTIFIED MEAN BEYOND WHICH A BATCH IS FAILED AND RE-ASSAYED

Recommended practice after Shewhart control charts; Smee, Bloom, Arne and Heberlein, GEEA, 2024

3

QUESTIONS, THREE INSERTIONS: REFERENCE MATERIALS TEST ACCURACY, BLANKS TEST CONTAMINATION, DUPLICATES TEST PRECISION

Smee et al., GEEA, 2024; CIM, 2018

ACCREDITATION

17025

ISO/IEC 17025:2017, THE COMPETENCE STANDARD BEHIND THE CIM PREFERENCE FOR AN ACCREDITED LABORATORY

ISO/IEC 17025:2017, General requirements for the competence of testing and calibration laboratories

DISCLOSURE

2012

JORC CODE EDITION WHOSE TABLE 1 REQUIRES A STATEMENT ON STANDARDS, BLANKS, DUPLICATES AND EXTERNAL CHECKS

JORC Code, 2012 Edition, Table 1, Section 1

HOW TO READ THIS PAGE

Two of these figures, the 1-in-20 insertion rate and the three-standard-deviation action limit, are widely used industry practice rather than a single code's mandated number; the binding guideline states the rule qualitatively and leaves the exact figure to program design. The paper flags every such figure as established practice where it appears.

How this paper is organised

CHAPTER	THE QUESTION IT ANSWERS	ANCHOR SOURCE
1 · QA/QC as a financial control	Why is a resource only as bankable as its data?	CIM, 2018; JORC, 2012
2 · Drilling and sample quality	Where in the chain is the largest error actually set?	CIM, 2018; GEEA, 2024
3 · The insertion scheme	What is inserted, where, and how often?	CIM, 2018; practice
4 · Certified reference materials	How is accuracy and bias measured against a known truth?	GEEA, 2024
5 · Precision, accuracy, contamination	How is a control chart read, and a failure triaged?	GEEA, 2024; practice
6 · Databases and chain of custody	Where do most errors enter, and how are they stopped?	CIM, 2018; ISO 17025
7 · The regulatory lens	What does a code reviewer flag first?	JORC, 2012; NI 43-101
8 · The Aurus standard	How does a team hold a program to all of the above?	Chapters 1 to 7

What this paper is, and is not

It is a practitioner's reading of the guidelines and codes that govern sampling and assay quality control in mineral exploration, and a method for executing against them. It is not a laboratory operating manual, a certification of any dataset, or a record of any specific client project. Every code and guideline is cited to its own published text; a small number of numeric conventions are flagged as established practice where the binding source is only qualitative. The signature control-chart triptych and the closing protocol are framework instruments built on illustrative values; they carry no client data.

1

EXPLORE · THE AUDIT TRAIL

QA/QC as a financial control

A resource estimate is a financial instrument built from measurements. Quality control is what lets a lender treat the grade as a number rather than a claim.

---

systematic
THE CIM TEST: EVERY CAMPAIGN, ALL DATA, THE FULL RANGE OF VALUES

if not, why not
THE JORC TABLE 1 DISCLOSURE STANDARD FOR QC PROCEDURES

3
QC INSERTIONS, THREE SEPARATE RISKS MEASURED

A mineral resource is not found in the ground. It is estimated from a set of measurements, each taken from a small physical sample, and the estimate is only as trustworthy as those measurements can be shown to be. Quality assurance and quality control is the engineering that turns a pile of assay certificates into a number a bank will lend against and a securities regulator will let a company publish.

Consider what actually happens between the deposit and the disclosure. A drill hole produces core or chips. An interval is selected and a sample is cut. That sample is dried, crushed and split, often more than once, until a few tens of grams of fine powder remain. The powder is digested in acid or fused in a furnace, and the solution or bead is read on an instrument that returns a concentration. That concentration is transcribed into a database, joined to a hole and a depth, composited, and fed into a geostatistical model that produces tonnes and grade. At every step, material is lost, split, moved and recorded, and at every step an error can enter that no downstream sophistication can remove. The resource model is a precise arithmetic on numbers whose quality was fixed long before the model ran.

Why the codes treat data as the asset

The reporting codes are explicit that the data, not the estimate, is what a Competent or Qualified Person stands behind. The CIM Mineral Exploration Best Practice Guidelines require that a quality-assurance program be in place and that quality-control measures confirm and document the accuracy and precision of results received from the laboratory. The program, the guidelines state, should be systematic, should apply to each drilling and sampling campaign and to all types of analytical data, and should span the full range of values measured rather than only the high or unusual results. That last clause matters more than it looks: a program that inserts reference materials only near expected ore grades leaves the low-grade and waste population, the population that defines the resource boundary, uncontrolled.

The same expectation appears in the disclosure codes as a public test. The JORC Code, 2012 Edition, requires that a public report address, on an "if not, why not" basis, the "nature of quality control procedures adopted (eg standards, blanks, duplicates, external laboratory checks) and whether acceptable levels of accuracy (ie lack of bias) and precision have been established". The "if not, why not" principle is the sharp edge: a criterion that is not addressed must be explained, so silence about QA/QC is not a neutral omission but a disclosed deficiency. Across the CRIRSCO family of codes, JORC 2012, SAMREC and PERC, together with the CIM Definition Standards that Canada's NI 43-101 incorporates, the through-line is identical: a resource is only as defensible as the data that can be audited from the report back to the raw field record.

EXHIBIT 1 · CAPABILITY INSTRUMENT · NO PROJECT DATA

The chain from ground to report, and where control has to sit at each link

Every stage loses or transforms material; a control inserted after a stage cannot recover what that stage got wrong.

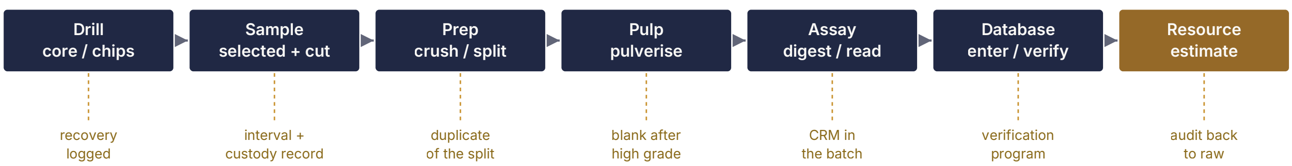

The reported resource sits at the far right; every control to its left is a chance to detect a problem while it is still cheap to fix. A missing control is a silent failure that only surfaces at diligence, when the sample that could have proven the grade is gone.

Structure after CIM, Mineral Exploration Best Practice Guidelines, 2018, §6 and §7; JORC Code, 2012 Edition, Table 1, Section 1. Framework schematic; illustrative, no project data.

The cost asymmetry that makes control pay

QA/QC looks like an expense until it is priced against what it prevents. The direct cost of inserting reference materials, blanks and duplicates, and of a modest program of umpire check assays at a second laboratory, is a small percentage on top of an analytical bill that is itself a small fraction of a drilling budget. The cost of an uncontrolled dataset is of a different order. If a bias is discovered at feasibility, the remedy is often a re-assay of pulp, if the pulp still exists, or a fresh drilling program, if it does not, together with the delay of a re-estimate and the credibility cost of a public restatement. The economics are asymmetric: control is a known small cost paid early, and its absence is an unknown large cost paid late, at the worst possible moment in a project's financing.

This is why the discipline belongs in the language of financial control rather than laboratory housekeeping. A control is a designed check that a process is doing what it claims, with a defined trigger when it is not. Inserting a certified reference material and plotting it against control limits is exactly that: a check on the accuracy of a batch, with a rule that fails the batch when the check is breached. Read that way, a QA/QC program is a system of internal controls over the most important number a resource company reports, and it should be designed, evidenced and audited with the seriousness that implies.

> **Two projects can report the same grade. One is bankable and one is not, and the difference is whether the number arrives with a control record or an assurance that everything was fine.**
>
> AURUS TECHNICAL COMMITTEE

Materiality: the reason a regulator cares

Securities regulators do not concern themselves with laboratory technique for its own sake. They concern themselves with materiality: whether information that could change an investor's decision has been disclosed fairly. QA/QC results are material in exactly that sense. A grade is a representation to the market, and the controls behind it are the evidence that the representation is sound. This is why the codes route QA/QC into public disclosure rather than leaving it in a technical annex, and why a Qualified or Competent Person, whose sign-off carries personal and professional liability, is required to be satisfied with it. A resource announced without disclosed controls is not merely technically thin; it is a representation the market cannot test, and a regulator reads an untestable representation as a risk to investors rather than a neutral omission.

The consequence for a program is that QA/QC has to be built to be disclosed, not merely to be done. Records that would satisfy a geologist that the data are fine are not the same as records that let an independent person confirm it: the second needs the control charts, the pass and fail criteria set in advance, the failures and the actions taken on them, and the chain that ties each result back to a sample and a certificate. A program that did good work but cannot evidence it in disclosable form has, for the purposes that matter, not controlled its data, because the market and its reviewers can only rely on what can be shown.

What the rest of this paper does

The chapters that follow build the program from the ground up. Chapter 2 starts where the error is largest, at the drill and in the sub-sample, because no laboratory control can rescue a sample that was never representative. Chapter 3 designs the insertion scheme, the what, where and how often of the reference materials, blanks and duplicates. Chapters 4 and 5 read the three insertions in turn and then together, in the control-chart triptych that is the paper's signature. Chapter 6 turns to the database and the chain of custody, the unglamorous half of data integrity where a large share of real errors enter. Chapter 7 puts the whole program under the reporting-code lens and names the red flags a reviewer looks for first. Chapter 8 assembles the result into a template protocol with its KPIs and a non-conformance workflow that closes the loop.

> Control is a small cost paid early against an unknown cost paid late. That is the whole economic case, and it is enough.

2

EXPLORE · WHERE THE ERROR IS SET

Drilling techniques and sample quality

The largest single error in the whole chain is usually the first one: the sample itself. Recovery and representativity decide what the laboratory can never fix.

---

recovery
LOGGED AS A CONTROL, NOT JUST A NUMBER IN THE LOG

RC vs core
TWO METHODS, TWO FAILURE MODES FOR SAMPLE QUALITY

sub-sample
THE SPLIT WHERE FUNDAMENTAL SAMPLING ERROR CONCENTRATES

It is tempting to think of quality control as a laboratory activity, because the laboratory is where the reference materials are read and the control charts are drawn. That instinct controls the wrong end of the chain. The recommended-practice literature is blunt that the largest single source of uncertainty in the whole process is usually the initial sampling, not the assay. A perfectly controlled laboratory reading an unrepresentative sample returns a precise, accurate answer to the wrong question.

Representativity: the property that cannot be added later

A sample is representative when its grade is an unbiased estimate of the grade of the material it was drawn from. Representativity is set physically, by how the sample was taken and reduced, and it cannot be restored by any later step. If drilling loses a disproportionate amount of the soft, high-grade fraction of a zone, the recovered sample is biased low, and every reference material and duplicate downstream will faithfully confirm the precision and accuracy of a biased number. This is the central reason drill-sample recovery is treated as a control in its own right, to be measured and logged interval by interval, rather than a passive observation.

The two dominant drilling methods carry different risks. Diamond core drilling recovers a continuous cylinder of rock that can be logged, photographed, measured for recovery against the drilled length, and split so that half is assayed and half retained as a physical record. Reverse-circulation, or RC, drilling returns chips through the rod string to a cyclone and a splitter, faster and cheaper per metre, but with sample quality that depends on a dry, well-managed return; wet RC samples, contamination between intervals, and splitter bias are its characteristic failure modes. Neither method is superior in the abstract. The control question is whether the method chosen is being run so that recovery is high and consistent and the sample is being split without bias, and whether that is being evidenced rather than assumed.

EXHIBIT 2 · CAPABILITY INSTRUMENT · NO PROJECT DATA

Two drilling methods, the sample they produce, and the control each demands

The right method is the one whose characteristic failure mode is being actively measured, not the one that is cheaper per metre.

ATTRIBUTE	DIAMOND CORE	REVERSE CIRCULATION (RC)
Sample form	Continuous cylinder of solid rock	Rock chips returned to a cyclone and splitter
Recovery control	Recovered length against drilled length, logged per run	Sample weight per metre against a theoretical mass, monitored for loss
Representativity risk	Core loss in soft or broken ground; grinding of the core	Wet returns, down-hole contamination, splitter bias
Physical record	Retained half core; a re-loggable, re-samplable archive	Chip trays and pulp reject; no re-drillable witness
Typical use	Resource-definition and geotechnical work; the reference method	Fast grade-control and early-stage drilling in suitable ground
The control question	Is recovery high, and is the half-core split unbiased?	Is the return dry, and is the splitter cutting a true fraction?

After CIM, Mineral Exploration Best Practice Guidelines, 2018, §6 and §7; JORC Code, 2012 Edition, Table 1, Section 1 (drilling technique; drill sample recovery; sub-sampling). Framework comparison; no project data.

Sub-sampling: the fundamental sampling error

Between the metres of core or the kilograms of RC chips and the tens of grams the instrument actually reads lies a series of mass reductions, and it is here that the fundamental sampling error concentrates. Each time a large mass is reduced to a smaller one, the smaller mass can only represent the larger if the material is crushed finely enough and split evenly enough that the fraction taken carries the same grade as the whole. Coarse gold is the textbook villain, because a single liberated grain in the wrong split moves the reported grade dramatically, but the principle is general: the finer and more variable the valuable mineral, the larger the sample mass that must be prepared and the finer it must be comminuted before it is split.

This is why sample preparation is a QA/QC concern and not a purely logistical one. The CIM guidelines require that where field-sample volume is reduced before shipping, the splitting procedures used to obtain representative sub-samples be tested, verified and then applied, and that preparation be appropriate to the material and the elements being determined. In practice the controls are a preparation duplicate, a second split taken at the coarse-reject or pulp stage to measure how much variance the preparation itself is adding, and a specification on crush and pulverise fineness, commonly expressed as a target percentage passing a stated sieve. A duplicate taken only at the pulp stage measures the analytical variance alone and is silent about the far larger error introduced upstream, which is the error that actually threatens the resource.

EXHIBIT 3 · CAPABILITY INSTRUMENT · NO PROJECT DATA

Where the variance lives: error accumulates, and the first stage dominates

Illustrative decomposition of total sampling-and-assay variance by stage. The shape, not the exact heights, is the point.

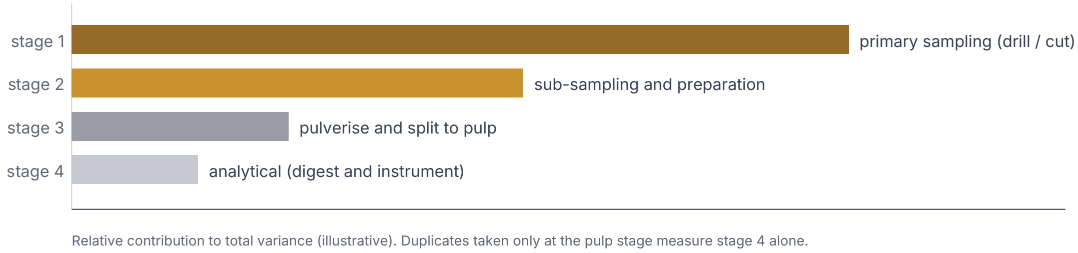

Illustrative, after the recommended-practice finding that the largest uncertainty arises at initial sampling (Smee, Bloom, Arne and Heberlein, GEEA, 2024). Schematic; not measured project data.

Recovery, and the grade it can hide

Recovery deserves a paragraph of its own because it is the control most often logged and least often interrogated. It is simple to record the recovered length of core against the drilled length, or the sample weight per metre of an RC hole against a theoretical mass, and most programs do. The discipline that separates a controlled program from a compliant one is examining whether recovery is related to grade. If the soft, weathered or fractured zones that drill poorly are also the zones that carry the metal, then low recovery is not random loss but a selective loss of grade, and the sampled population is biased in a direction no laboratory control will reveal. The JORC Table 1 asks specifically whether the relationship between recovery and grade has been considered and whether sample bias may have occurred, which is why recovery is a grade-integrity control and not a drilling statistic.

The remedy is designed at the program level. Where core loss is a risk, drilling technique is adjusted, triple-tube core barrels and shorter runs among the options, and recovery is monitored interval by interval so a developing problem is seen while the hole is still being drilled. Where an RC return turns wet, the log records it, because a wet sample split is a different and less reliable measurement than a dry one. None of this is exotic; it is the ordinary practice of treating the first measurement in the chain with the seriousness its position deserves, because it is the one measurement that every later number depends on and the one that cannot be repeated once the ground is drilled.

Chain of custody starts at the core tray

Sample security is not a step that happens after sampling; it begins the moment the sample exists. The CIM guidelines require that a geoscientist supervise sample collection and that a chain of custody be established and recorded, with secure logging, sampling, storage and preparation facilities and prompt, secure, direct shipping to the laboratory. The reason is both quality and integrity. A recorded custody trail lets a reviewer confirm that the pulp read as hole twelve at forty metres is that sample and not a mislabelled neighbour, and the QC insertions themselves double as a custody check, because a reference material that returns the wrong certified value can reveal a batch that was mixed up as readily as one that was mis-assayed.

Twinned holes and the independent check

One control sits outside the insertion scheme entirely and answers a question the reference materials cannot: is the sampling itself reproducible in the ground? A twinned hole is a second hole drilled close to an original, sometimes by a different method, so that two independent samples of nearly the same volume of rock can be compared. Where the twin reproduces the grade profile of the original, it gives confidence that the drilling and sampling are capturing the deposit faithfully; where it does not, it exposes a problem, poor recovery, a nugget effect, or a method bias, that no laboratory control would ever have shown, because the laboratory only ever saw one of the two samples. The JORC Table 1 names the use of twinned holes explicitly under verification of sampling and assaying, which places it among the disclosures a reviewer expects to see addressed.

Twinning is used sparingly, because a twin hole is an expensive control, and it is aimed where the stakes are highest: at an early conversion from RC to core, at a change of contractor or method, or where a resource depends heavily on a small number of high-grade intercepts. The point is not to twin everything but to have twinned enough, in the right places, that the reproducibility of the primary sampling is evidenced rather than assumed. A program that inserts reference materials diligently but never once checks that its drilling reproduces in the ground has controlled the laboratory and left the largest question, whether the samples themselves are right, unanswered.

THEMATIC ASIDE

Portable XRF: a field-decision tool, never a reported grade

Handheld X-ray fluorescence gives a geologist a fast, non-destructive reading in the field and can guide a decision to extend or infill a hole. The CIM guidelines treat it exactly as that, a rapid indicator whose strengths and shortcomings the user must understand and convey with the data, and they caution that its results are not a substitute for a laboratory assay. In this paper a portable-XRF number is treated as field triage only. It informs where to sample; it is never carried into a resource database as a grade, and it is never reported as an assay value.

> No control in the laboratory can rescue a sample that was never representative. Spend the attention where the variance actually lives.

3

EXPLORE · THE SCHEME, BY DESIGN

Designing the QC insertion scheme

A control cannot be inserted after the batch is read. The scheme, what goes in, where in the stream, and how often, is a design decision made before the first sample ships.

1 in 20
COMMON PER-TYPE INSERTION RATE; ABOUT 5 PERCENT EACH

every batch
AT LEAST ONE REFERENCE MATERIAL IN EACH ANALYTICAL BATCH

full range
REFERENCE GRADES THAT BRACKET THE DEPOSIT, NOT JUST ORE

The single property that distinguishes a QA/QC program from a pile of assay certificates is that the controls were designed in advance and inserted into the sample stream before it reached the laboratory. A certified reference material cannot be added to a batch after the batch has been read, and a field duplicate cannot be taken once the sample has been consumed. The scheme is therefore a design artifact, written into the sampling procedure and the sample-submission sheet before the program starts.

Three insertions, three questions

A well-formed scheme inserts three kinds of control sample, because there are three separate things to measure. Certified reference materials, inserted with a known certified grade, measure accuracy and bias: whether the laboratory is returning the true value, and whether it drifts over time or fails at the batch level after a recalibration. Blanks, made of material certified to be barren, measure contamination: whether grade is carrying over from a high sample to the next through the crusher or the pulveriser. Duplicates, second samples taken at a defined stage, measure precision: how repeatable the whole process is from the point the duplicate was split. The three questions are independent, and a scheme that inserts reference materials generously but takes few blanks and no coarse duplicates has left contamination and the largest component of sampling variance unmeasured.

EXHIBIT 4 · CAPABILITY INSTRUMENT · NO PROJECT DATA

What each insertion tests, what it detects, and where it belongs in the stream

A scheme is complete when all three questions are answered across the full grade range, not when one is answered thoroughly.

INSERTION	QUESTION	WHAT IT DETECTS	PLACEMENT
Certified reference material (CRM)	Accuracy, bias	Laboratory returning a value off the certified mean; drift over time; a batch failure after recalibration	Inserted blind in the stream; grades bracketing the deposit; at least one per batch
Blank (coarse and pulp)	Contamination	Carry-over between samples through preparation or analysis equipment	Immediately after a high-grade or visually mineralised sample
Field duplicate	Total precision	Combined variance from sub-sampling onward; the real repeatability of the result	A second split taken in the field or at the coarse reject
Preparation / pulp duplicate	Preparation and analytical precision	Variance added by preparation (coarse split) or by the instrument alone (pulp split)	Taken by the laboratory at the reject or pulp stage
Umpire (check) assay	Inter-laboratory bias	Systematic difference between the primary and a second accredited laboratory	A subset of pulps sent to a second laboratory

After CIM, Mineral Exploration Best Practice Guidelines, 2018, §7(d) (external blanks, certified reference materials, duplicates and third-party check sampling); Smee, Bloom, Arne and Heberlein, GEEA, 2024. Framework; no project data.

How often, and where the numbers come from

The binding guideline is qualitative. The CIM guidelines require that blanks and certified reference materials be included frequently enough to provide statistical confidence in the results, and that duplicates and third-party check sampling be part of the program, but they do not mandate a percentage. Industry practice has converged on a working default: each control type inserted at roughly one sample in twenty, about five percent, so that reference materials, blanks and duplicates together make up on the order of fifteen to twenty percent of the samples submitted, with at least one reference material in every analytical batch so that no batch is read without an accuracy check. This is a starting point, not a rule from a code, and it should be tuned to the deposit: a high-nugget gold system needs more duplicates to characterise its sampling variance, while a homogeneous bedded deposit may justify fewer.

Because the one-in-twenty figure is practice rather than a mandated number, it is flagged as established practice wherever it appears in this paper, and the reader is pointed to the qualitative requirement that actually binds. What is not optional is that the rate be chosen deliberately, documented in the QA/QC plan, and high enough that a control chart has the points it needs to distinguish a real shift from ordinary scatter. A program that inserts one reference material every hundred samples cannot detect a batch-level failure, because it has almost no chance of placing a control in the failing batch.

EXHIBIT 5 · CAPABILITY INSTRUMENT · NO PROJECT DATA

A batch of twenty, laid out: where the controls sit in the stream

Illustrative layout of a submission batch at a one-in-twenty per-type rate, with a reference material early so a batch failure is caught before the run completes.

Grey cells are ordinary samples; the three coloured positions are the inserted controls. A blank is placed to follow a likely high-grade sample so any carry-over shows; positions are varied between batches so the laboratory cannot recognise the pattern.

Illustrative layout at established-practice rates (about one in twenty per type); qualitative basis CIM, 2018, §7(d). Schematic; not project data.

Coarse and pulp blanks answer different questions

Not all blanks are the same, and a scheme that inserts only one kind measures only one kind of contamination. A coarse blank, submitted as barren rock or a certified coarse material and taken through the full preparation sequence, tests for carry-over introduced by the crusher and the pulveriser, the equipment that reduces one sample after another and can smear grade from a rich sample into the next. A pulp blank, inserted at the analytical stage, tests only for contamination introduced in the laboratory instrument and reagents. The coarse blank is the more demanding control because it exposes the preparation stage, which is where most physical carry-over actually happens, and it is the one more often omitted because it consumes a preparation slot. A program serious about contamination inserts coarse blanks, and places them immediately after the samples most likely to be high grade so that any carry-over has the best chance of showing.

Reference-material selection carries a parallel subtlety. A material chosen for convenience, one already on the shelf at a familiar grade, may not match the project mineralogy, and a material read against a digestion different from the one its certificate assumes will report a bias that belongs to the method mismatch and not to the laboratory. The scheme therefore specifies both the materials and the method they are read by, and it stores enough of each material, from a single homogenised batch, to last the program, so that a mid-program change of reference lot does not introduce a step in the control chart that looks like laboratory drift but is really a change of the ruler.

A last design point is blindness. The controls only work if the laboratory cannot tell a reference material or a blank from an ordinary sample, so they are submitted under the same numbering scheme, with no giveaway in the sample identifier, and their positions are varied from batch to batch. A reference material that the laboratory recognises can be treated with extra care, which defeats the purpose of measuring routine performance. Blindness is cheap to design in at the submission stage and impossible to add afterward.

You cannot insert a control into a batch that has already been read. The scheme is a decision made before the first sample ships, or it is not made at all.

AURUS TECHNICAL COMMITTEE

The cheapest control is the one you designed before you needed it.

AURUS INSTITUTE FOR RESOURCE DEVELOPMENT

4

EXPLORE · MEASURING AGAINST A KNOWN TRUTH

Certified reference materials

A reference material is a sample whose grade is already known to a stated uncertainty. Inserted blind, it turns accuracy from an opinion into a measurement.

consensus
CERTIFIED MEAN FROM A ROUND-ROBIN OF ACCREDITED LABORATORIES

bracket
GRADES CHOSEN TO SPAN THE DEPOSIT, NOT JUST ORE

drift
THE SLOW BIAS A SINGLE CONTROL POINT CANNOT SEE

Accuracy is the question of whether a measurement returns the true value. It cannot be answered by repeating a measurement, because a biased process repeats its bias faithfully. It can only be answered by measuring something whose true value is already known, and that is what a certified reference material provides: a prepared, homogenised sample whose grade has been established to a stated uncertainty and against which the project laboratory's result can be judged.

What "certified" actually means

A certified reference material carries a certified value and an uncertainty that were established, in the usual case, by a round-robin exercise in which a number of accredited laboratories each analysed the material by agreed methods. The certified mean is the consensus of those results, and the associated standard deviation captures the between-laboratory spread. When a project inserts that material and reads it, the deviation of the reading from the certified mean is a direct measurement of the laboratory's bias for that grade and matrix. The recommended-practice review is explicit that reference materials monitor accuracy and bias against consensus round-robin values, that they track drift over time within a laboratory, and that they identify batch-level failures caused by an abrupt change such as a recalibration or a procedural change.

EXHIBIT 6 · CAPABILITY INSTRUMENT · NO PROJECT DATA

The anatomy of a reference-material certificate, read as a control specification

Each field on the certificate becomes a parameter in the control chart the material will be plotted on.

CERTIFICATE FIELD	WHAT IT FIXES FOR THE CONTROL
Certified mean	The centre line of the control chart for that material
Standard deviation	The spacing of the warning (two SD) and action (three SD) limits
Certifying round-robin	The accredited laboratories and methods behind the consensus value; the basis for its authority
Matrix and mineralogy	Whether the material resembles the project ore, so the bias it measures is the relevant one
Recommended method	The digestion and finish the certified value assumes; a partial digest read against a total-digest certificate is a false failure
Homogeneity and stability	Confidence that every insertion of the material is the same material

After recommended practice, Smee, Bloom, Arne and Heberlein, GEEA, 2024; laboratory competence framed by ISO/IEC 17025:2017. Framework; no project data.

Choosing the materials: bracket the deposit

A single reference material controls accuracy at a single grade. A resource spans a range, from the low grades that set the cut-off boundary to the high grades that carry the metal, and a laboratory's bias is not guaranteed to be constant across that range. Good practice is therefore to insert several reference materials whose certified grades bracket the deposit: at least one near the expected cut-off, one near the average grade, and one high, so that accuracy is controlled across the population the resource is estimated from rather than at a single convenient point. This is the operational meaning of the CIM requirement that the program span the full range of values and not just high or unusual results.

Matrix match matters as much as grade. A reference material certified in a matrix unlike the project ore, a different host rock or a different mineralogy of the element of interest, may respond differently to the chosen digestion and report a bias that is an artifact of the mismatch rather than a property of the laboratory. The material should resemble the ore in the ways that affect the assay, and where a well-matched commercial material does not exist, a matrix-matched material prepared from project reject and certified through its own round-robin is the stronger, if more involved, control.

EXHIBIT 7 · CAPABILITY INSTRUMENT · NO PROJECT DATA

Bracketing accuracy across the grade range the resource is built on

Three reference grades placed against a schematic grade distribution; the low material controls the cut-off boundary, where the tonnes are most sensitive.

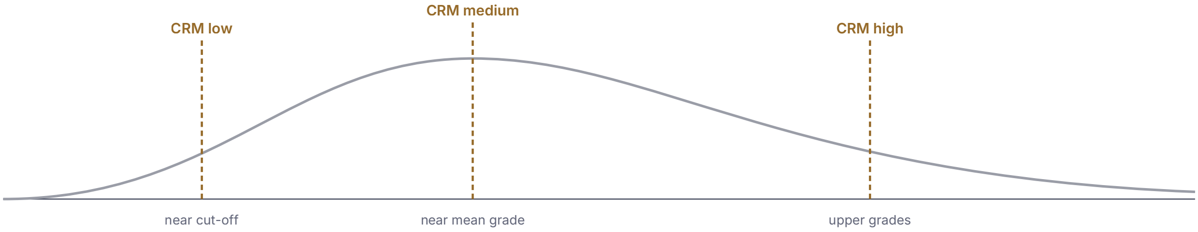

Horizontal axis: grade. Curve: schematic frequency of sample grades. Accuracy is controlled where the resource is most sensitive, not merely where the metal sits.

After CIM, 2018, §7(d) (full range of values); recommended practice, GEEA, 2024. Schematic distribution; no project data.

Where a well-matched commercial material is used, its certificate should be kept with the project record, because the control chart is only meaningful against the certified mean and standard deviation that certificate carries. Where a matrix-matched material is prepared from project reject, the work is larger but the control is stronger: a bulk of representative reject is homogenised, split, and sent through a round-robin of accredited laboratories to establish a consensus value and a between-laboratory standard deviation, after which it behaves exactly as a commercial reference material does. Either way the governing requirement is the same, that the material resemble the ore in the ways that affect the assay and that enough of a single homogeneous lot exists to control the whole program without a mid-stream change of reference.

Drift and the batch-level failure

A single reference-material result tells you whether one batch was accurate. The sequence of results over time tells you something a single point cannot: whether the laboratory is drifting, slowly biasing high or low as a reagent ages or an instrument falls out of calibration, and whether a discrete event has shifted the process, such as a recalibration that moved every result after it. These are patterns in the run of control points, and they are the reason reference materials are plotted as a time-ordered control chart rather than merely checked one by one against a tolerance. A program that records only pass or fail for each insertion, without plotting the sequence, throws away the drift signal that is often the earliest warning of a developing problem.

When the reference material itself is the problem

A reference material is a control, and controls can fail. A material can degrade or segregate if it is poorly stored, so that late insertions no longer match the certified value; a certificate can be misread, most often by comparing a partial-digestion result against a value certified for a total digestion; and a commercial lot can simply be a poor match for the project mineralogy. When a reference material begins to fail consistently, the disciplined first question is whether the laboratory has drifted or whether the control has, and the answer usually comes from the other insertions: if the blanks and duplicates remain clean and only one reference material fails, the material or its handling is the more likely cause, whereas a broad failure across several materials points to the laboratory. Treating a failing control as automatically a failing laboratory is how a good batch gets needlessly re-assayed and a real bias, hidden behind a mismatched reference, gets missed.

This is why more than one reference grade is inserted and why their certificates are kept with the project record. Redundancy in the controls is what lets a program tell a laboratory problem from a control problem, and the cost of a second and third reference material is small against the cost of misreading which one has failed. The same logic argues for a modest program of umpire check assays at a second accredited laboratory: an independent laboratory reading the same pulps is the cleanest test of whether an apparent bias belongs to the primary laboratory or to the controls being used to judge it.

> Repetition proves precision, never accuracy. Only a known truth, inserted blind, can tell you whether the number is right.

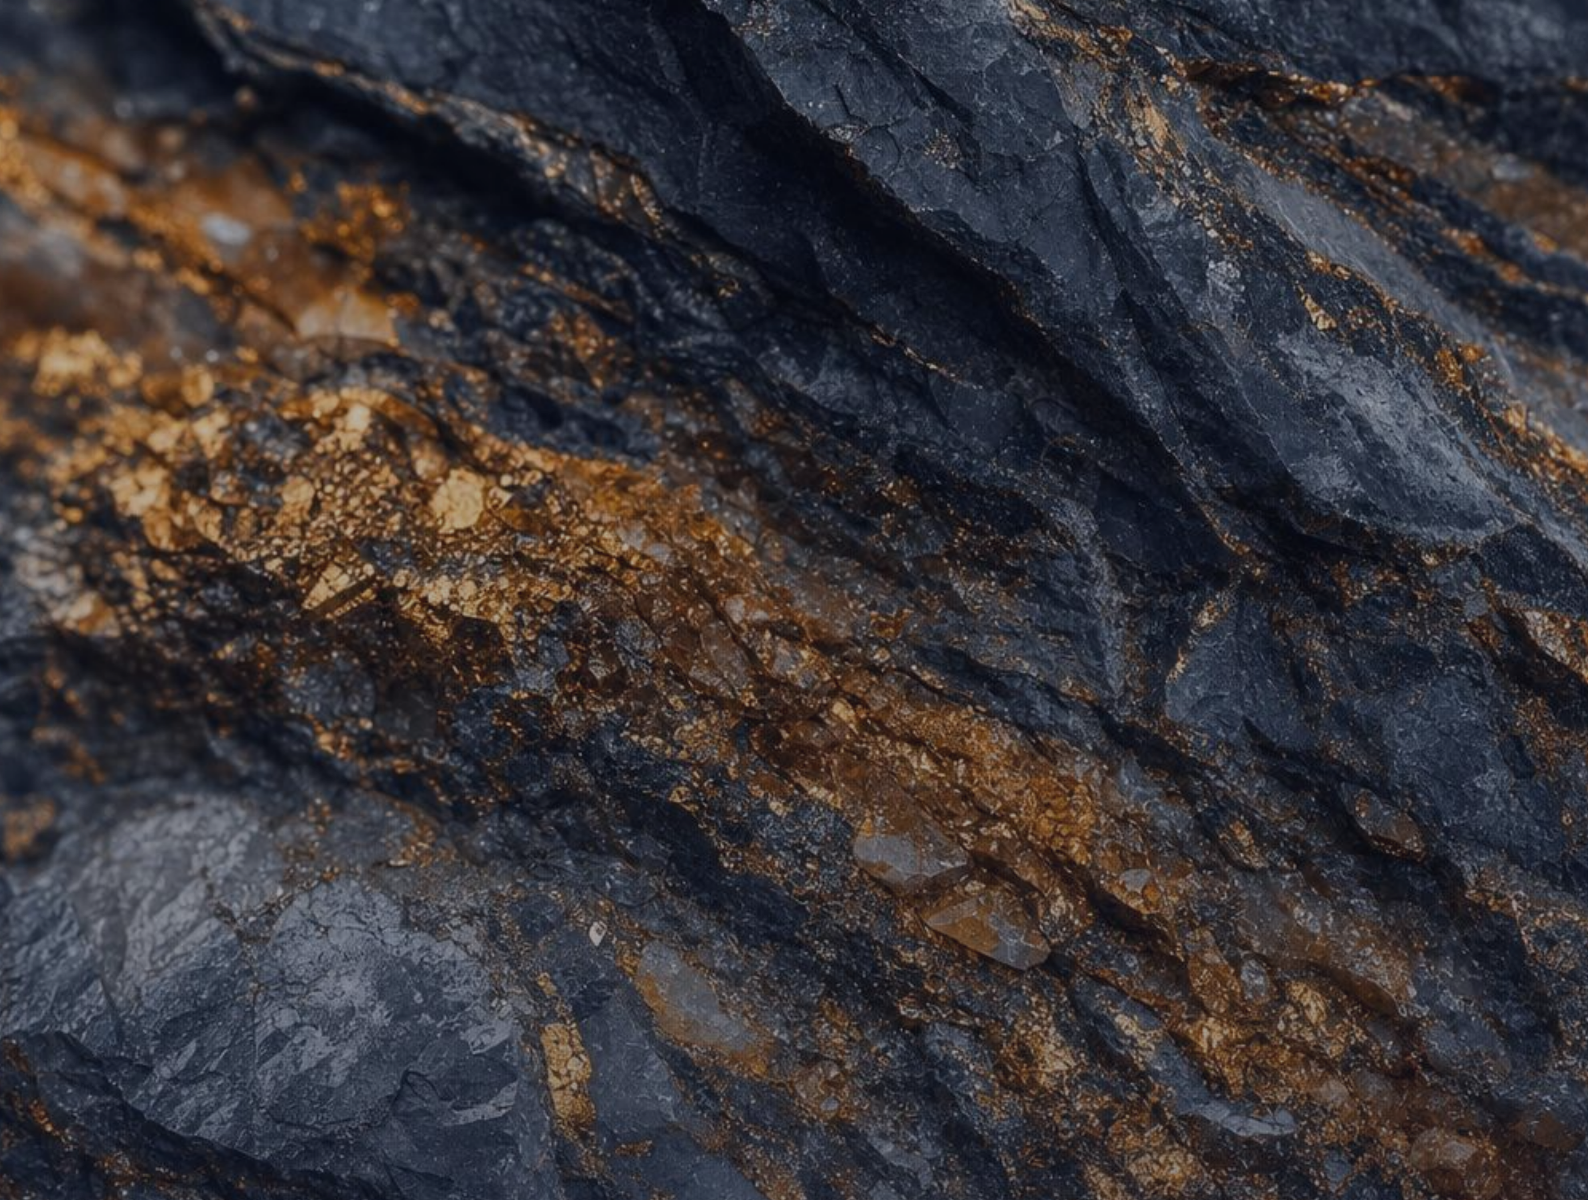

5

EXPLORE · THE SIGNATURE EXHIBIT

Precision, accuracy and contamination

Three control charts, read side by side, tell the whole story of a dataset's quality. This chapter reads each one, and then reads a failure.

±2 / ±3 SD
WARNING AND ACTION LIMITS ON THE ACCURACY CHART

HARD
HALF ABSOLUTE RELATIVE DIFFERENCE, THE PRECISION MEASURE

triage
A DEFINED RESPONSE, NOT A NEGOTIATION, WHEN A CHART FAILS

The three insertions come together in three control charts, and the discipline of QA/QC is largely the discipline of reading them honestly and acting on them by a rule set in advance. This chapter is the paper's centre. The signature exhibit puts the three charts side by side, because a dataset is only as good as its worst chart, and a program that watches accuracy while ignoring precision or contamination is watching one door of three.

Reading the three charts

The accuracy chart plots each reference-material result against time, with the certified mean as the centre line and control limits at two and three standard deviations. The convention, drawn from Shewhart control charts and used throughout the recommended-practice literature, is that a result beyond two standard deviations is a warning and a result beyond three is an action: the batch is failed and re-assayed. A common batch-failure rule also fails on two consecutive results beyond two standard deviations on the same side, because that pattern signals a bias rather than random scatter even when no single point has crossed the action limit. The exact rule is a program decision, written into the QA/QC plan and confirmed against the laboratory's own data-quality objectives, and this paper flags the two-and-three-standard-deviation convention as established practice rather than a figure any single code mandates.

The contamination chart plots blank results against a threshold set a small multiple above the analytical detection limit. A blank should read at or near nothing; a blank that carries grade after a high-grade sample is direct evidence of carry-over through the crusher, the pulveriser or the instrument, and it condemns not the blank but the samples around it. The precision chart plots paired data, original against duplicate, and summarises their agreement. A widely used measure is the half absolute relative difference of each pair, and a common acceptance is that the large majority of pairs fall below roughly ten percent, with the tolerance relaxing near the detection limit where relative differences are naturally large. Precision can also be summarised across a whole program by an average coefficient of variation, for which the recommended-practice review supplies a calculation tool.

EXHIBIT 8 · SIGNATURE EXHIBIT · CAPABILITY INSTRUMENT · NO PROJECT DATA

The control-chart triptych: accuracy, contamination and precision, read together

Illustrative charts. The point is how each is read and where its action line sits, not the specific plotted values.

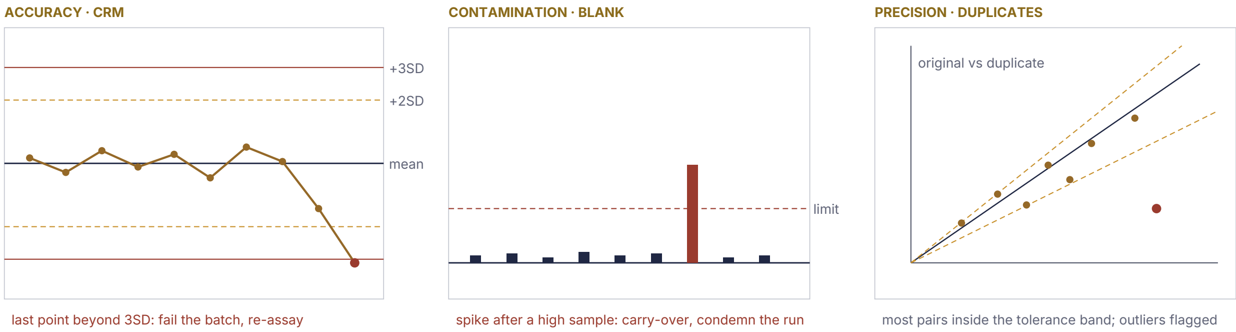

Illustrative control charts after recommended practice (Shewhart limits; HARD precision), Smee, Bloom, Arne and Heberlein, GEEA, 2024; CIM, 2018, §7(d). Not project data; plotted values are schematic.

Precision as a scatterplot, and the average CV

Precision is best seen, not just tabulated. The standard picture is a scatterplot of each duplicate pair, the original grade on one axis and the duplicate on the other, with the line of perfect agreement drawn through the origin. Well-controlled pairs cluster tightly around that line; poor precision spreads them, and the spread usually widens toward the low-grade end where relative differences are naturally larger. Reading the plot rather than a single summary number matters because it distinguishes ordinary low-grade scatter, which is expected, from a genuine loss of precision across the whole range, which is not, and because it exposes the occasional gross outlier, a pair that disagrees by an order of magnitude, that a sample swap or a transcription error would produce.

For a single figure that travels well between programs, precision can be summarised as an average coefficient of variation, the mean over all pairs of the standard deviation of each pair divided by its mean. The recommended-practice review provides a calculation for this average coefficient of variation and treats it as the headline precision statistic for a dataset, alongside the pair plot that shows where the variance lives. A program reports both: the plot, which localises the problem, and the single statistic, which lets one program be compared against another and against the data-quality objective set at the outset.

Triage: what a failure sets in motion

A control chart is only worth the response it triggers. When the accuracy chart fails, the primary action is to quarantine the affected batch and re-assay from retained pulp; if the pulp re-assay passes, the original run was the problem, and if it fails too, the problem is upstream in preparation and the coarse reject must be revisited. When a blank spikes, the samples bracketing it are suspect for carry-over and are re-run, and the preparation sequence is examined for a cleaning failure. When duplicates show poor precision, the stage at which the duplicate was taken localises the cause: poor field-duplicate agreement with acceptable pulp-duplicate agreement points to sampling and preparation, not the instrument. The value of writing this triage down in advance is that a failing batch triggers a defined action rather than a discussion about whether the failure was real.

EXHIBIT 9 · CAPABILITY INSTRUMENT · NO PROJECT DATA

The failure-response table: each signal has one pre-defined first action

Written before the data arrives, so a breach is triaged, not argued.

CONTROL SIGNAL	PRE-DEFINED FIRST ACTION
CRM beyond three standard deviations	Fail and quarantine the batch; re-assay from retained pulp
Two consecutive CRM beyond two SD, same side	Treat as a bias signal; hold results, investigate calibration
Blank above the contamination limit	Re-run the bracketing samples; audit preparation cleaning
Field duplicate precision poor, pulp duplicate acceptable	Investigate sampling and preparation, not the instrument
Umpire laboratory shows systematic offset	Open an inter-laboratory bias investigation; widen the check batch

After recommended practice, GEEA, 2024; CIM, 2018, §7(d) (documented pass/fail criteria and actions on failures). Framework; no project data.

One discipline underlies all of it: the controls are read in near real time, batch by batch, not assembled at the end of the program. A failure caught within a batch can be fixed from pulp that still exists and re-run before the next batch compounds it. The same failure discovered at resource estimation, months later, is discovered when the pulp may be gone and the only remedy is to re-drill. Timeliness is itself a control.

A grade you cannot audit is a rumour with a decimal point.

AURUS INSTITUTE FOR RESOURCE DEVELOPMENT

6

EXPLORE · WHERE ERRORS REALLY ENTER

Databases and chain of custody

The best-controlled assay in the world is worthless if it is entered against the wrong hole. Most real errors are clerical, and the database is where they are caught or missed.

relational
PREFERRED TO SPREADSHEETS FOR CONTROL AND SECURITY

verify
A PROGRAM OF DATA VERIFICATION, NOT A ONE-TIME CHECK

17025
THE LABORATORY'S SCOPE OF ACCREDITATION, METHOD BY METHOD

A great deal of QA/QC attention goes to the assay, because that is where the reference materials are read. Yet a large share of the errors that actually corrupt a dataset are not analytical at all. They are clerical: a result typed against the wrong sample number, a depth interval transposed, a unit misread, two files merged on a mismatched key. These errors are invisible to the control charts, because the reference materials and duplicates were themselves recorded correctly; only the ordinary samples were mis-entered. The database is where they are caught, or where they are missed.

Why the guidelines prefer a database to a spreadsheet

The CIM guidelines are direct on this point. For discovery-stage and delineation-stage properties, they state, storage of drill-hole data in a relational database that provides proper control and security is preferred to spreadsheets, which cannot be secured as effectively and are prone to an increased likelihood of errors during data manipulation. They require a database-management protocol that includes a data-verification program to confirm that accurate, error-free information is entered, and they note that, given the cost of drilling and analysis, the cost of proper data storage is justified to protect the integrity of the project data. The preference is not aesthetic. A relational database enforces referential integrity, so an assay cannot be entered against a hole that does not exist; it applies validation rules at entry; and it keeps an audit trail of who changed what, which a spreadsheet cannot.

EXHIBIT 10 · CAPABILITY INSTRUMENT · NO PROJECT DATA

Why the reporting-grade record lives in a database, not a spreadsheet

The comparison the CIM guidelines draw, made concrete as a set of controls a database enforces and a spreadsheet cannot.

CONTROL	RELATIONAL DATABASE	SPREADSHEET
Referential integrity	An assay cannot attach to a non-existent hole or interval	Any value can be typed into any cell
Validation at entry	Ranges, types and code lists enforced on input	Optional, easily overridden or deleted
Audit trail	Records who changed what and when	No native change history
Concurrent access	Controlled, with permissions	File copies diverge and are merged by hand
Verification	Systematic re-checks against the primary certificates	Ad hoc, if done at all

After CIM, Mineral Exploration Best Practice Guidelines, 2018, §6(f) and §7(d). Framework comparison; no project data.

Verification as a program, and the laboratory audit

Verification is not a one-time acceptance at the end. The guidelines frame it as an ongoing program: a defined share of the database is checked back against the source laboratory certificates and the primary logs, discrepancies are investigated and corrected, and the check is repeated as new data loads. The same discipline extends outside the company's walls. The CIM guidelines call for audits of the methods and procedures used by the primary assay laboratory on a periodic basis, and they require that analysis be done by a reputable and preferably ISO-accredited laboratory, with results supported by signed certificates and a statement of the methods used.

The strongest single reduction in clerical error is to remove the keyboard from the path. Where the laboratory delivers results as a structured electronic file that loads directly into the database against the submitted sample numbers, the transcription step that produces so many errors is eliminated, and the load can be validated automatically against the submission list so that a missing or extra sample is caught at import. The QC results ride in the same file, so the control charts update as the batch loads and a failure is visible immediately. This does not remove the need for verification, because a file can still be loaded against the wrong survey or joined on a mismatched key, but it moves the residual risk from hundreds of hand-typed values to a small number of well-defined load operations that can each be checked.

ISO/IEC 17025:2017, the standard behind that accreditation, sets the general requirements for the competence of testing and calibration laboratories, and a laboratory's accreditation is not a blanket badge but a scope: a schedule of the specific methods, elements and ranges for which competence has been demonstrated. A reviewer confirms not merely that the laboratory is accredited, but that the method actually used for the project sits inside its accredited scope. An accredited laboratory running a project on a method outside its schedule is, for that work, an unaccredited laboratory.

THEMATIC ASIDE

The chain of custody is a document, not a good intention

Custody is evidenced, not asserted. A defensible trail records each transfer of the samples, from the core shed to the preparation facility to the laboratory, with dates, quantities and the person or party accepting them, so that a reviewer can confirm the sample read as a given hole and depth is that sample. The QC insertions reinforce the trail: a reference material that returns the wrong certified value can expose a mix-up as readily as a mis-assay. Custody records and control samples are two halves of the same assurance that the number belongs to the sample it claims.

The audit trail is the deliverable

It is worth stating plainly what all of this produces, because the deliverable of a QA/QC program is not a clean control chart but an audit trail: an unbroken record that lets an independent person start at the reported grade and walk back, through the database entry, the laboratory certificate, the QC results for that batch, the custody record and the sampling log, to the interval in the ground. Each link is a document that exists because a control required it. The reference-material result exists because the scheme inserted it; the certificate exists because the guidelines require a signed one; the custody record exists because sample security is a recorded procedure; the verified database entry exists because a data-verification program checked it. A reviewer does not audit the grade. The reviewer audits the trail, and the grade is trustworthy exactly to the extent the trail is complete.

A program built this way is expensive to fake and cheap to defend, which is the property that makes it worth its cost. Its opposite, a dataset assembled without the trail, is cheap to produce and impossible to defend, and it fails at precisely the moment defence is needed, when a lender's or an acquirer's technical adviser asks to see the controls and finds an assurance where the evidence should be. The whole of this paper is, in the end, an argument for building the trail deliberately and from the start, because a trail cannot be added to data that was collected without it.

Most errors that reach a resource model never touched a crusher. They were typed. The database is the control that decides whether they are caught.

AURUS TECHNICAL COMMITTEE

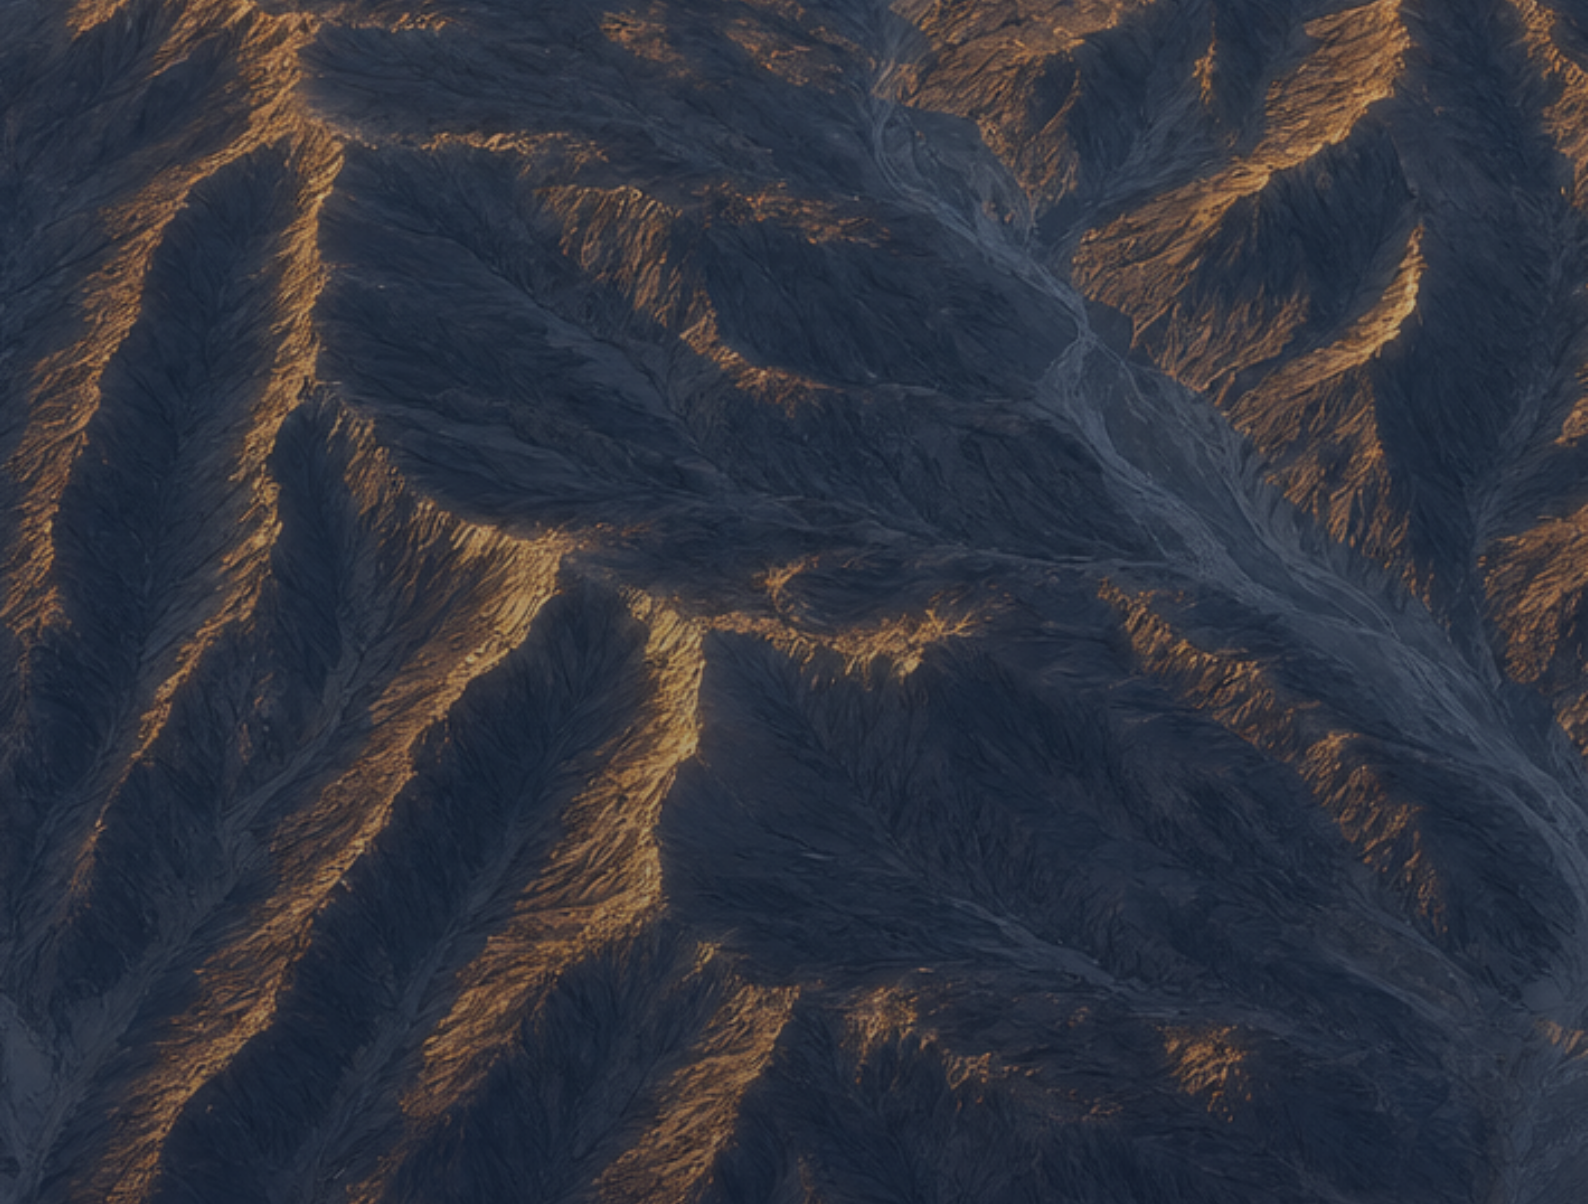

7

EXPLORE · THE DISCLOSURE TEST

The regulatory lens

A code reviewer reads a public report looking for the QA/QC that is missing. Table 1 turns silence into a disclosure, and silence about controls is the first red flag.

if not, why not
EVERY TABLE 1 CRITERION IS ADDRESSED OR EXPLAINED

Section 1
SAMPLING, SUB-SAMPLING, ASSAY QUALITY, VERIFICATION, SECURITY

QP / CP
A NAMED PERSON SIGNS AGAINST THE AUDIT TRAIL

Everything in the preceding chapters exists so that a report can pass a specific test: the scrutiny of a Qualified or Competent Person, acting for a lender, an exchange or a regulator, who reads the disclosure looking for the control that should be there and is not. Understanding what that reviewer looks for first is the fastest way to see why the program is built the way it is.

Table 1 as a checklist the reviewer already holds

The JORC Code, 2012 Edition, requires that a public report of exploration results, mineral resources or ore reserves be accompanied by Table 1, a structured set of criteria addressed on an "if not, why not" basis, so that any criterion left unaddressed is an explicit and visible gap rather than a quiet omission. Section 1 of Table 1, "Sampling Techniques and Data", is the QA/QC section, and it reads almost as a specification for the program this paper describes. It calls for statements on sampling technique, drilling technique, drill-sample recovery, logging, sub-sampling techniques and sample preparation, the quality of assay data and laboratory tests, the verification of sampling and assaying, sample security, and audits or reviews.

The criterion a reviewer turns to first is the quality of assay data: "nature of quality control procedures adopted (eg standards, blanks, duplicates, external laboratory checks) and whether acceptable levels of accuracy (ie lack of bias) and precision have been established". A report that describes drilling in detail and then treats this criterion with a single sentence about the laboratory's accreditation, without reference materials, blanks, duplicates or check assays, has answered the question by not answering it. The adjacent criterion, verification of sampling and assaying, asks specifically about the use of twinned holes and about the documentation of primary data, data-entry procedures, data verification and data storage, which is why the database discipline of Chapter 6 is a disclosure matter and not merely good housekeeping.

EXHIBIT 11 · CAPABILITY INSTRUMENT · NO PROJECT DATA

The Table 1 Section 1 criteria, read as the red flags a reviewer looks for

For each criterion, the disclosure a reviewer expects, and the absence that reads as a red flag.

TABLE 1 CRITERION	WHAT THE REVIEWER EXPECTS	THE RED FLAG
Drill sample recovery	Recovery measured and logged; relationship to grade examined	Recovery not reported, or not related to grade
Sub-sampling and preparation	Split procedure, crush and pulverise specification, prep duplicates	No preparation duplicates; unstated split method
Quality of assay data	CRMs, blanks, duplicates, external checks; accuracy and precision established	Only laboratory accreditation cited; no control samples
Verification of sampling and assaying	Twinned holes where relevant; documented data entry, verification, storage	No independent verification; spreadsheet-only storage
Sample security	A recorded chain of custody from field to laboratory	Custody asserted but not evidenced
Audits or reviews	Periodic laboratory audits; independent QA/QC review	No audit; QA/QC compiled only at report time

Criteria from JORC Code, 2012 Edition, Table 1, Section 1 (artifact on disk); expectations after CIM, 2018, §6 and §7. Framework; no project data.

The instrument behind the signature

Codes differ in their machinery but agree on the principle. In the CRIRSCO family, JORC 2012, SAMREC and PERC, together with the CIM Definition Standards that Canada's NI 43-101 incorporates, a named professional takes responsibility for the disclosure. NI 43-101 is a disclosure instrument rather than a reporting code: it defines the Qualified Person and requires that a technical report be prepared by or under the supervision of one, resting on the CIM Definition Standards for its resource and reserve categories. The Qualified Person is not certifying that the grade is correct in some absolute sense; the person is certifying that the data supporting it were collected and controlled to a standard that makes the estimate reasonable, and that the QA/QC results support it. The audit trail this paper has described is precisely what that person signs against.

One discipline, two settings: exploration and grade control

The same principles apply at two moments in a project's life, and the reviewer knows to weigh them differently. During exploration and resource definition, the QA/QC program controls the data that a public resource estimate rests on, and the disclosure lens is Table 1 Section 1: reference materials, blanks, duplicates and check assays, disclosed with their pass and fail outcomes. During production, grade control runs on faster, cheaper, higher-volume sampling to decide, day by day, what is ore and what is waste, and its controls are tuned to that decision: more duplicates to characterise short-range variability, a tighter feedback loop, and reconciliation of the predicted grade against what the plant actually receives. The recommended-practice review is explicit that a QA/QC program should be adjusted over time to meet the changing data-quality requirements at different stages of development and exploitation, rather than run as a single fixed scheme.

For a reviewer, the practical consequence is that the questions asked of a dataset depend on what the dataset is for. A grade-control stream that reconciles well against the plant is doing its job even if its per-sample precision is looser than a resource-definition program would accept, and a resource-definition program is judged on the accuracy and precision established across the full grade range rather than on throughput. Confusing the two, holding grade control to a resource standard or reporting a resource on grade-control data, is itself a red flag, because it signals that the purpose of the sampling was not matched to the control designed around it.

THEMATIC ASIDE

Reporting codes and disclosure instruments are not the same thing

Precision of language matters here. The CRIRSCO-family reporting codes are JORC (2012), SAMREC and PERC, alongside the CIM Definition Standards. NI 43-101 is a Canadian disclosure instrument that incorporates the CIM Definition Standards and defines the Qualified Person; it is not itself a CRIRSCO code. Data are reported under, or delivered to, these codes; a code is applied, never possessed. A report that speaks of being compliant with a code it has mis-classified invites a reviewer to doubt the rest.

> Table 1 does not ask whether you did QA/QC. It asks you to show it, and it turns every silence into a statement.

8

EXPLORE · THE STANDARD, ASSEMBLED

The Aurus QA/QC standard

A template protocol, the small set of KPIs that show at a glance whether it is holding, and the non-conformance workflow that closes the loop when it is not.

---

design first
THE SCHEME WRITTEN BEFORE THE FIRST SAMPLE SHIPS

KPIs
A SHORT DASHBOARD, READ EVERY BATCH

close the loop
EVERY NON-CONFORMANCE TO A DOCUMENTED RESOLUTION

The preceding chapters can be assembled into a single standard: a protocol a team can hold a program against, a small set of indicators that show whether it is holding, and a workflow that turns a control failure into a documented resolution. This chapter sets out that standard as a framework instrument. It carries no client data; the numbers in it are targets and defaults, not results.

The protocol, in six stages

The protocol follows the sample. Before the program starts, the QA/QC plan is written: the reference materials chosen to bracket the deposit, the insertion rates, the blank and duplicate scheme, the control limits and the batch-failure rule, and the triage response for each signal. During drilling, recovery is logged as a control and the chain of custody begins at the core tray. At preparation, the split method and comminution specification are enforced and preparation duplicates are taken. At the laboratory, an accredited method within scope is used, and the blind controls travel in the stream. On receipt, the controls are plotted and read batch by batch, and any failure triggers the triage. Continuously, the database is verified against the primary certificates, and the laboratory is audited periodically. The stages are not sequential paperwork; they are a loop that runs for the life of the program.

EXHIBIT 12 · CAPABILITY INSTRUMENT · NO PROJECT DATA

The QA/QC dashboard: a short set of indicators, read every batch

Illustrative targets. A program is holding when these stay green; each maps to a chapter and a control.

INDICATOR	DEFINITION	ILLUSTRATIVE TARGET
QC insertion rate	Control samples as a share of samples submitted	About 15 to 20 percent
CRM pass rate	Reference materials inside the action limits	Investigate below about 95 percent
Blank pass rate	Blanks below the contamination limit	Near 100 percent
Duplicate precision	Share of pairs within the precision tolerance	Majority within about 10 percent HARD
Umpire agreement	Primary against second-laboratory bias	No systematic offset
Verification coverage	Database checked back to primary certificates	A defined, disclosed share

Illustrative targets at established-practice levels; qualitative basis CIM, 2018, §7(d); precision and limit conventions after GEEA, 2024. Framework; no project data.

The non-conformance workflow

A control failure is a non-conformance, and the point of the standard is that every non-conformance runs to a documented close rather than a shrug. The workflow is short and fixed: detect the signal on the chart, quarantine the affected data, investigate to a root cause, act to correct it, verify that the correction worked, and record the whole event so a reviewer can see both that it happened and that it was resolved. The CIM guidelines ask precisely for this, that the QA/QC disclosure include the pass and fail criteria and the actions taken on results outside the limits, which is only possible if those actions were defined and logged as they occurred.

EXHIBIT 13 · CAPABILITY INSTRUMENT · NO PROJECT DATA

Closing the loop: a non-conformance from signal to documented resolution

The same six steps for every failure, so the response is a procedure and the record is auditable.

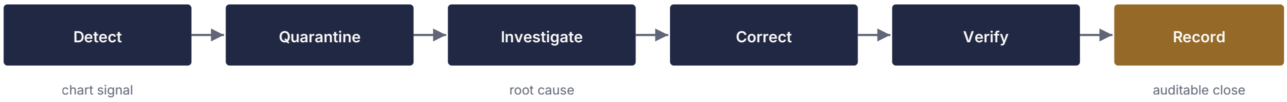

After CIM, Mineral Exploration Best Practice Guidelines, 2018, §7(d). Framework workflow; no project data.

What a reviewer will ask for

The standard can be stated a final way, as the short list of things a lender's or an acquirer's technical adviser asks to see, because a program that can produce this list on request has met the bar and a program that cannot has not. The adviser asks for the QA/QC plan, dated before the program started. They ask for the control charts, the reference-material sequences, the blank record and the duplicate plots, read to the end of the program rather than assembled for the meeting. They ask for the non-conformance log, because a program with no recorded failures across a large campaign is not a program that never failed but a program that was not watching. They ask for the laboratory's scope of accreditation, to confirm the method used sits inside it. And they ask to trace one number, chosen at random, from the report back to the sample in the ground.

Each of these has appeared in this paper as a control, and together they are simply the audit trail made into a request. A team that has worked to the standard hands the list over without ceremony, because every item already exists as a by-product of doing the work properly. That, in the end, is the test of a QA/QC program: not whether it can be described, but whether it can be produced, in full, on the day someone with capital at risk decides to look.

EXHIBIT 14 · CAPABILITY INSTRUMENT · NO PROJECT DATA

The diligence request, and the chapter that already produced each item

A program built to this paper can answer every row from records that already exist; the right-hand column is where each was made.

WHAT THE TECHNICAL ADVISER ASKS FOR	WHERE THE PROGRAM ALREADY MADE IT
The QA/QC plan, dated before the program started	Chapter 3, the insertion scheme designed in advance
Control charts read to the end of the program	Chapter 5, the triptych read batch by batch
The non-conformance log, with actions taken	Chapter 8, the workflow closed to a record
The laboratory scope of accreditation	Chapter 6, the method confirmed within scope
The chain-of-custody record	Chapters 2 and 6, custody recorded as it moved
One number traced from report to the ground	Chapter 1, the audit trail through every link

After CIM, Mineral Exploration Best Practice Guidelines, 2018, §7 (disclosure, custody, accreditation); JORC Code, 2012 Edition, Table 1, Section 1. Framework; no project data.

How Aurus applies the standard

The standard above is a framework. Where Aurus has applied its data-quality discipline in the field, the record is stated here at the altitude the work supports, and no further. These are practice notes, not case studies, and they carry no project-specific data or results beyond what is set out.

AURUS PRACTICE NOTE

A ground magnetic and radiometric reconnaissance campaign over a manganese exploration permit in Central Africa was delivered under Competent-Person supervision to the CRIRSCO-family codes (JORC 2012, SAMREC) and the CIM Definition Standards incorporated by NI 43-101, with full data-quality acceptance of the delivered dataset. The mandate is cited here as evidence of a delivered data-governance discipline, not as a resource statement; no grade or instrument value is reported, and the survey's scope is described only at this level.

AURUS PRACTICE NOTE · SERVICE LINE

For a multi-permit exploration portfolio, Aurus designed a turnkey monthly technical-reporting program built on ISO 17025 laboratory chains, a QA/QC insertion discipline of about five percent, alignment to the CRIRSCO-family codes and the CIM Definition Standards incorporated by NI 43-101, and legally opposable e-signed deliverables. This is a designed reporting architecture offered as a service line; it is described in the proposal-stage terms that its status supports, not as an operated program of stated duration.

AURUS PRACTICE NOTE · SERVICE LINE

A drilling-campaign supervision mandate was structured to place a full sample-quality and QA/QC discipline around a diamond and reverse-circulation program: a recorded chain of custody, a designed insertion scheme of reference materials, blanks and duplicates, and an ISO-accredited analytical laboratory. The mandate is cited as evidence of the service line and its method; no delivery, duration or outcome is claimed for it here.

The thread through all three is the argument of this paper. Data quality is designed before the drilling starts, evidenced as the samples move, and read while there is still time to act. A team that works this way produces a number a reviewer can audit and a lender can underwrite, which is the only kind of number worth reporting.

> **Build the control in from the first hole. A resource is not made bankable at the report; it is made bankable, or not, in the sampling.**
>
> AURUS TECHNICAL COMMITTEE

REFERENCES

Codes, guidelines and reviews

Every figure and framework in this paper traces to one of the sources below, cited by institution, publication and year. The three primaries that carry the paper (the CIM guidelines, the JORC Table 1 template and the recommended-practice review) are archived on disk in the series evidence dossier; the standards catalogued below are cited by their published records. Where a numeric convention is widely used industry practice rather than a single code's mandated figure, the paper marks it as established practice at the point of use, and the Method and evidence note lists every such row.

1. Canadian Institute of Mining, Metallurgy and Petroleum (CIM), Mineral Exploration Best Practice Guidelines, prepared by the Mineral Resource and Mineral Reserve Committee, 23 November 2018. (Primary; §6 to §7 the QA/QC spine.)
2. Canadian Institute of Mining, Metallurgy and Petroleum (CIM), CIM Definition Standards on Mineral Resources and Mineral Reserves, 2014.
3. Joint Ore Reserves Committee (JORC) of the AusIMM, AIG and Minerals Council of Australia, Australasian Code for Reporting of Exploration Results, Mineral Resources and Ore Reserves (The JORC Code), 2012 Edition, including Table 1. (Primary; Section 1 the regulatory-lens spine.)
4. Smee, B.W., Bloom, L., Arne, D. and Heberlein, D., 2024, "Practical applications of quality assurance and quality control in mineral exploration, resource estimation and mining programmes: a review of recommended international practices", Geochemistry: Exploration, Environment, Analysis, vol. 24, no. 2, article geochem2023-046, Geological Society of London (Reviews in Exploration Geochemistry collection). (Primary review; abstract and record archived.)
5. International Organization for Standardization / International Electrotechnical Commission, ISO/IEC 17025:2017, General requirements for the competence of testing and calibration laboratories, 2017.
6. Canadian Securities Administrators, National Instrument 43-101 Standards of Disclosure for Mineral Projects (disclosure instrument incorporating the CIM Definition Standards; defines the Qualified Person).
7. CRIRSCO (Committee for Mineral Reserves International Reporting Standards), International Reporting Template, 2019 (the family within which JORC, SAMREC and PERC sit).
8. Abzalov, M., 2011, quality-assurance and quality-control discussion for exploration sampling, as cited in CIM (2018), §7(d).
9. Stanley, C.R. and Lawie, D., 2007, work on relative error and precision in geochemical determinations, as discussed in Smee et al. (2024).
10. Gazley, M. and Fisher, L., 2014, review of the reliability and validity of portable-XRF data, as cited in CIM (2018), §7(b).

Archived copies of the three primaries (CIM 2018 guidelines, JORC 2012 Table 1 template, and the bibliographic record and abstract of the GEEA review) are held in the series evidence dossier and were read directly in preparing this paper. The ISO and NI 43-101 items are cited from their published catalogue records; confirm any clause number against the official text at the point of use.

Exhibit source index

EXHIBIT	SUBJECT	ANCHOR SOURCE
1	The chain from ground to report	CIM 2018; JORC 2012
2	Diamond core against RC, and the control each demands	CIM 2018; JORC 2012 Table 1
3	Where the variance lives, by stage	Smee et al., GEEA, 2024
4	What each QC insertion tests	CIM 2018 §7(d); GEEA 2024
5	A batch of twenty, controls placed	CIM 2018 §7(d); practice
6	Anatomy of a reference-material certificate	GEEA 2024; ISO/IEC 17025:2017
7	Bracketing accuracy across the grade range	CIM 2018 §7(d); GEEA 2024
8	Signature: the control-chart triptych	GEEA 2024; CIM 2018 §7(d)
9	The failure-response table	GEEA 2024; CIM 2018 §7(d)
10	Database against spreadsheet	CIM 2018 §6(f), §7(d)

EXHIBIT	SUBJECT	ANCHOR SOURCE
11	Table 1 Section 1 red flags	JORC 2012 Table 1; CIM 2018
12	The QA/QC dashboard	CIM 2018 §7(d); GEEA 2024
13	The non-conformance workflow	CIM 2018 §7(d)
14	The diligence request, mapped to chapters	CIM 2018 §7; JORC 2012 Table 1

How to read the evidence in this paper

This paper is a reading of published guidelines and codes, not a report of measured project data. Its statements fall into three classes, and the class of each is made explicit so a reader relying on the paper for diligence knows exactly how firm each claim is.

Frameworks against statistics

Every exhibit in this paper is a framework instrument: a schematic, a comparison table, a control-chart layout or a workflow. None reports a measured grade or a project result. The plotted values in the control-chart triptych (Exhibit 8) and the targets in the dashboard (Exhibit 12) are illustrative defaults chosen to show how a chart is read or a target is set, and they are labelled as such on each exhibit.

Established-practice rows, flagged

A small number of numeric conventions are widely used industry practice rather than a figure mandated by any single code. Three appear in this paper and each is flagged where it is used: the one-in-twenty (about five percent) per-type insertion rate; the two-standard-deviation warning and three-standard-deviation action limits on the accuracy chart; and the roughly ten percent half absolute relative difference used as a duplicate-precision tolerance. The binding source in each case (the CIM guidelines) states the requirement qualitatively, that controls be inserted frequently enough for statistical confidence and that acceptable accuracy and precision be established, and leaves the exact figure to program design. These conventions should be confirmed against the full recommended-practice review or the project's own data-quality objectives before any exact number is quoted as a code requirement.

Open items

Two cited sources are not archived in full. The recommended-practice review published in Geochemistry: Exploration, Environment, Analysis (Smee and others, 2024) is subscription-gated; its abstract and complete bibliographic record are archived and carry every claim attributed to it, and no exact numeric threshold is attributed to it as a verbatim figure. ISO/IEC 17025:2017 is a paid standard cited by its published title and year; it anchors the accreditation concept, not a number. The three primaries that carry the paper, the CIM 2018 guidelines, the JORC 2012 Table 1 template and the review's record, are held on disk and were read directly.

A note on the Aurus practice references

The three practice notes in Chapter 8 are stated at the altitude the underlying work supports and no further. One describes a delivered data-governance discipline and reports no grade, instrument value or resource statement. Two describe service lines in the terms their stage supports: a designed reporting architecture, and a mandate structured around a QA/QC discipline, neither claimed as an operated program of stated duration or a delivered outcome. No client is named, no project is located beyond the region, and no figure specific to any mandate is disclosed.

How to test this paper

Every figure and framework statement carries a source tag in the underlying document, keyed to the series evidence dossier (rows WP07-01 to WP07-24), so that each can be traced to its guideline, code or review and audited for orphans. The exhibit source index above maps every exhibit to its anchor source. A reader who wishes to verify a claim should read it against the cited primary, two of which are freely available and archived on disk.

REVIEW INSTRUMENT

The assay-release control sheet

A batch is not ready because a laboratory file arrived. It is ready when sample identity, control performance, failures, rework and database loading form one reviewable record.

REVIEW LINE	EVIDENCE TO INSPECT	HOLD CONDITION
Dispatch	Sample sequence, seal record, dispatch sheet, carrier handoff and laboratory receipt	Broken or unexplained custody
Preparation	Method code, crushing and pulverising specification, split method and retained fractions	Method differs from instruction
Blanks	Results against the program's contamination rule, reviewed in batch order	Contamination or carry-over unresolved
Reference materials	Certified value, method compatibility, warning and action limits, time sequence	Accuracy failure unresolved
Duplicates	Duplicate type, pair identity, grade range and precision response	Precision outside the stated objective
Laboratory QC	Laboratory controls, repeats and analyst comments beside client controls	Contradictory control evidence
Re-assay	Failed interval, selected rework, replacement values and reason code	Original and replacement values confused
Umpire work	Selected pulps, second-laboratory method and bias comparison where required	Material inter-laboratory bias open
Database load	Import checksum, validation exceptions, immutable source file and load approval	Unresolved load or identity exception
Release	Named reviewer, release date, batch status and linked non-conformance record	No accountable release decision

Framework instrument derived from the CIM Mineral Exploration Best Practice Guidelines, Section 7, and the JORC Table 1 sampling and assay criteria. Thresholds remain program-specific unless explicitly identified as established practice.

> The result enters the resource workflow only after the controls around it have earned release.

The vocabulary of sampling and assay control

Accuracy · how close a measurement is to the true value; the property a certified reference material tests. Distinct from precision.

Precision · how repeatable a measurement is; the property a duplicate tests. A precise result can still be inaccurate if the process is biased.

Bias · a systematic error that shifts results consistently in one direction; the thing accuracy control exists to detect.

Certified reference material (CRM) · a homogenised sample with a certified grade and uncertainty, established by a round-robin of accredited laboratories; inserted blind to control accuracy.

Blank · material certified to be barren, inserted to detect contamination or carry-over between samples through preparation or analysis.

Duplicate · a second sample taken at a defined stage (field, coarse reject or pulp) to measure the precision of the process from that stage onward.

Umpire (check) assay · a subset of pulps re-assayed at a second accredited laboratory to detect inter-laboratory bias.

Control chart · a time-ordered plot of a control result against a centre line and limits (commonly two and three standard deviations), used to detect drift and batch failures.

HARD · half absolute relative difference; a common measure of the agreement between duplicate pairs, with acceptance often taken near ten percent for most pairs.

Coefficient of variation · the standard deviation as a fraction of the mean; an average value summarises precision across a program.

Fundamental sampling error · the irreducible error that arises when a large mass is reduced to a small one; minimised by adequate sample mass and fine comminution before splitting.

Diamond core drilling · drilling that recovers a continuous rock cylinder that can be logged, measured for recovery and split; the resource-definition reference method.

Reverse circulation (RC) · drilling that returns rock chips to a cyclone and splitter; faster than core, with sample quality dependent on a dry, well-managed return.

Recovery · the proportion of the drilled interval actually returned as sample; logged as a control because low or variable recovery can bias grade.

Pulp · the fine powder produced by pulverising, from which the analytical portion is taken and which is retained for possible re-assay.

Fire assay · a fusion method, typically on a nominal 30 or 50 gram charge, standard for gold, finished by instrument.

Four-acid digestion · a near-total acid digestion for multi-element base-metal work, read by ICP-MS or ICP-OES.

Aqua regia · a partial acid digestion; useful for some elements but not total, so the certified method of a reference material must match the digestion used.

Chain of custody · the recorded sequence of transfers of a sample from field to laboratory, evidencing that the sample assayed is the sample it claims to be.

ISO/IEC 17025 · the standard for the competence of testing and calibration laboratories; accreditation is granted for a defined scope of methods, not as a blanket badge.

Qualified Person / Competent Person · the named professional who takes responsibility for a public technical disclosure under NI 43-101 (QP) or a CRIRSCO-family code (CP).

CRIRSCO family · the international family of reporting codes (JORC 2012, SAMREC, PERC and others) with the CIM Definition Standards; NI 43-101 is a disclosure instrument that incorporates the latter.

AURUS INSTITUTE FOR RESOURCE DEVELOPMENT

Data You Can Bank: Drilling, Sampling and Assay QA/QC as the Foundation of Resource Confidence.

White paper number seven in the Aurus white paper library. Explore pillar. Prepared by the Aurus Technical Committee, July 2026.

Citation policy: every statistic and framework in this paper is attributed in prose or in an exhibit source line to its institution, publication and year. Numeric conventions that are established industry practice rather than a single code’s mandated figure are marked as such where they appear, and the binding qualitative requirement is named. No source, figure or date in this paper is invented.

Typefaces: Fraunces (display), Source Serif 4 (text) and Inter (labels and figures).

Imagery generated and art-directed in-house with the OpenAI image model gpt-image-2, anchored to the series style set; no photograph depicts an identifiable place or person, and no image contains people.

Working draft 1: wave-3 Technical Paper; adversarial certification pending.

MINING · INFRASTRUCTURE · ENGINEERING · ENVIRONMENT