loft

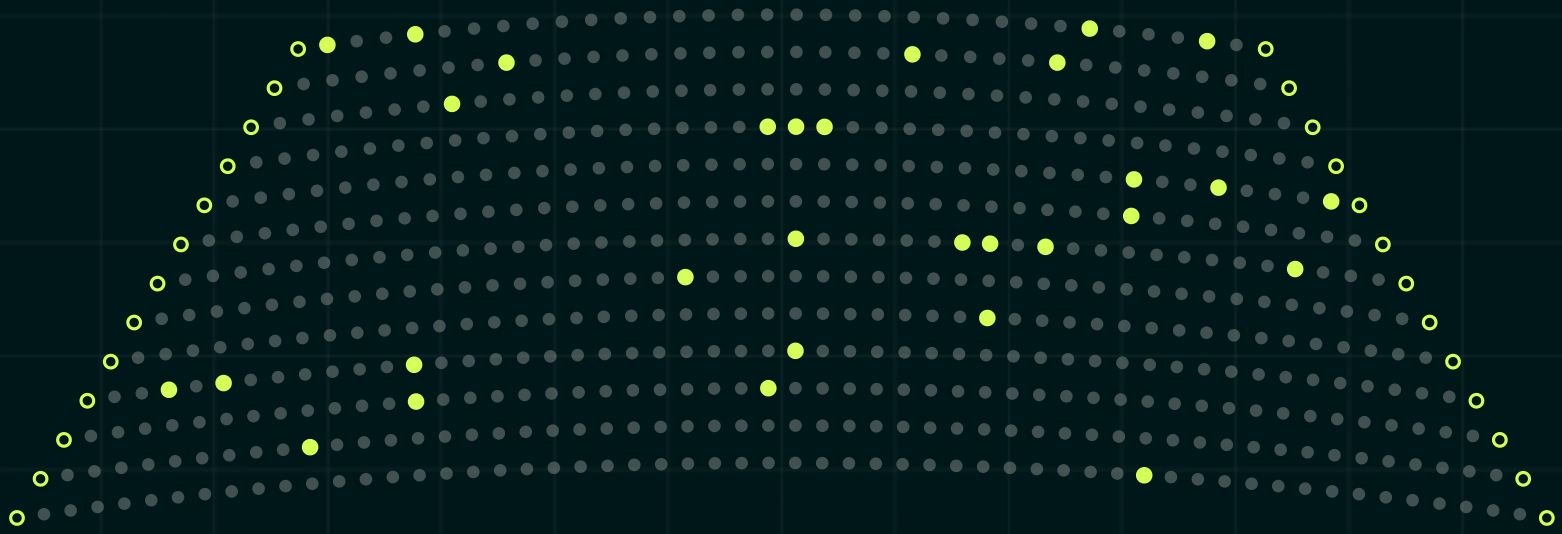
WHITE PAPER · VERSION 1.0

SEATING AND INVENTORY

Structure-aware adaptive sampling for venue seat inventory verification

A risk-based method for establishing confidence in digital seating inventories, with a worked application to the 20,096 seat records of the Akii Bua stadium inventory.

PUBLISHED FOR REVIEW AND FIELD VALIDATION · ALL THRESHOLDS ARE PROVISIONAL



PREPARED BY
Loft Technologies Ltd

PUBLISHED
October 2026

READ ONLINE
loftuganda.tech/research/savs

ABOUT THIS PAPER

PREPARED BY	Loft Technologies Ltd, Kampala, Uganda
STATUS	Published for review and field validation. Thresholds, weights and switching rules are provisional until the validation programme in Section 11 is complete.
CASE DATA	The Akii Bua figures were recomputed from Loft's seat plan v2 for the venue, prepared from the venue's seat placement drawing of August 2026. They illustrate the method and are not a certification of the venue.
CASE STUDY	Akii-Bua Olympic Stadium, Lira City, Uganda: 20,096 seat records in 93 blocks and 909 rows
VERSION	1.0, October 2026

CONTENTS

	Executive summary	3
1	Introduction	4
2	The verification problem	5
3	Why random sampling of seats is not enough	6
4	Statistical foundations	7
5	The SAVS framework	8
6	Building the inspection plan	10
7	Findings, escalation and acceptance	13
8	Case study: Akii Bua	17
9	Evidence, states and certification	20
10	Implementation in TKT	24
11	Validation programme	25
12	Limitations	26
13	Conclusion	27
A	Formulas	27
B	Glossary	29
	References	30

HOW TO CITE

Loft Technologies (2026). Structure-Aware Adaptive Sampling for Venue Seat Inventory Verification. Technical White Paper, version 1.0. Kampala: Loft Technologies Ltd.

Read online and download: loftuganda.tech/research/savs

296	21%	909	1,592
RANDOM SEATS THAT FIND A 1% ERROR RATE 95% OF THE TIME	THE CHANCE THE SAME SAMPLE FINDS A 16-SEAT NUMBERING SHIFT	ROWS, EACH TESTED BLIND FOR ITS END LABELS AND CHAIR COUNT	BASE FIELD TASKS IN THE ILLUSTRATIVE AKII BUA PLAN

ABSTRACT

A digital seat that does not match a physical seat can be sold twice, sold when it does not exist, or send a spectator to the wrong place. This paper proposes Structure-Aware Adaptive Sampling for Venue Seat Inventory Verification (SAVS): a method for establishing, before a seating inventory goes on sale, that its seats correspond to the venue at a stated confidence, without inspecting every seat.

SAVS combines five inspection mechanisms: statistical random sampling of seats with a guaranteed minimum inclusion probability; stratification by venue structure; deterministic, blind structural tests of blocks and rows; direct and risk-weighted inspection of known exceptions; and adaptive escalation when a discrepancy is found.

The case for structure rests on one exact property. A simple random sample of seats detects a defect with a probability that depends only on how many seats the defect affects, not on where they are. That is adequate for widespread error but weak for the errors that matter most in venues: a missing chair, a shifted numbering sequence, a misnumbered skybox, and physical seats absent from the digital list, which a sample drawn from that list can never select. In the 20,096–record Akii Bua inventory, a random sample of 296 seats detects a 1% defect rate with 95% probability but finds a numbering shift affecting 16 seats only 21% of the time, while one blind test of a row's end labels and chair count checks the row's whole numbering sequence.

For a first acceptance, SAVS therefore tests every block and every row structurally, samples seats for the defects that structural tests cannot see, inspects known exceptions directly, and escalates by rule. The paper defines the framework, its acceptance rules and the exact meaning of a "Sample Verified" result, applies it to Akii Bua, and specifies the simulation and field trials needed to calibrate its thresholds before any certification claim is made.

KEYWORDS
seat inventory verification; reserved seating; acceptance sampling; discovery sampling; stratified sampling; adaptive sampling; ticketing; venue operations

Executive summary

SAVS lets an operator accept a venue's digital seat inventory for sale on documented, reproducible evidence, at a stated confidence, without a seat-by-seat survey.

Three numbers from Akii Bua

20,096	seat records in the Akii Bua inventory, in 93 blocks and 909 rows
1.5%	chance that a 296–seat random sample finds one missing chair
864 of 909	rows to test at random for a 95% chance of finding one defective row

- **Why random sampling alone fails.** Detection depends on how many seats a defect touches. Small but serious defects (a missing chair, a shifted row, a misnumbered skybox) are rarely found, and seats missing from the digital list can never be found by sampling that list.
- **The design principle: census the structure, sample the seats.** For a first acceptance, every block and every row gets a blind structural test of its end labels and chair count. Seats are then sampled for the defects structural tests cannot see.
- **A guarantee that holds wherever defects are.** Seats are drawn at random within each block, and every seat's chance of selection is at or above a floor set by the assurance target. Risk weighting can raise that chance, never lower it.
- **Known problems are inspected, not sampled.** Every exception found by digital validation or raised by the venue gets its own task.
- **Escalation follows published rules.** A finding triggers its neighbourhood, its row, a tightened block sample and, if needed, a full inspection of the block.
- **"Sample Verified" is a precise claim.** All blocks and rows passed their structural tests, all exceptions are closed, and a seat sample that finds no defect bounds the remaining seat-level error rate, for example to about 1% at 95% confidence. It never means that every seat was inspected.
- **Akii Bua, illustrative plan.** About 1,600 field tasks (103 exception tasks, 930 structural tests, 487 sampled seats and 72 independent re-inspections). A full survey would read every label in the same 909 rows; the field trial will show which costs less. The row census also yields a measured physical seat count for the open capacity question (19,996 sellable seats drawn against 20,009 in the drawing's table).
- **Status.** All thresholds are provisional. Simulation and a field trial must calibrate them before certification claims are made.

> **Census the structure, sample the seats.**
>
> Structural units are few and each structural test checks a whole sequence, so testing all of them is affordable and decisive. Seats are many and each check sees one seat, so seats are sampled, with a guarantee that holds wherever the errors are.

§ 1 Introduction

Digital reserved-seating systems rest on one assumption: every sellable digital seat corresponds to the correct physical seat in the venue.

That assumption is usually inherited from architectural drawings, CAD files, spreadsheets, contractor schedules, legacy ticketing databases or hand-made seating plans. A digital seat record does not by itself establish that the physical seat exists, and physical existence does not establish that the seat is represented correctly. A seat may:

- not exist;
- exist under a different number;
- appear more than once in the digital inventory;
- belong to a different row or section;
- be assigned to the wrong gate;

- be inaccessible, obstructed or unusable;
- carry an unreadable or incorrect label;
- be physically installed but absent from the inventory;
- sit in a numbering sequence that differs from the digital plan.

Each of these is an operational and commercial risk: duplicate sales, disputes at the gate, spectators routed through the wrong entrance, and capacity figures that do not match the venue. A systematic verification step is therefore needed between the import of a digital seating inventory and its acceptance as a production seat plan.

Complete physical verification of every seat gives the strongest assurance, but it is costly in venues of tens of thousands of seats and must be repeated whenever the inventory changes. Simple random sampling reduces the effort but, as Section 3 shows, misses exactly the errors that matter most. This paper proposes a middle path, Structure-Aware Adaptive Sampling (SAVS), and shows how it can be built into TKT, Loft's ticketing platform.

§ 2 The verification problem

Verification asks whether a digital inventory corresponds to the physical venue closely enough for reserved-seat sales and access control. It is a detection problem before it is an estimation problem.

A venue holds N digital seat records, each carrying attributes such as:

$$s_i = (\text{stand}, \text{tier}, \text{block}, \text{row}, \text{seat}, \text{gate}, \text{category}, x_i, y_i, \ldots)$$

This paper uses the following definitions.

- **Digital inventory:** the seat records of one inventory version.
- **Physical configuration:** the chairs, labels and signage installed at the time of inspection.
- **Correspondence:** a one-to-one mapping between in-scope digital records and physical seats that preserves identity (block, row and seat label) and the attributes the operation relies on (usability, view, category and gate).
- **Discrepancy:** any departure from that mapping: a record with no seat, a seat with no record, two records on one seat, or a record mapped to a seat that carries a different label.
- **Material discrepancy:** one that can cause a duplicate or impossible sale, a wrong seat assignment, a wrong entry route or a wrong capacity figure.

Two populations are in play: the digital records and the physical seats. Verification must find records without seats and seats without records, and a sample drawn from the digital list can only ever see the first.

The principal question is therefore not only "what percentage of seats is wrong?" The operational questions are whether there is a material error anywhere in the venue; whether whole rows or sections are represented incorrectly; whether numbering sequences, row boundaries and numbering directions are correct; whether

seats are assigned to the correct block and row; whether errors are concentrated in particular areas; and whether the remaining risk is low enough to sell reserved seats. A method must answer for individual defects and for structural defects, and state what it can claim about each.

§ 3 Why random sampling of seats is not enough

A simple random sample of seats gives an exact guarantee, but only about how many seats are wrong, never about where. For venue inventories, that guarantee misses the defects that matter most.

What a random seat sample guarantees. If D of N seat records are defective and n records are drawn at random, the chance of finding at least one defect, assuming that a check of any affected seat reveals it, is:

$$P(\text{detect} \geq 1) = 1 - \frac{\binom{N-D}{n}}{\binom{N}{n}}$$

The formula depends on D alone. It does not matter whether the defective seats are scattered or packed into one row: clustering does not weaken a random sample. What clustering changes is how much one well-placed test can find, and many material defects involve very few seats.

Small defects are found rarely. At Akii Bua (N = 20,096), a 296-seat sample finds a missing chair 1.5% of the time, a misnumbered skybox (12 seats) 16% of the time and a numbering shift affecting 16 seats 21% of the time. Quadrupling the sample to 1,168 seats raises those figures only to 6%, 51% and 62% (Figure 1).

A 296-seat sample finds a 1% error rate 95% of the time, but a 16-seat shift only 21%

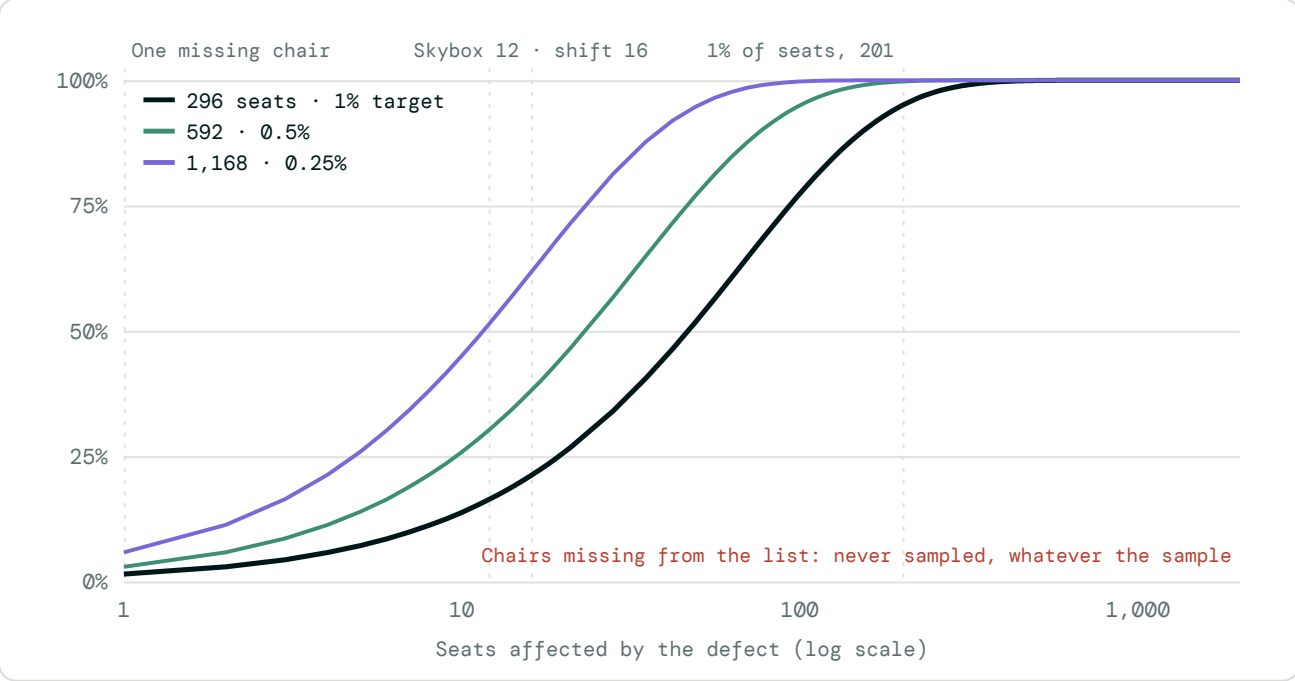


Figure 1. Chance that a simple random sample of seats finds at least one affected seat, by the number of seats a defect affects, for the samples that meet 1%, 0.5% and 0.25% targets at 95% confidence, and one you set (lime; 1,005 seats is "inspect 5%"). Exact hypergeometric probabilities for the Akii Bua v2 inventory (20,096 seat records).

One observation can test a whole sequence. Suppose a row is recorded as A1 to A30 and physical seat A7 is missing, so every later label has shifted. Random checks of A2, A4 and A6 find nothing. One check of the row's last seat, or one count of its chairs, finds the problem at once.

Seats missing from the list can never be drawn. A sample drawn from the digital inventory cannot select a physical seat that is not in it: a row extension, an unrecorded back row, a temporary stand. Only a physical walk of the venue's structure finds them.

A venue-wide sample says little about any one block. 296 seats spread over 93 blocks is about three per block, and a 12-seat skybox expects 0.18 sampled seats. A clean venue-level result bounds the venue-wide error rate; it does not show that any particular block is correct.

Known ambiguity should not be left to chance. When digital validation or the venue has already flagged an uncertain row or area, inspecting it directly is cheaper and more certain than hoping the sample reaches it.

Random seat sampling stays the right tool for one job: bounding the rate of seat-level errors that structural tests cannot see, such as a wrong label in mid-row, a broken chair or a wrong category. SAVS keeps it for that job and adds the tests that do the rest.

§ 4 Statistical foundations

SAVS introduces no new statistics. It combines established sampling methods in a way fitted to how venue errors behave.

Acceptance sampling. ISO 2859-1:2026 sets sampling schemes for a continuing series of lots, indexed by an acceptance quality limit, with switching rules between normal, tightened and reduced inspection as evidence accumulates. ISO 2859-2:2020 covers isolated lots, indexed by limiting quality. One inventory version resembles an isolated lot, so SAVS borrows limiting-quality thinking for a single acceptance and the switching idea across versions and venues (Section 7.4). Neither standard defines stadium seat verification; SAVS adopts their principles, not their tables.

Discovery sampling with zero acceptance. When any defect matters, the sample is sized to find at least one defect with high probability if the defect rate is at or above a stated level, and the lot is accepted only if none is found, an acceptance number of zero (Squeglia, 2008). Section 6.2 gives the calculation.

Bounds from zero findings. Finding no defects in n random seats does not show that there are none; it bounds the rate. The one-sided 95% upper bound is close to $3/n$, the "rule of three" (Hanley and Lippman-Hand, 1983), and is computed exactly from the binomial (Clopper and Pearson, 1934) or the hypergeometric distribution.

Stratification. Dividing a population into strata and sampling each one guarantees coverage of every stratum and usually improves precision (Neyman, 1934; Cochran, 1977). SAVS stratifies by block.

Unequal probabilities. Units may be selected with different probabilities, provided each is known and greater than zero; estimates then weight each unit by the inverse of its probability (Horvitz and Thompson, 1952). Risk weighting in SAVS is unequal-probability sampling with a floor.

Adaptive sampling. In adaptive designs, what is observed changes what is selected next, which pays where the feature of interest is clustered (Thompson, 1990; Thompson and Seber, 1996). Adaptively added units are not a random sample, so an error rate computed naively over everything inspected overstates the true rate. SAVS therefore keeps adaptive additions out of its statistical statement (Section 5.1).

Rectifying inspection. When a lot fails its sample, classic schemes inspect it in full and correct it, which limits the quality that leaves inspection, the average outgoing quality of Dodge and Romig (1959). SAVS applies the same logic per block: a failing block is inspected and corrected, not merely rejected.

Sequential decisions. Deciding as evidence arrives can reach the same error rates with fewer observations than a fixed sample (Wald, 1945). The SAVS escalation ladder is a simple sequential rule; formal sequential designs are a candidate for a later version.

§ 5 The SAVS framework

A SAVS inspection plan is the union of four samples, each allowed to support a different kind of claim, run in five layers.

5.1 Four samples, five layers

$$S = S_R \cup S_S \cup S_K \cup S_A$$

Table 1. The components of a SAVS plan and what each may claim

COMPONENT	LAYER	HOW UNITS ARE CHOSEN	WHAT IT MAY CLAIM
S_R **Random statistical sample**	C	Seats drawn at random within each block; every seat's inclusion probability known and at or above a floor	An upper bound on the seat-level defect rate
S_S **Mandatory structural sample**	B	Every block; every row for a first acceptance, or random rows for a re-verification	A census statement, or a row-level bound
S_K **Risk-weighted sample**	A and C	Every registered exception, inspected directly; extra seats drawn at a raised, known rate in elevated blocks	Each exception closed; the extra seats join the probability sample
S_A **Adaptive sample**	E	Triggered by a finding, under the escalation rules	The extent of each defect, found and corrected
Re-inspection	D	A random share of completed tasks, repeated blind by a second inspector	The reliability of the observations themselves

Three rules keep the claims honest:

1. Every inspected unit records why it was selected.
2. Only units drawn with known probabilities enter the statistical statement. Exception tasks and adaptive additions support closure statements, never the error-rate bound.
3. A statement is made only at the level its evidence supports: a census statement for what was tested in full, a bound for what was sampled.

5.2 Venue hierarchy and inspection units

SAVS treats a venue as a hierarchy with a test at each level, not as a flat list of seats. A typical hierarchy runs venue, stand, tier, block (or section), row, seat; the levels are configurable because venues name them differently.

Every block and row is tested; only seats are sampled

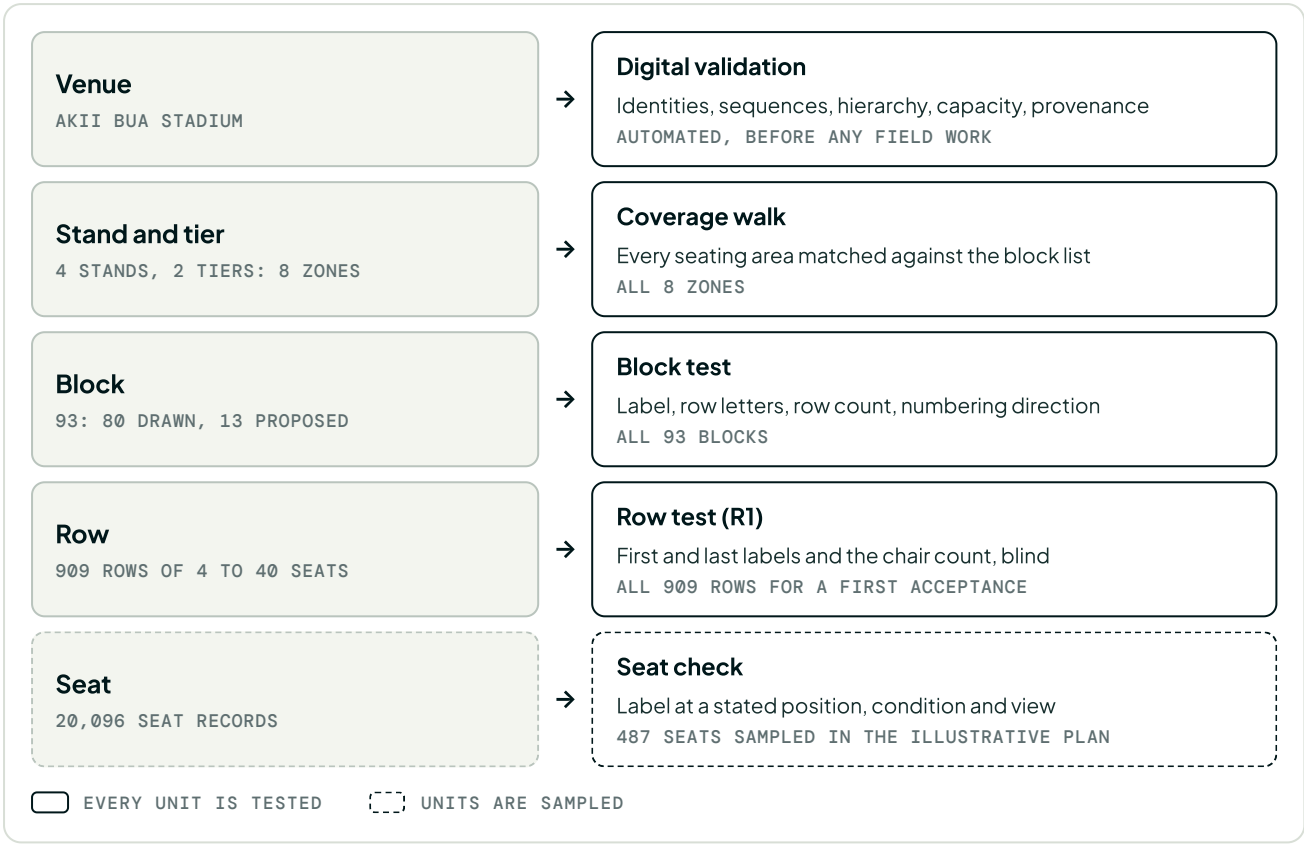


Figure 2. The venue hierarchy with Akii Bua's count at each level and the test that checks it. Only the seat level is sampled; every level above it is tested in full for a first acceptance.

The sampling engine must know each seat's row, each row's block and each block's stand, and also the order inside them: rows from front to back, seats along each row, the numbering direction, aisle positions and which blocks adjoin. That order is what makes blind tests possible, because a position can be described physically ("the third chair from the aisle on your left, facing the pitch") without naming the expected label.

Table 2. Inspection units and what each finds

TEST	UNIT	THE INSPECTOR RECORDS	IT FINDS
Coverage walk	Stand and tier zone	Every seating area seen, with its label, against the block list	Unrecorded areas, missing blocks
Block test	Block	Block label and signage, front and back row letters, number of rows, numbering direction	Wrong block label, missing or extra rows, reversed numbering
Row test (R1)	Row	First label, last label and chair count, blind	Shifts, truncation, extension, missing or extra chairs, wrong row letter
Row read (R2)	Row	Every label in order, with gaps, duplicates and unreadable labels	All that R1 finds, plus mid-row swaps and duplicates
Seat check	Seat	The label found at a stated position, the chair's condition, view and category marking	Mid-row label errors, unusable chairs, restricted view, wrong category

TEST	UNIT	THE INSPECTOR RECORDS	IT FINDS
Special test	Registered exception	What the exception needs: a wheelchair space and companion seat, a view photograph, a block boundary	The registered question

§ 6 Building the inspection plan

A SAVS plan is built in a fixed order: validate the data, size the seat sample, allocate it across blocks, add the structural census, make every task blind, and weight the sample by risk.

6.1 Pre-sampling digital validation

Physical inspection starts only after automated checks of the imported data, so that no field effort is spent finding what a computer can find. The checks below are the minimum.

- **Unique identifiers.** Every sellable seat has a unique canonical identifier, for example W101-H-17. A duplicate is an immediate exception, never left for the field to discover.
- **Row sequences.** Missing or duplicated numbers, unexpected resets, unusual starting numbers, discontinuities, and row lengths out of line with neighbouring rows.
- **Hierarchy consistency.** Every seat references a valid stand, tier, block, row and gate.
- **Capacity reconciliation.** Seat records are summed and compared with every capacity the sources declare: row, block, stand and venue, and by seat model or category where the drawing gives them.
- **Spatial anomalies.** Where coordinates exist: overlapping seats, unusual gaps, seats outside their block's outline, abnormal spacing, and ordering that contradicts the numbering direction.
- **Provenance.** Which records come straight from the source, and which were inferred, joined or proposed during import.

Validation produces two outputs. The **exception register** holds one entry per finding, with its class (Section 7.1), the question it raises and who can answer it: the venue, the drawing's author or a field inspection. The **risk inputs** per block feed the weighting in Section 6.6. Exceptions are never fixed silently: each is either corrected with a recorded decision or carried into physical verification as a task.

Loft's drawing conversion toolkit already runs the first four checks. The Akii Bua v2 files pass the identifier, sequence and hierarchy checks with one documented exception (S233 row X), and the capacity check finds a 13-seat gap that is registered in Section 8.2.

6.2 Sizing the seat sample

The seat sample is sized from an assurance target fixed before any seat is drawn: find at least one defect with probability $1 - \beta$ whenever the defect rate is p^* or more. Read the other way, a sample that finds nothing shows that the rate is below p^* at confidence $1 - \beta$.

Table 3. Seat sample sizes for 95% confidence, Akii Bua inventory (N = 20,096)

TARGET DEFECT RATE p^*	DEFECTIVE SEATS AT THAT RATE	EXACT SAMPLE (HYPERGEOMETRIC)	BINOMIAL APPROXIMATION	SHARE OF SEATS INSPECTED
2%	402	148	149	0.7%

TARGET DEFECT RATE P^*	DEFECTIVE SEATS AT THAT RATE	EXACT SAMPLE (HYPERGEOMETRIC)	BINOMIAL APPROXIMATION	SHARE OF SEATS INSPECTED
1%	201	296	299	1.5%
0.5%	100	592	598	2.9%
0.25%	50	1,168	1,197	5.8%
0.1%	20	2,795	2,995	13.9%

- **The sample barely depends on venue size.** It is close to $3/p^*$ seats, so a 40,000–seat venue needs almost the same sample as Akii Bua for the same target.
- **Halving the target rate doubles the sample.** Assurance at low defect rates is expensive, which is why SAVS does not rely on seat sampling for structural defects.
- **No observed errors does not mean no errors.** If 296 random seats are all correct, the 95% upper bound on the defect rate is 1.00%; at 592 seats it is 0.49%, and at 1,168 seats 0.24%.

The target applies to critical and major defects together; minor defects are recorded but do not fail a plan (Section 7.1). The calculation assumes the inspector recognises every defect seen. If an inspector catches a real defect with probability s, the effective sample is about $s \times n$, so the sampling floor in Section 6.3 is divided by s. Re-inspection gives a lower bound for s (Section 6.5).

6.3 Stratified allocation with a floor

SAVS samples seats block by block and sets a floor on every seat's chance of selection, not only a minimum count per block. The floor is what lets the venue-level guarantee survive risk weighting. The usual allocation spreads a total sample n in proportion to block size N_h and risk weight w_h:

$$n_h = n \, \frac{N_h \, w_h}{\sum_j N_j \, w_j}$$

On its own, this lets heavy weights in some blocks starve others. SAVS instead fixes a base inclusion probability π_0 and allocates:

$$n_h = \max(m, \lceil \pi_0 \, w_h \, N_h \rceil), \qquad w_h \geq 1, \qquad \pi_0 = 1 - \beta^{1/D^*}$$

where m is a minimum per block (default 2) and $D^* = p^* \, N$ is the number of defective seats the target must catch. With seats drawn by simple random sampling inside each block, independently between blocks, every seat's chance of selection is at least π_0, and for any D^* defective seats, wherever they sit, the chance of missing them all is at most $(1 - \pi_0)^{D^*} = \beta$ (Appendix A). Risk weights can only add inspection; they can never weaken the guarantee or leave a block unvisited.

The guarantee depends on that design. A plan that selects seats in groups, such as whole rows, can keep every seat's chance above the floor and still miss a defective row far more often, so systematic and spatially balanced selection are not used for the seat sample.

At Akii Bua, for a 1% target at 95% confidence, $\pi_0 = 1.479\%$. The sampling frame, the seats not already read in full under Layer A, holds 19,556 seats in 80 blocks. With $m = 2$ and a weight of 2 in the 43 elevated blocks, the sample is 487 seats, 2 to 14 per block; without weighting it would be 326. Placing 201 defective seats where they are least likely to be found still leaves a 96.0% chance of finding at least one, if inspectors miss nothing.

6.4 The structural census

SAVS censuses the structure and samples the seats. Structural units are few (Akii Bua has 93 blocks and 909 rows against 20,096 seats, 22 seats to a row on average) and each structural test checks a whole sequence, so testing all of them is affordable and decisive.

Rows. The row test records the first label, the last label and the chair count, blind. One test catches a shift, a truncation, an extension, a missing or extra chair and a wrong row letter anywhere in the row. When the inspector also records which physical end carries the first label, it checks the numbering direction.

Why every row, not a sample of rows. A single defective row is material, and a random sample of rows finds one only in proportion to its size: catching one bad row among 909 with 95% probability takes 864 row tests. Testing all 909 costs little more and leaves no row untested, inspector error aside. A row sample is enough only when the target is a share of rows: 257 random rows find a problem affecting 1% of rows (9 rows) with 95% probability. SAVS therefore tests every row for a first acceptance and samples rows only for re-verification with a clean history (Section 7.4).

Blocks. Every block gets a block test: label and signage, the front and back row letters, the number of rows and the numbering direction. This catches block-level errors, such as a swapped label or an unrecorded back row, in one visit.

Coverage. A walk of each stand and tier matches every seating area seen against the block list and records any area the inventory lacks, such as press seating or a temporary stand.

Positions inside rows. Seat checks drawn in the sample cover interior positions, and registered exceptions cover numbering transitions, aisle crossings, unusual gaps and block boundaries. A row recorded as seats 1 to 35 is thus tested at both ends and, if it falls in the seat sample, at a random interior position, rather than at three arbitrary interior seats.

The field trial must time each kind of test. If reading every label proves little slower than counting chairs, full enumeration becomes the better plan for a first acceptance, and SAVS sampling applies to re-verification and to changed areas only.

6.5 Blind verification and independent re-inspection

SAVS tasks are blind by default: the inspector records what is there before seeing what the inventory expects. A directed instruction such as “Confirm that Row H contains seats 1 to 34” hands the inspector the answer and invites confirmation bias. A blind task reads instead: *Locate Block W101, Row H. Record the first visible seat number, the last visible seat number and the total physical seat count.* The application compares the observation with the inventory only after it has been submitted.

- **Blind by default:** row ends and counts, block labels, numbering direction, unusual sequences, and seat checks, where the task names a position and the inspector reports the label found there.
- **Directed only where a blind question is impossible,** for example “Is the view from W101–H-27 obstructed? Photograph it.”

- **Enforced by the application:** expected values stay hidden, an observation locks on submission, and any later change is kept as a revision with its reason.

Independent re-inspection (Layer D). A random 5% of completed structural tests and seat checks are repeated blind by a second inspector who does not see the first result, and a disagreement is settled by a third observation or a photograph. Because both inspectors record values rather than verdicts, disagreements measure the recording error rate e whatever the defect rate. A real defect is missed only when an error happens to reproduce the inventory's value, so the sensitivity s in Section 6.2 is at least $1 - e$. Directed yes-or-no tasks are different: at low defect rates two inspectors agree even if neither notices anything, so their sensitivity is measured on known cases from the exception register. If agreement falls below a provisional 98%, that inspector's work is re-inspected at a higher rate.

Separation of duties. Inspectors should not be the people who built or imported the digital inventory. The issued plan stays hidden from the venue and the import team until field work ends, so nobody can correct only the seats that will be checked.

6.6 Risk weighting

Risk weighting sends more inspection where an error is more likely or would cost more, under rules fixed before sampling so that the plan cannot be steered after the fact. The likelihood of error in block h can be scored from the validation outputs:

$$L_h = a_1 D_h + a_2 C_h + a_3 M_h + a_4 H_h + a_5 U_h$$

where D_h counts detected digital anomalies, C_h measures capacity discrepancy, M_h marks manual modification or inferred data, H_h records past verification failures and U_h marks unresolved venue or drawing ambiguity. The coefficients are set by the verification policy. Consequence is scored separately, because an unlikely error in an accessible or VIP area can matter more than a likely one elsewhere. For auditability, SAVS turns scores into tiers rather than using raw scores.

Table 4. Risk tiers

TIER	APPLIES WHEN	TREATMENT
Exception	A registered question about a specific unit	Inspected directly in Layer A, by a special test or a full row read
Elevated	The block holds an exception, inferred or joined data, a capacity difference, past failures, or high-consequence seats (VIP, premium, accessible)	Seat sampling weight $w_h = 2$
Standard	No flags	$w_h = 1$

Weights only raise selection probabilities above the floor π_0. At Akii Bua, the illustrative plan elevates 43 of the 80 drawn blocks: those holding registered exceptions, joined or renamed rows, restricted-view seats or possible companion seats. Elevating on consequence as well, for VIP and premium categories, would add more.

§ 7 Findings, escalation and acceptance

Every finding is classified by consequence, the class decides how far inspection escalates, and fixed acceptance rules decide when an inventory may be published.

7.1 Defect classification

Table 5. Defect classes and their effect

CLASS	EXAMPLES	EFFECT
Critical	A seat identity that could be sold twice; a seat on sale for a live event that does not exist; a materially wrong block or row assignment; a numbering shift across a row or more; an error that would route spectators wrongly or unsafely	Immediate escalation to full block inspection (E4); affected seats withheld from sale until corrected and re-verified
Major	A wrong seat number; a missing or unexpected physical seat found before sale; a wrong row boundary or count; a significant labelling discrepancy	Escalation from the neighbourhood upwards (E1)
Minor	A partly damaged label; a cosmetic numbering issue where identity stays unambiguous; a non-critical metadata inconsistency	Logged and corrected in the next version; no escalation

Condition is not correctness. A broken chair that is correctly represented in the inventory is not an inventory defect. It is recorded as a condition finding and the seat is made unavailable for that reason, without triggering escalation. Keeping the two apart stops maintenance issues from inflating the measured error rate, and stops inventory errors from hiding among maintenance issues.

7.2 Adaptive escalation

Inspection grows where defects are found, by fixed rules, until the defect's extent is known, it is corrected, and the block's assurance is restored. This is the principal difference between SAVS and a static sample.

A finding climbs only as far as the evidence pushes it

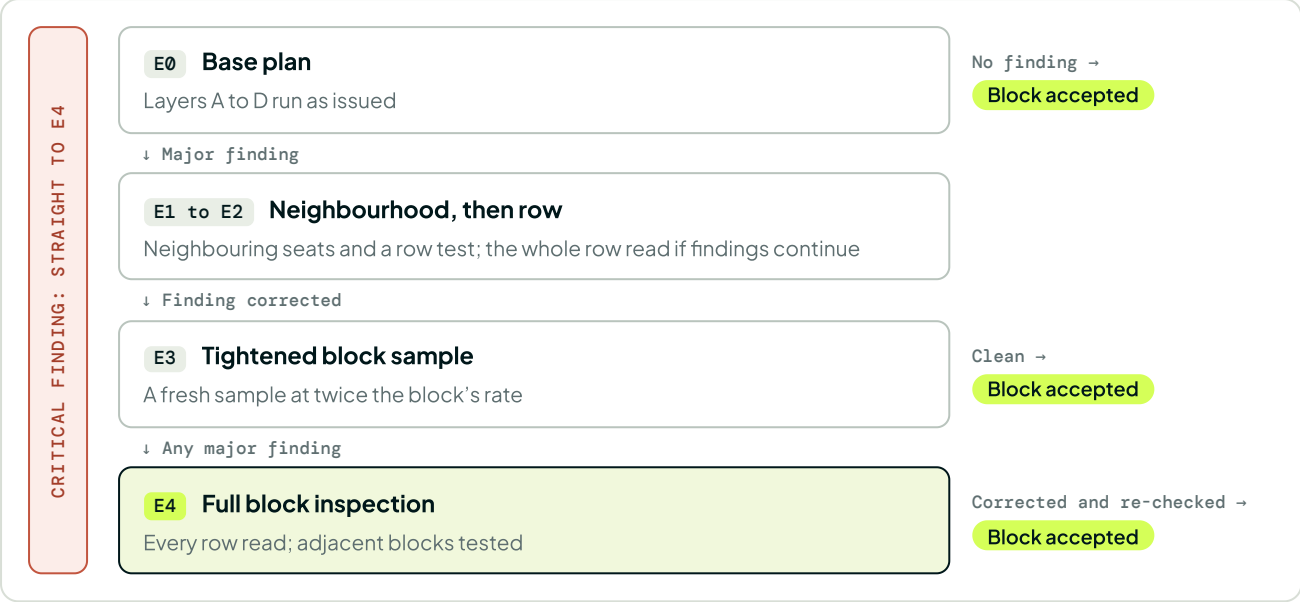


Figure 3. The escalation ladder for one block. A critical finding goes straight to full block inspection; a major finding climbs one step at a time; minor and condition findings do not escalate. All thresholds are provisional.

For example, if seat G17 in W101 fails, the inspector checks G16 and G18 and tests row G's ends and count. If those are clean, G17 is corrected and the block takes a tightened re-sample. If G16 is also wrong, the whole of row G is read and rows F and H are tested. A failure in the tightened sample sends W101 to full inspection.

- **An isolated finding still costs the block its clean result.** After correction, a fresh sample at twice the block's rate, plus row tests on any untested rows, must come back clean before the block is accepted.
- **Block errors spill over.** Full inspection of a block includes block tests on its neighbours, because a swapped label or an offset in row lettering tends to cross a boundary.
- **Venue-level switching.** If three or more blocks in one stand reach full inspection, or a tenth of all blocks reach the tightened sample, the plan for the rest of that stand, or the venue, is re-issued at the Tightened level (Section 7.4). The statistical statement stays at the issued target until the re-issued plan has been drawn and passed. This mirrors the switch from normal to tightened inspection in ISO 2859-1.
- **No de-escalation inside a verification.** An escalated block is accepted only after its corrections are re-verified.

WORKED EXAMPLE (HYPOTHETICAL)

The blind seat check in E126 row C asks for the label on the seventh chair from the left, facing the pitch. The inventory expects 21; the inspector records 20. Under E1, the neighbouring chairs read 22 and 21, so labels 20 and 21 have been swapped. The row test of E126 C had passed, because a swap leaves the end labels and the count unchanged: this is the kind of defect only the seat sample can see.

Nothing else is found, so the finding is isolated. The two labels are corrected, and the block takes a fresh sample of 13 seats, drawn at twice its rate of 7, which comes back clean; the block is accepted. Because the probability sample did find a major defect, the report withdraws the "about 1%" wording: with one finding, the bound computed from the floor (Appendix A) becomes at most 318 of the 19,556 frame seats, about 1.6%. The census statement for blocks and rows is unaffected.

Adaptive tasks establish the extent of defects; they never enter the error-rate bound (Section 5.1). All thresholds here are provisional until simulation and field trials calibrate them (Section 11).

7.3 Error processes and the tests that find them

SAVS assumes that many seat-plan errors arrive in runs, an assumption that hypothesis H2 in Section 11 tests. If they do, finding one changes the odds for its neighbours, and each kind of run has a test that finds it directly.

Table 6. Error processes, their size, and the test that finds each

ERROR PROCESS	WHAT HAPPENS	SEATS AFFECTED	A RANDOM SEAT SAMPLE FINDS IT	SAVS TEST THAT FINDS IT
Numbering shift	A chair is missing and every later label moves down one	Tens (16 from seat 7 of a 22–seat row)	Sometimes	Row test: last label and count
Row truncation	The row ends before its recorded end	1 to a few	Rarely	Row test: last label and count
Row extension	Chairs exist beyond the recorded end	None in the list	Never	Row test: count
Row letter error	A whole row carries another letter	A row (22 at the median)	Sometimes	Row test, blind label
Section displacement	Seats belong to the neighbouring block	A run (10 at E124 rows G and H)	Rarely	Block-boundary special test
Installation omission	A run of planned chairs was never installed	The run	Rarely to sometimes	Row test: count
Drawing interpretation error	A drafting or transcription error across an area	Up to a block (222 at the median)	Usually	Block test and row tests
Two records on one chair	One chair carries two identities	2 records	Rarely	Digital validation, or a row count below the record count
Mid-row label swap	Two neighbouring labels are reversed	2	Rarely	Seat sample only

Figure 1 gives the random sample's chance for each defect size. Every clustered process has a structural signature, which is why the structural census does most of the detection work. The last row has none: isolated label errors inside a row are found only by sampling seats, which is the job SAVS keeps for the seat sample.

7.4 Acceptance rules and assurance levels

An inventory passes when four conditions hold together; the assurance level decides how much evidence each must rest on.

1. **Exceptions closed.** Every registered exception is closed by a field result or by a recorded decision, for example seats withheld from sale until the venue answers. Reports count the two kinds of closure separately.
2. **Structure passed.** Every block test and row test in the plan has passed, or its failure was corrected and the escalated scope re-verified.

3. **Seat sample clean.** No critical or major defect remains unresolved, and every block accepted on its sample has a clean final sample, tightened where escalation required it.
4. **Inspection reliable.** Independent re-inspection agrees at or above the threshold, provisionally 98%.

The headline bound in Section 9.4 holds only if the probability sample found no critical or major defect at all. If it found some, they are corrected and escalated as usual, and the bound is recomputed from the sample as drawn with those findings counted (Appendix A). That bound is weaker, and the report says so.

Table 7. Assurance levels

LEVEL	USED FOR	STRUCTURE	SEAT SAMPLE AT THE FLOOR (AKII BUA, BEFORE RISK WEIGHTING)
Reduced	Re-verifying an accepted version with a clean history and no reported physical change	Every block; random rows sized to find a problem in 1% of rows (257 of 909)	2% target (195 seats)
Standard	First acceptance of a new inventory, or after a structural change	Every block and every row	1% target (326 seats)
Tightened	After venue-level switching, or for high-consequence events	Every block and row; full row reads in escalated stands	0.25% target (1,178 seats)
Full enumeration	When the venue or a regulator requires it, or tightened inspection keeps failing	Every seat checked in full	None needed

Registered exceptions are inspected in full at every level. Movement between levels follows evidence, as in ISO 2859-1:

- **Standard to Tightened:** the venue-level switching rule in Section 7.2.
- **Tightened to Standard:** the next verification passes with no block escalated beyond the neighbourhood step.
- **Standard to Reduced:** two consecutive accepted verifications with no major or critical findings and no reported physical change.
- **Reduced to Standard:** any major or critical finding, a reported physical change, or a new version that changes structure.

§ 8 Case study: Akii Bua

The worked example uses Loft's seat plan v2 for the Akii-Bua Olympic Stadium, a 20,000-capacity stadium under construction in Lira City, Uganda (The Independent, 2024; Uganda Radio Network, 2026). The plan was built on 29 October 2026 from the venue's seat placement drawing of August 2026.

8.1 The inventory

The inventory holds 20,096 seat records in 93 blocks and 909 rows. Every figure in this section was recomputed from the v2 seat list for this paper.

Table 8. The Akii Bua inventory by stand

STAND	BLOCKS	ROWS	SEATS	PROPOSED GATE ASSIGNMENT (SEATS)
West	31	212	4,158	Gate 1 West (3,688); VIP Gate (470)
North	23	261	5,712	Gate 2 North (5,712)
East	16	175	4,529	Gate 3 East (4,529)
South	23	261	5,697	Gate 4 South (5,697)
Total	93	909	20,096	5 gates

- **Tiers:** 13,073 upper-tier and 7,023 lower-tier seats.
- **Blocks:** 80 drawn blocks plus 13 areas with proposed numbering (12 skyboxes and the VVIP box). Blocks hold 12 to 445 seats (median 222) in 2 to 16 rows (median 8).
- **Rows:** 871 drawn rows plus 38 in the proposed areas, holding 4 to 40 seats (median 22).
- **Flags:** 100 restricted-view seats in 23 runs across 11 blocks; 44 seats in 6 rows that look like wheelchair spaces with companion seats, to be confirmed; 173 seats whose numbering is proposed rather than drawn.

The gates need care. The drawing shows none, so the five gates are Loft's proposed assignment by stand, and their positions are placeholders. Verification treats them as a planning attribute: the venue confirms its real gates first, and route checks from each gate to the blocks it serves follow. These characteristics show why Akii Bua is not a homogeneous population of 20,096 interchangeable seats.

8.2 Known exceptions

The v2 workbook and import guide already name twelve questions. Each becomes a direct task rather than something left for the sample to stumble on: known ambiguity is not randomised away.

Table 9. The Akii Bua exception register

EXCEPTION	WHAT THE SOURCE SHOWS	FIELD TASK	CLOSED BY
Capacity	19,996 sellable seats drawn (20,096 records less 100 restricted-view) against 20,009 in the drawing's table: GRAND 101 −3, GRAND 105 −11, GRAND VIP 101 +1	The row census gives a measured physical count by row and block	Venue confirmation against the measured count
S233 row X	Chairs 11 to 30 drawn; numbers 1 to 10 written with no chairs	Blind full read: first and last labels, count, any trace of positions 1 to 10	Field result, then the venue's choice of 20 or 30 seats
Corner rows W202 T and W207 S	18 corner chairs continue a row past the aisle but are lettered as the rows behind; v2 reletters them (W202 T 1 to 37, W207 S 1 to 35)	Full reads recording each chair's physical row letter and number	Field result and venue agreement on lettering
E124 rows G and H	Seats 22 to 26 run past E124's aisle; a drawing note puts them in E125	Boundary test: full reads of E124 G and H, row tests on E125 G and H, signage at the aisle	Field result and venue decision

EXCEPTION	WHAT THE SOURCE SHOWS	FIELD TASK	CLOSED BY
Pending drawing revisions	The drawing says the N213, N215 and N217 back rows and the W104 to W106 front rows need revising	Full reads of the 9 rows affected	Revised drawing and field result
Back-row reruns	Numbering restarts at 1 inside the same block on 23 back rows, named X2, X3 or U2 in v2	Row tests recording the physical row label and the restart	Field result
Proposed numbering	12 skyboxes (rows A to C of 4 seats) and the VVIP box (rows of 15 and 14): 173 seats in 38 rows	Full reads of all 38 rows; block tests of the 13 areas	Venue sign-off of the numbering, then field result
Possible wheelchair spaces	44 GRAND 105 chairs in W102 F, W108 F, E122 E, E127 E, N115 X and S134 X stand in pairs with 18 gaps of about 2.3 m	Accessibility test: spaces, companion seats, step-free route; full row reads	Venue confirmation and field result
Restricted view	100 seats in 23 runs across 11 blocks	View test per run: photographs from the run's end seats, where the obstruction stops	Field result; sell as restricted or withhold
Press seating	192 seats in separate detail drawings, outside the inventory	The coverage walk records the press areas	Venue decision: accreditation or new sections
Gates	Positions are placeholders; the drawing shows no gates	Route checks once the venue names its gates	Venue confirmation
Numbering direction	v2 assumes seat 1 at the right-hand end for a seated spectator, numbers running clockwise round the bowl	Recorded in every block test	Field result

The capacity difference is not hidden by the method. It stays a registered exception until a measured count and the venue's answer close it.

8.3 An illustrative inspection plan

An illustrative Standard-level plan for Akii Bua comes to 1,592 base field tasks before any escalation. It is computed from the v2 seat list with a 1% seat-level target at 95% confidence; a final plan follows approval of the verification policy.

Table 10. Illustrative Standard-level plan for Akii Bua

LAYER	WHAT IT COVERS	FIELD TASKS
A. Exceptions	58 full row reads (540 seats), 22 flagged row tests and 23 view tests, plus 5 desk tasks (capacity, press, gates, proposed numbering, revised drawing)	103
B. Structure	8 coverage walks (one per stand and tier), 93 block tests, 829 row tests	930
C. Seat sample	Blind seat checks: floor 1.479%, 2 to 14 per block, weight 2 in 43 elevated blocks, frame of 19,556 seats	487
D. Re-inspection	5% of the 1,431 structural tests and seat checks, by a second inspector	72

LAYER	WHAT IT COVERS	FIELD TASKS
E. Adaptive	As findings trigger it (Section 7.2)	As needed
Total		1,592

In the field this means reading every label of 540 seats, reading both end labels of 851 more rows, checking 487 sampled seats, and counting the chairs in every one of the 909 rows. The counts add up to a measured physical capacity for the open question in Section 8.2. A full survey would read every label in the same 909 rows instead of the two end labels and the count; which costs less depends on task times the field trial must measure.

If the plan passes, Loft can state that every block and every row was verified structurally, every exception was closed, and, provided the seat sample finds no critical or major defect, that at most 200 of the frame's 19,556 seats (about 1%) carry one, at 95% confidence.

A flat instruction such as "inspect 5% of Akii Bua" (1,005 random seats) would find a 16-seat numbering shift only 56% of the time, and could never find a chair missing from the list. The SAVS instruction is different: execute the base plan, then extend inspection wherever evidence of risk or discrepancy appears.

8.4 Example field tasks

A field task names a place and asks what is there; it never states the expected answer. Two examples from the Akii Bua plan:

BLIND ROW TEST — BLIND

VENUE	STAND	BLOCK	ROW
Akii Bua Stadium	West	W101	H

SELECTED BECAUSE
Structural census, Layer B

1. Record the block identifier exactly as the signage shows it.
2. Find row H by its physical row label.
3. Facing the pitch, record the seat number on the chair at the left-hand end of the row.
4. Record the seat number on the chair at the right-hand end.
5. Count the chairs in the row.
6. Report any missing, duplicated or unreadable labels you notice.
7. Photograph both end chairs with their labels visible, and the row from the aisle.

Expected values stay hidden until you submit — **Submit observation**

BLIND SEAT CHECK — BLIND

VENUE	STAND	BLOCK	ROW
Akii Bua Stadium	East	E126	C

SELECTED BECAUSE
Seat sample, Layer C

Facing the pitch, count chairs from the left-hand end of row C. At the seventh chair, record its seat label, whether it is present and usable, and whether the view of the pitch is clear. Photograph the label.

Expected values stay hidden until you submit — **Submit observation**

The expected values stay hidden until submission. The application then compares the observation with the canonical inventory and, if they differ, classifies the finding and raises the escalation tasks it requires.

§ 9 Evidence, states and certification

Every observation is kept as evidence, every plan can be regenerated, and every result is reported at exactly the level its evidence supports.

9.1 Verification evidence

Every observation keeps enough evidence for someone who was not there to review it. A verification record holds the inventory and sampling-plan versions; the unit inspected and why it was selected; the inspector and, for a re-inspection, the second inspector; the timestamp; the observation exactly as submitted, locked on submission; photographs where the task requires them; location evidence (GPS is unreliable inside a stadium bowl, so a photograph of the block and row signage is the default); the discrepancy classification; the escalation tasks it generated; and the resolution, with who made it.

Records are append-only. A correction adds a revision with its reason and nothing is overwritten, which gives a chain of evidence from the source inventory, through physical inspection, to acceptance. No customer, payment or contact data belongs in a verification record.

9.2 Reproducibility

Once a plan is issued, its random selection is fixed and anyone with the inputs can regenerate it exactly. The plan stores the seed; the algorithm version and the named pseudo-random generator, with published test vectors so that other implementations draw the same sample; the inventory version, identified by a hash of the canonical seat list; the assurance target, the floor, the tiers, weights and minimum per block; the exception register as it stood at issue; and the selection rules and generation timestamp.

Given the same inventory, algorithm, parameters and seed, the initial sample comes out identical. The seed is fixed before the plan is issued and is not chosen by the team being verified, so that nobody can redraw until a convenient sample appears; committing to it in advance, for example by publishing its hash, makes that checkable. Adaptive additions are recorded separately, each with the finding that triggered it, because they depend on what inspectors observed.

9.3 Verification states

An inventory version moves through seven states rather than a binary verified or unverified flag, and the two end states are never equivalent.

Sample Verified and Fully Enumerated are different end states

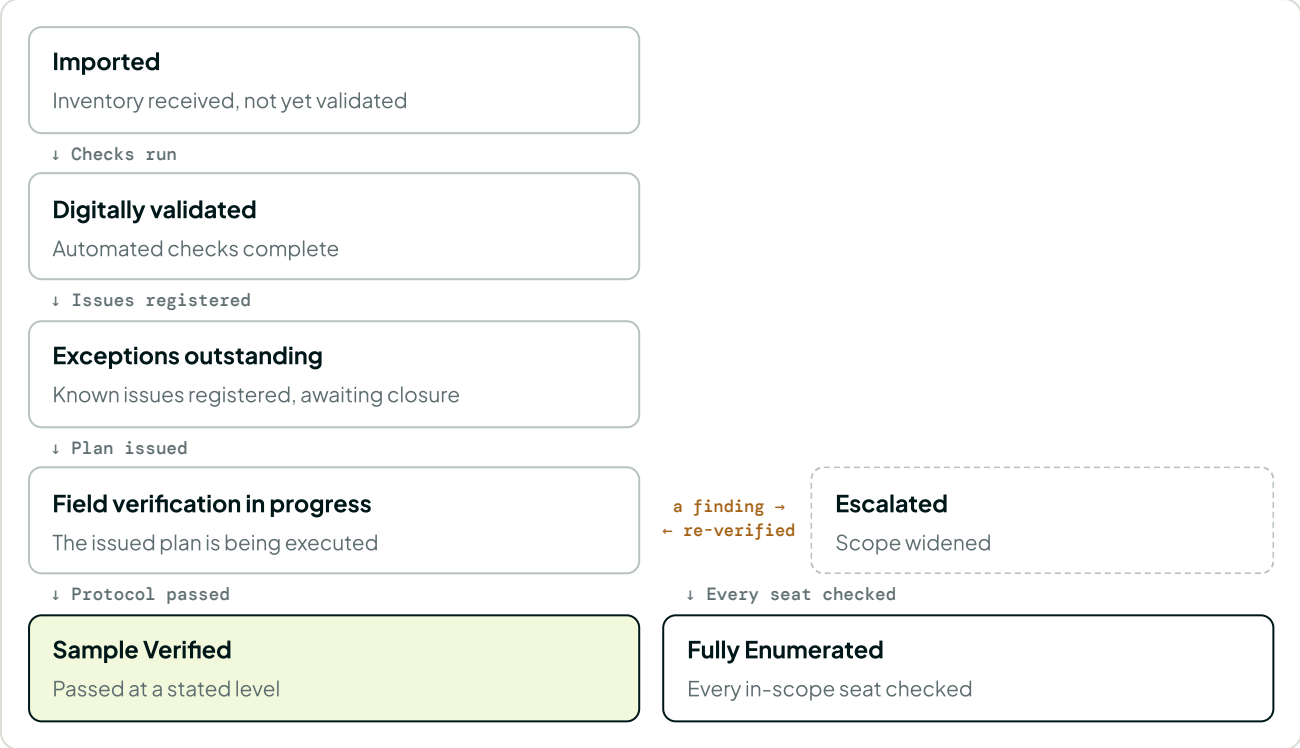


Figure 4. Verification states of one inventory version. Field verification loops through Escalated as often as findings require and ends in one of two end states that are never interchangeable.

Sample Verified rests on structural tests, closed exceptions and a statistical bound; Fully Enumerated means that every in-scope seat was checked. Reports, the TKT interface and any marketing must keep them apart. Corrections made during verification produce a new inventory version, and the certificate is issued for that corrected version: unchanged units keep their evidence and the statistical bound, and changed units are re-verified (Section 10.2). Two administrative states sit outside this flow: Rejected, for a version that will not be accepted, and Superseded, for one replaced by a newer version.

9.4 What “Sample Verified” means

Sample Verified means that the inventory satisfied a documented, reproducible sampling and acceptance protocol at a stated assurance level. It does not mean that every physical seat was inspected; no sampling method can show that without full enumeration, and every report must say so. For the illustrative Akii Bua plan, a passing result would read:

STATEMENT

Inventory version [v] of Akii Bua Stadium is Sample Verified at the Standard level under SAVS 1.0. All 93 blocks and all 909 rows passed blind structural tests, and all registered exceptions are closed ([n] by field result, [n] by recorded decision). Every label of the 540 seats in fully read rows was read. A probability sample of 487 seats, drawn with every seat’s inclusion probability at least 1.479%, found no critical or major defect, so at 95% confidence at most 200 of the 19,556 seats in the sampling frame (about 1%) carry such a defect.

The bound holds wherever those defects might sit, because it rests on the floor probability rather than on the average (Section 6.3). It rests on four assumptions, which the report repeats:

1. seats were drawn by simple random sampling within each block, independently, at no less than the floor;
2. the sample found no critical or major defect (otherwise the weaker bound in Appendix A is reported);
3. the floor was adjusted for the measured lower bound on inspector sensitivity;
4. the version certified is the version sampled, or inherits its evidence under Section 10.2.

Every report also lists what the result does not cover: individual uninspected seats, defects that cannot be seen, changes after the inspection date, and anything beyond the correspondence between inventory and venue (Section 12).

9.5 The certification report

A completed verification produces one report that states the target, the evidence and the limits together. It is generated from the evidence records, never typed by hand, so that every figure can be traced to the observations behind it.

Table 11. Certification report fields, with Akii Bua values where they are known in advance

FIELD	AKII BUA (ILLUSTRATIVE)
Venue	Akii Bua Stadium
Inventory version	[version and hash of the canonical seat list]
Digital population	20,096 seat records; 93 blocks; 909 rows
Method and level	SAVS 1.0; Standard
Assurance target	Critical or major seat-level defects: 1% at 95% confidence; floor 1.479%
Plan	Seed commitment [hash]; issued [date]
Exception tasks	[n] of 103 completed
Structural tests	[n] of 930 completed
Seat checks	[n] of 487 completed
Re-inspections	[n] of 72; agreement [x]%
Adaptive tasks	[n]
Findings	Critical [n]; major [n]; minor [n]; condition [n]
Resolved and outstanding	[n] resolved; [n] outstanding
Blocks passing the protocol	[n] of 93
Rows passing structural tests	[n] of 909
Registered exceptions closed	[n] by field result and [n] by decision, of 12
Statistical statement	As in Section 9.4, with the observed figures
Verification state	[state]

FIELD	AKII BUA (ILLUSTRATIVE)
Limitations	The standard text from Sections 9.4 and 12

§ 10 Implementation in TKT

SAVS fits TKT as a verification stage between importing a venue and publishing it for sale, so that only an accepted revision can become the canonical source of event seat inventory.

10.1 Workflow

1. **Import.** TKT's venue import, or Loft's drawing conversion toolkit, produces the venue structure (sections, rows, seats and gates) as a draft revision.
2. **Validate.** The checks in Section 6.1 run on TKT's canonical seating model, and each finding opens an entry in the exception register.
3. **Plan.** The sampling engine issues a plan at the chosen assurance level, with its seed committed (Sections 6 and 9.2).
4. **Verify.** A field application gives inspectors blind tasks, works offline inside the bowl, captures observations, photographs and signage evidence, and synchronises when connected.
5. **Escalate.** Each finding is classified and its escalation tasks are generated automatically (Section 7.2).
6. **Reconcile.** Corrections and decisions are applied as a new revision. Seats with open questions are held or blocked with TKT's existing seat actions until they are closed.
7. **Accept and publish.** A reviewer accepts the revision as Sample Verified or Fully Enumerated, and only then can it be published; events adopt it as they adopt any revision today.

What has to be built: the exception register and a verification state on each venue revision; a deterministic, versioned plan generator; the field application, with blind task flow, photo capture and re-inspection assignment; the comparison and classification engine; the report generator (Section 9.5); a publish gate that blocks or warns on an unverified revision, set by policy; and delta verification over TKT's revision differences.

Roles. A verification lead approves the policy, issues the plan and commits the seed. Inspectors carry out tasks. A venue representative answers exceptions. A reviewer, independent of the import team, signs acceptance.

10.2 Version control and delta verification

Verification belongs to one inventory version. When the inventory changes, only the evidence for unchanged units carries over:

$$V_1 \rightarrow \text{verification} \rightarrow \text{accepted}, \qquad V_2 \neq V_1 \Rightarrow \text{evidence for } V_1 \text{ is not inherited wholesale}$$

TKT computes the difference between versions at seat, row and block level and applies these provisional rules:

- **Unchanged units keep their evidence** when their records (identities, labels, order and the attributes a test examined) are identical and no physical change has been reported.
- **A changed seat reopens its row:** the row test and any seat checks in that row are repeated.

- **A changed block label, row list or numbering direction reopens the block:** its block test and all its row tests.
- **Added or removed seats change the sampling frame,** so the seat sample is topped up until every seat in a changed block meets the floor again.
- **Metadata-only changes** (category, ticket type, colour) need a desk check, unless the attribute is physically marked in the venue.

Physical works reopen the areas they touch whatever the digital version shows, and evidence ages: a policy may require periodic re-verification at the Reduced level. The certification statement names the date of each piece of evidence it relies on. Delta verification means that a change to one section does not force a new inspection of the whole stadium: the affected section and anything that depends on it are re-verified, and the rest keeps its evidence.

§ 11 Validation programme

SAVS should not be presented as better than simpler methods until simulation and a field trial show it. Both are designed here so that the claim can be tested rather than asserted.

11.1 Simulation

- **Venues:** synthetic venues of 5,000, 20,000, 40,000 and 80,000 seats with realistic rows, blocks and aisles, plus the real Akii Bua structure.
- **Injected defects:** independent random errors; single missing seats; contiguous missing runs; numbering shifts; wrong row boundaries; wrong row and section assignments; duplicated records; unrecorded physical seats; and mixtures of these. A clustering parameter controls how strongly defects bunch.
- **Inspectors:** a sensitivity below 1 and a false-alarm rate, both varied.
- **Methods compared:** simple random sampling; fixed-percentage sampling; stratified random sampling; systematic row sampling; SAVS without adaptation; full SAVS; and complete enumeration as the benchmark.
- **Equal effort:** methods are compared at equal field time, using task times measured in the field trial, not at equal task counts, because a row test costs more than a seat check.
- **Precision:** about 10,000 runs per scenario and method, which keeps the 95% interval on each detection probability within about one percentage point.
- **Sensitivity analysis:** escalation thresholds, risk weights, the floor and the re-inspection rate.

11.2 Field trial

Run the Standard plan at a venue and, as ground truth, fully enumerate a random subset of its blocks, so that what the plan missed can be measured. The trial also times each kind of task, measures inspector agreement, and tests whether blind tasks stay blind in practice. Its outputs (calibrated thresholds and weights, an effort model and an estimate of inspector sensitivity) replace the provisional values in this paper.

11.3 Evaluation metrics

Each method is judged on assurance first and effort second: a method that inspects fewer seats but misses whole defective sections is not better because it is cheaper.

The measures	
P_D	P(a material defect is detected)
P_{FC}	P(accepted | a material defect is present): the false certification probability
$E[T]$	Expected field time, including escalation
R_D	Defects detected ÷ defects present
R_U	Affected structural units detected ÷ affected structural units
AOQ	Defect rate remaining after correction
C_D	Inspection effort ÷ material defects detected

- The false certification probability P_{FC} is the primary measure: it is the risk the method exists to control.
- Structural unit recall R_U shows whether a method finds every affected row and block, not just many defective seats.
- The average outgoing quality AOQ measures what the venue is left with after correction, which is what spectators experience.
- $E[T]$ and C_D measure cost, and are compared only between methods with acceptable P_{FC}. The false alarm rate is tracked alongside, because false alarms waste escalation effort.

11.4 Hypotheses

The programme tests three propositions, chosen so that the results can count against SAVS as well as for it.

1. **H1.** For venue inventories with spatially or structurally correlated errors, SAVS has a lower probability of false certification than simple random sampling at equal field time.
2. **H2.** That advantage grows as defect clustering increases.
3. **H3.** When errors are independent and confined to labels inside rows, the advantage disappears: stratified seat sampling of equal field time performs at least as well.

H1 is supported only if SAVS's false certification probability is lower in every clustered scenario, with non-overlapping 95% Monte Carlo intervals at equal expected field time. One result needs no test because it holds by construction: a sample drawn from the digital list can never find a seat missing from it, so for that defect any structural census beats any seat sample. Until simulation and field validation are complete, these are propositions to test, not conclusions to cite.

§ 12 Limitations

SAVS reduces the risk of accepting a wrong inventory; it does not remove it, and its reports must say where it stops.

1. Sampling cannot prove that any individual uninspected seat is correct.
2. Results depend on inspectors' competence and independence. Blind tasks and re-inspection reduce that dependence; they do not remove it.
3. Some defects cannot be seen, such as a wrong digital category or wrong map coordinates.

4. The method needs a structured inventory, with rows, blocks and their order. A flat list must be structured before it can be verified.
5. Risk weights, escalation thresholds and assurance levels need empirical calibration.
6. The statistical bound holds only if the plan is executed as issued, against the inventory version it was drawn from.
7. The structural census assumes that every row can be reached and its labels read. Poor labelling raises the cost and may force full reads.
8. Venues change after verification, and evidence ages.

Finally, SAVS verifies the correspondence between a defined digital inventory and a physical configuration at a point in time. It does not certify structural engineering, fire safety, crowd capacity or regulatory compliance.

§ 13 Conclusion

Digital seat inventories are operating infrastructure: they drive ticket sales, spectator routing, capacity management and access control. Verifying them as a spreadsheet reconciliation is not enough, and inspecting every seat for every revision spends effort on areas that are already right.

SAVS takes the middle path by being explicit about what is known, what is uncertain, what is sampled, what must be inspected, and what needs more inspection once a defect is found. Its central move is simple: census the structure, sample the seats, inspect known questions directly, and escalate by rule.

Akii Bua shows why this matters. Its 20,096 seat records in 93 blocks and 909 rows come with a 13–seat capacity question, a row whose first ten numbers have no chairs, joined corner rows, proposed numbering for 173 seats and 44 possible accessibility positions. A random percentage would leave most of these to chance; SAVS turns each into a task.

The next stage is empirical. Simulation, sensitivity analysis and a field trial will set the thresholds and show whether SAVS meets its objective: lowering the probability that a materially wrong seat inventory is accepted for use, without inspecting areas that are already right.

Appendix A Formulas

The figures in this paper come from the formulas below, computed exactly for Akii Bua's 20,096 seat records.

Detection by a simple random sample of n from N seats, D of them defective:

$$P(\text{detect} \geq 1) = 1 - \frac{\binom{N-D}{n}}{\binom{N}{n}} = 1 - \prod_{i=0}^{D-1} \frac{N-n-i}{N-i}$$

Sample size for a target (p^*, β): the smallest n with detection of at least $1-\beta$ when $D^* = p^* N$, rounded. For large N:

$$n \approx \frac{\ln \beta}{\ln(1 - p^*)} \approx \frac{3}{p^*} \qquad (\beta = 0.05)$$

Upper bound after zero findings: the largest D whose miss probability is above β. In binomial form:

$$p_u = 1 - \beta^{1/n} \approx \frac{3}{n}$$

Bound after one or more findings ($x \geq 1$): the largest D for which a binomial count with D trials and success probability π_0 would show x or fewer findings with probability above β. Sampling every seat at exactly the floor makes this exact; for stratified sampling without replacement it is a close approximation, and it is always weaker than the zero-finding bound. At Akii Bua's floor, one finding gives at most 318 defective seats in the frame, and two give at most 423.

The floor guarantee. Sampling without replacement inside block h, each factor in the product above is at most $1 - n_h/N_h$, and blocks are sampled independently. If every block's sampling fraction is at least π_0, then for any D defective seats, wherever they are:

$$P(\text{miss all}) = \prod_h \frac{\binom{N_h - D_h}{n_h}}{\binom{N_h}{n_h}} \leq \prod_h \left(1 - \frac{n_h}{N_h}\right)^{D_h} \leq (1 - \pi_0)^D$$

and with $\pi_0 = 1 - \beta^{1/D^*}$, $(1 - \pi_0)^{D^*} = \beta$.

Rows use the same hypergeometric formula, with R rows of which D_r are defective: at Akii Bua ($R = 909$), finding one defective row with 95% probability takes 864 random row tests, and finding 9 defective rows takes 257.

Inspector sensitivity s lowers detection to about:

$$P(\text{detect}) \approx 1 - (1 - s\,\pi)^D$$

so the floor that keeps the target is $\pi_0 = (1 - \beta^{1/D^*})/s$. At Akii Bua, a measured s of 0.9 would raise the floor from 1.479% to 1.644%.

Estimating the number of defective seats from a probability sample with inclusion probabilities π_i (Horvitz and Thompson, 1952):

$$\hat{D} = \sum_{i \in S_R,\, i \text{ defective}} \frac{1}{\pi_i}$$

Monte Carlo precision: the runs needed for a 95% interval of $\pm h$ on a probability near 0.5:

$$n_{\text{runs}} = \frac{1.96^2 \times 0.25}{h^2} = 9{,}604 \qquad (h = 0.01)$$

Appendix B Glossary

Terms and their meaning in this paper

TERM	MEANING IN THIS PAPER
Assurance level	Reduced, Standard, Tightened or Full enumeration: how much evidence acceptance rests on (Section 7.4)
Assurance target	The defect rate p^* the seat sample must catch, with probability $1 - \beta$
Blind task	A task that names a place and asks what is there, without stating the expected answer
Block	A seating section bounded by aisles; the stratum for seat sampling
Block test	A blind check of a block's label, front and back row letters, row count and numbering direction
Census statement	A claim about every unit of a kind, supported by testing all of them
Condition finding	A physical problem with a correctly represented seat, such as a broken chair
Coverage walk	A walk of a stand and tier that matches every seating area seen against the block list
Critical, major, minor	The three defect classes and their consequences (Section 7.1)
Delta verification	Re-verifying only what changed between inventory versions (Section 10.2)
Escalation (E1 to E4)	Neighbourhood, full row, tightened block sample, full block inspection (Section 7.2)
Exception register	The list of known questions found by validation or raised by the venue, each closed by a task or a decision
Floor (π_0)	The minimum inclusion probability of every seat in the sampling frame
Fully Enumerated	The state in which every in-scope seat was individually checked
Inclusion probability	The chance that a unit is selected under the issued plan
Layers A to E	Exceptions, structure, seat sample, re-inspection, adaptive escalation
Material discrepancy	One that can cause a duplicate or impossible sale, a wrong seat assignment, a wrong route or a wrong capacity figure
Row test (R1)	A blind record of a row's first label, last label and chair count
Row read (R2)	A record of every label in a row, in order
Sample Verified	The state in which the inventory passed the protocol at a stated assurance level; not every seat was inspected
Sampling frame	The seats eligible for the seat sample: all seats not already read in full

TERM	MEANING IN THIS PAPER
Seat check	A blind check of the label, condition, view and marking at one stated position

References

STANDARDS

International Organization for Standardization (2026). ISO 2859-1:2026. Sampling procedures for inspection by attributes, Part 1: Sampling schemes indexed by acceptance quality limit (AQL) for lot-by-lot inspection. Edition 3, replacing the 1999 edition and its amendments. Geneva: ISO. iso.org ↗

International Organization for Standardization (2020). ISO 2859-2:2020. Sampling procedures for inspection by attributes, Part 2: Sampling plans indexed by limiting quality (LQ) for isolated lot inspection. Edition 2. Geneva: ISO. iso.org ↗

STATISTICS

Clopper, C. J. and Pearson, E. S. (1934). The use of confidence or fiducial limits illustrated in the case of the binomial. Biometrika, 26(4), 404–413.

Cochran, W. G. (1977). Sampling Techniques, 3rd edition. New York: Wiley.

Dodge, H. F. and Romig, H. G. (1959). Sampling Inspection Tables: Single and Double Sampling, 2nd edition. New York: Wiley.

Hanley, J. A. and Lippman-Hand, A. (1983). If nothing goes wrong, is everything all right? Interpreting zero numerators. JAMA, 249(13), 1743–1745.

Horvitz, D. G. and Thompson, D. J. (1952). A generalization of sampling without replacement from a finite universe. Journal of the American Statistical Association, 47(260), 663–685.

Neyman, J. (1934). On the two different aspects of the representative method. Journal of the Royal Statistical Society, 97(4), 558–625.

Squeglia, N. L. (2008). Zero Acceptance Number Sampling Plans, 5th edition. Milwaukee: ASQ Quality Press.

Thompson, S. K. (1990). Adaptive cluster sampling. Journal of the American Statistical Association, 85(412), 1050–1059.

Thompson, S. K. and Seber, G. A. F. (1996). Adaptive Sampling. New York: Wiley.

Wald, A. (1945). Sequential tests of statistical hypotheses. Annals of Mathematical Statistics, 16(2), 117–186.

CASE STUDY SOURCES

The Independent (2024). Plans finalized for 20,000-seater Akii-Bua Olympic Stadium in Lira. 10 December 2024. independent.co.ug ↗

Uganda Radio Network (2026). Akii-Bua Olympic Stadium nears completion as installation of seats kicks off. 15 May 2026. ugandaradionetwork.net ↗

Loft Technologies (2026). Akii Bua seat plan v2: seat list, workbook and import guide. 29 October 2026. Internal working files; all Akii Bua figures in this paper were recomputed from its seat list.

LOFT TECHNOLOGIES LIMITED
Plot 67A, Spring Road, Level 2, The Guardian, Bugolobi, Kampala, Uganda
info@loftuganda.tech · +256 784 094 309 · loftuganda.tech

Structure-aware adaptive sampling for venue seat inventory verification. White paper, version 1.0, October 2026. The latest version is at loftuganda.tech/research/savs.