loft

WHITE PAPER

STADIUM ACCESS INTELLIGENCE

What we have learned about stadium access control

A field-led paper on delayed admissions, late-arrival pressure, gate throughput, boids-based crowd modelling and practical access-control decisions for African stadiums.



PREPARED BY

Mugalula Nickson, Tendo Michael and Samuel Mabonga

PUBLISHED

June 2026

READ ONLINE

loft.ug/research/access-control

ABOUT THIS PAPER

PREPARED BY	Mugalula Nickson, Tendo Michael and Samuel Mabonga
ORGANISATION	Loft Technologies, June 2026
SERIES	White paper · Stadium access intelligence
PUBLISHED	June 2026

CONTENTS

Executive note 3
Reader value 3
Access pressure 3
Formula framework 5
Worked example 5
Throughput sensitivity 6
Queue compounding 7
Lane planning 8
Time to clear 9
Boids translation 9
Boids formula 10
Spatial risk 11
Delayed admissions 12
Turnstile behaviour 12
Live control model 13
Intervention playbook 14
African stadium realities 14
Implementation roadmap 15
Data requirements 15
KPIs 16
Limitations 16
Closing position 17
References 18

HOW TO CITE

Mugalula, N. and Tendo, M. (2026). What We Have Learned About Stadium Access Control. White paper. Loft Technologies.

Read online and download: loft.ug/research/access-control

60%	18,000	600/min	56 → 112
Of attendees assumed to arrive in the final 30 minutes	People to admit in the final 30 minutes of a 30,000-person match	Peak arrivals across the stadium under that assumption	Lanes needed as the admission cycle slows from 4.5 to 9 seconds

EXECUTIVE NOTE

This paper is about learning from real stadium access pressure

This white paper is written from the practical side of stadium access control. It is not presented as a laboratory paper on simulation. It is a lessons paper: what we have learned from planning and operating digital ticketing, scanning, turnstiles, handheld validation and live access control for high-pressure stadium events, and what other venues can take from that experience.

Our strongest lesson is that access control is not only a question of whether a ticket is valid. It is a question of whether a valid ticket can be turned into a physical admission quickly enough, at the right gate, under real arrival pressure, with enough stewarding and hardware capacity to keep queues moving.

Boids-based crowd modelling is useful because it translates individual movement into system behaviour. A fan chooses a route, slows down near a queue, follows a group, reacts to signage, avoids dense areas when possible, and is influenced by nearby people. These simple behaviours can produce pressure pockets, blocked approaches and cross-flows that are not obvious from ticket sales data alone.

WORKING ASSUMPTION USED IN THIS PAPER

For a high-demand football event, we assume **60 percent of attendees arrive in the final 30 minutes before kickoff**. This is a deliberately stressful assumption. It forces the access-control plan to answer the real question: can the stadium process the late surge without unsafe density?

60%	scan-to-admit	boids + logs
ARRIVAL SHOCK IN FINAL 30 MINUTES	THE REAL ACCESS METRIC	MOVEMENT MODEL GROUNDED IN OPERATIONS DATA

READER VALUE

What the reader can pick from this paper

The reader does not need to adopt every concept in this paper. A stadium operator, federation, ticketing provider, security commander or venue consultant can pick one practical improvement and apply it immediately.

For stadium operators

Stop planning with total capacity alone. Break every event into gate-by-gate demand, gate-by-gate throughput and gate-by-gate risk. A stadium with enough total lanes can still fail at one overloaded gate.

For federations and event owners

Ticket categories, complimentary allocations, accreditation flows and sales closure rules are crowd-control decisions. They must be approved early enough to influence gate design and communication.

For security and stewarding teams

Failed admissions are not only ticketing problems. Every failed scan consumes lane time, creates argument time and pushes pressure backwards into the queue.

For technology teams

The access-control dashboard should show admission pressure, not only scan totals. The useful question is what is forming now and what will happen in the next 10 minutes.

The best stadium access-control system is not the one that only counts entries. It is the one that helps the operations team understand pressure early enough to act.

ACCESS PRESSURE

The central lesson: pressure forms when demand outruns admission capacity

Crowd risk near gates is rarely caused by attendance alone. It is caused when too many people arrive at the same place faster than the venue can validate, admit and disperse them. The failure point is the gap between arrival demand and effective admission capacity.

Late-arrival stress assumption: 60% in final 30 minutes

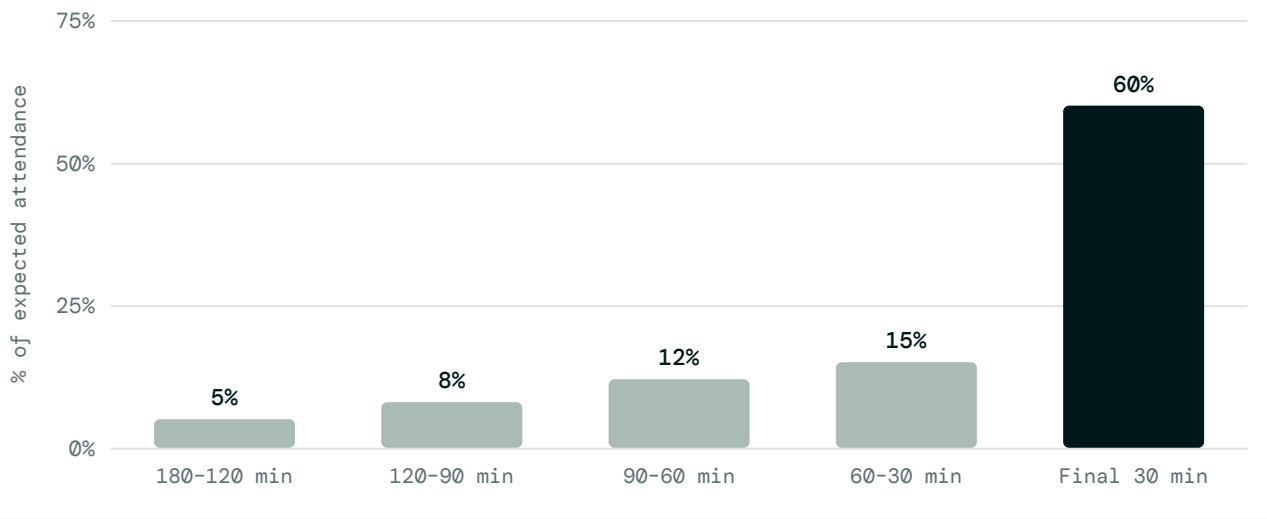


Figure 1. The late-arrival assumption concentrates 60 percent of demand into the final 30 minutes. This is the stress condition that reveals whether the access plan is realistic.

For a 30,000-attendee match, this assumption places 18,000 people into the last 30 minutes. That is 600 arrivals per minute across the whole stadium footprint. If those arrivals are unevenly distributed, one gate can become dangerous even if total stadium capacity looks adequate.

Planning views and the questions they miss

PLANNING VIEW	WHAT IT MISSES	BETTER QUESTION
Total tickets sold	When people arrive and which gate they use	How many people arrive per gate per minute?
Total turnstiles installed	Open lanes, device uptime and human passage time	How many actual admissions can each gate process now?
Scan count	Failed attempts, wrong-gate pressure and physical passage delays	How many scans become admitted people?

FORMULA FRAMEWORK

The practical maths of stadium ingress

The formulas below are not meant to make crowd control look complicated. They are meant to force the right operational questions. Every variable should be measurable from ticketing logs, device logs, steward reports or field observation.

$$\lambda_{g,t} = \frac{A_{g,t}}{\Delta t}$$

Arrival rate at gate g in time window t. This converts attendance expectation into operational pressure.

$$\tau_g = t_p + t_v + t_o + t_m + t_e$$

Total admission cycle time. This is the full time from ticket presentation to the lane being ready for the next person.

$$\mu_{g,t} = L_g \times \frac{60}{\tau_g} \times S_g \times U_g \times C_g$$

Effective gate throughput. It combines active lanes, cycle time, success rate, uptime and crowd-readiness conditions.

$$Q_{g,t+1} = \max(0, Q_{g,t} + \lambda_{g,t} - \mu_{g,t})$$

Queue growth. If arrivals exceed throughput, the queue does not merely remain. It compounds over time.

INTERPRETATION

A stadium can have many turnstiles and still experience dangerous pressure if the admission cycle is slow, success rate is weak, device uptime is low, or the crowd arrives too late.

WORKED EXAMPLE

A 30,000-person match under late-arrival pressure

Assume a high-demand football match with 30,000 expected attendees and a late-arrival pattern where 60 percent arrive in the final 30 minutes. The stadium therefore needs to absorb 18,000 people in 30 minutes.

Inputs for the worked example

INPUT	ASSUMPTION	OPERATIONAL MEANING
Expected attendance	30,000	Total people to admit
Late-arrival share	60%	18,000 people in final 30 minutes
Peak arrival rate	600/min	Demand across the footprint
Active lanes	63	Turnstiles and equivalent admission lanes
Cycle time	6.4 seconds	Presentation, validation, unlock, pass and reset
Success rate	92%	Share of attempts that become clean admissions
Uptime	95%	Devices and lanes available
Crowd-readiness factor	92%	Tickets ready, phone brightness, clear movement, low confusion

ILLUSTRATIVE RESULT

63 lanes at 6.4 seconds can theoretically handle about 590 people per minute before losses. After success-rate, uptime and readiness factors, effective throughput falls to roughly 500 people per minute. Against a 600 people/minute arrival shock, the stadium accumulates a deficit of about 100 people per minute unless operations intervene.

THROUGHPUT SENSITIVITY

Small quality losses become large capacity losses

The admission system is sensitive to scan-to-admit success rate. When a venue is under late-arrival pressure, a few percentage points of failure can be the difference between flow and backlog.

Success-rate sensitivity: small drops become large capacity losses

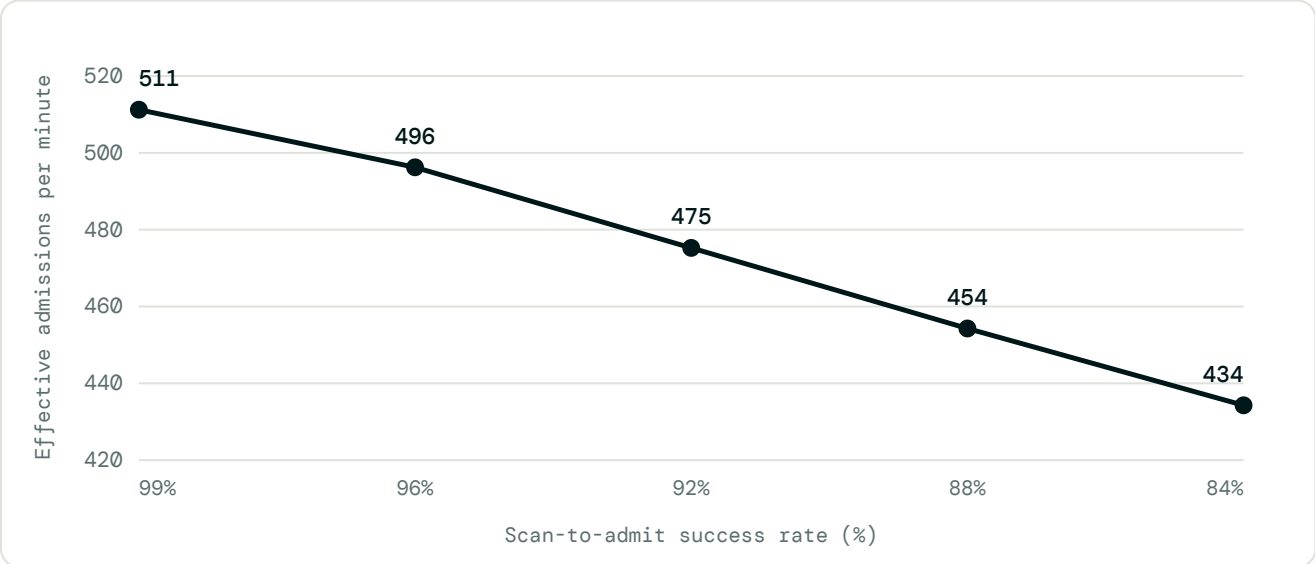


Figure 2. Illustrative model using 63 lanes, a 6.4 second admission cycle, 95 percent uptime and 92 percent crowd-readiness factor.

Success-rate loss is not only caused by fraudulent or invalid tickets. In field operations it can come from wrong-gate presentation, fans searching for tickets at the scanner, unreadable screens, duplicate screenshots, category mismatch, slow network response, steward uncertainty, or physical hesitation at the turnstile.

Operational leaf

Report success rate by gate and by 10-minute interval. A stadium-wide average hides the gate where pressure is forming.

Technology leaf

Treat every failed scan as a queue event. Failed attempts should feed the live pressure model, not only the audit report.

QUEUE COMPOUNDING

Delayed admission turns small deficits into visible crowd pressure

A gate can appear stable for a few minutes and then become difficult quickly. The reason is compounding. When capacity is below demand, the unprocessed people do not disappear. They remain in the approach, making every next minute harder to manage.

Queue growth under a 600 people/minute peak arrival shock

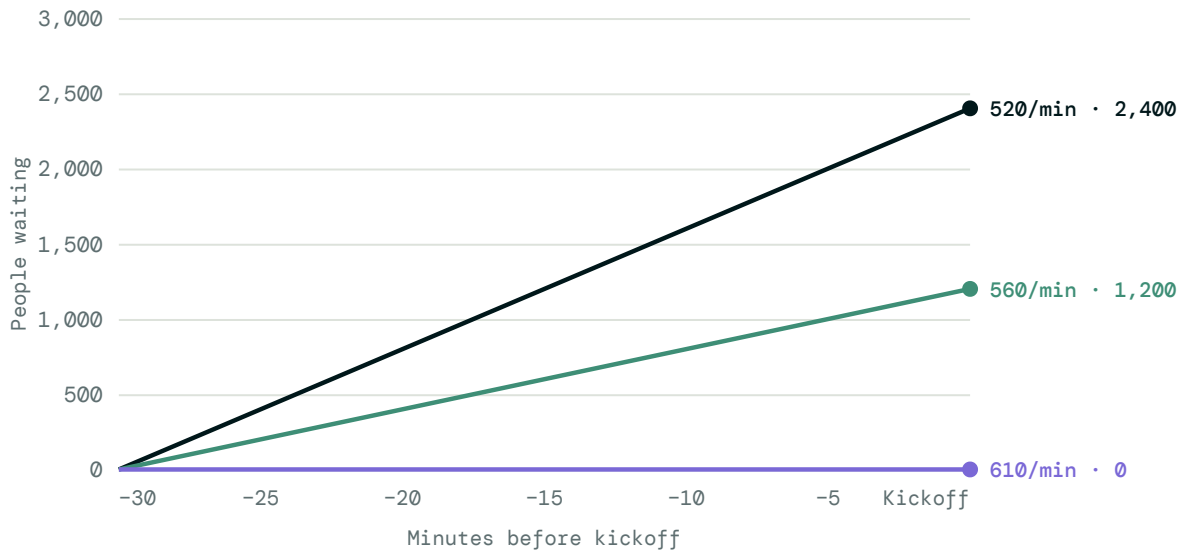


Figure 3. Illustrative queue accumulation when peak arrival demand is 600 people per minute. The difference between 560/min and 610/min is operationally decisive.

$$P_{g,t} = \frac{\lambda_{g,t} + Q_{g,t}/\Delta t}{\mu_{g,t}}$$

Admission pressure ratio. Values above 1.0 indicate that the gate is receiving more pressure than it can process.

PRACTICAL RULE

When pressure ratio stays above 1.0 for more than one reporting interval, the control room should not wait for a visual crisis. It should trigger a prepared intervention.

LANE PLANNING

The number of open lanes must follow the peak, not the average

The lane requirement formula helps convert peak demand into an operational staffing and equipment plan. It also shows why slow physical admission is expensive. Each extra second in cycle time increases the number of lanes required.

$$N_{\text{lanes}} = \left\lceil \frac{\lambda_{\text{peak}} \times \tau_g}{60 \times S_g \times U_g \times C_g} \right\rceil$$

Minimum active lanes required for a target peak arrival rate. This should be calculated per gate, not only for the whole stadium.

Lane requirement rises sharply as the admission cycle slows

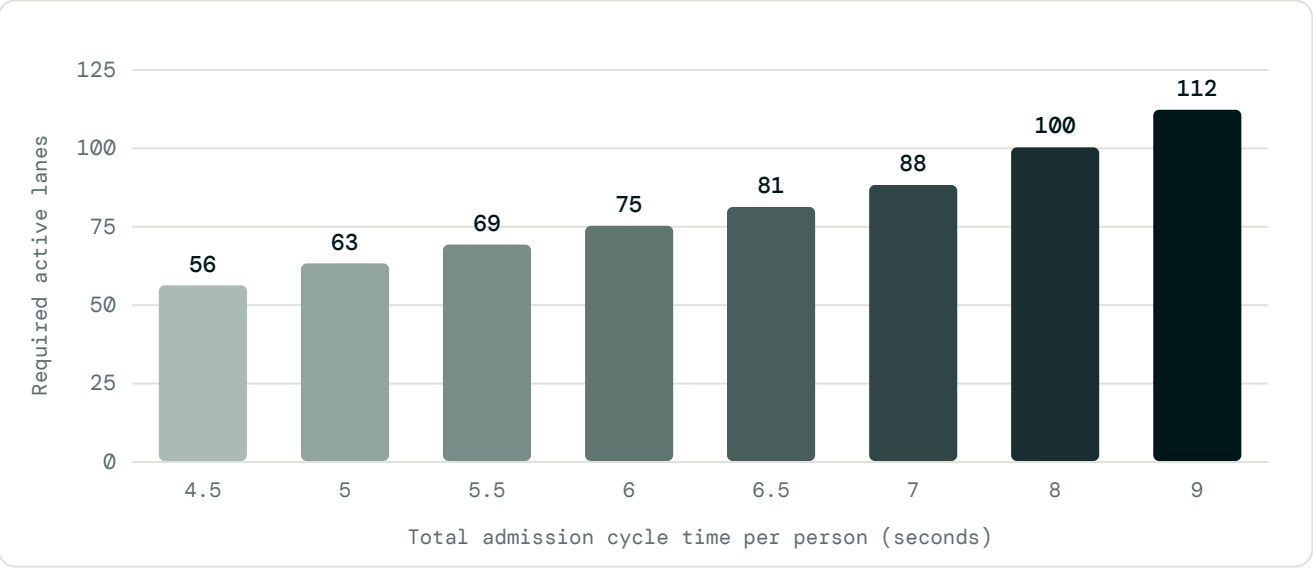


Figure 4. Illustrative lane requirement at 600 people/minute peak arrival demand, 92 percent success, 95 percent uptime and 92 percent crowd-readiness.

This chart gives a simple lesson: improving cycle time can be as valuable as adding hardware. Faster ticket retrieval, better feedback screens, upstream filtering and clearer gate assignment can reduce the number of lanes required to handle the same crowd.

TIME TO CLEAR

Backlog must be translated into minutes, not only people

Queue length matters, but decision makers often respond better to time-to-clear. A queue of 2,000 people means different things depending on whether the gate has a net clearing capacity of 50 people per minute or 200 people per minute.

$$TTC_{g,t} = \frac{Q_{g,t}}{\mu_{g,t} - \lambda_{g,t}} \quad \text{where } \mu_{g,t} > \lambda_{g,t}$$

Time-to-clear estimate. This applies only when throughput is greater than arrival rate. If it is not, the queue is still growing.

Time-to-clear makes backlog visible to decision makers

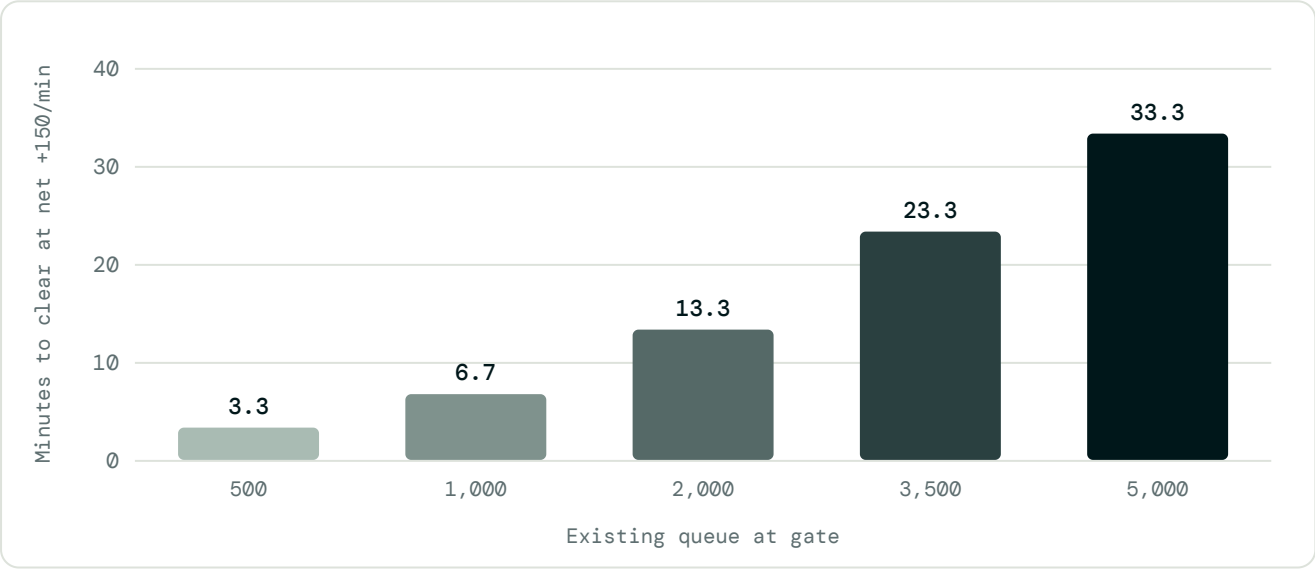


Figure 5. Illustrative time-to-clear using a net clearing capacity of 150 people per minute after arrivals are reduced or additional capacity is opened.

Time-to-clear should be shown on the live dashboard because it converts technical conditions into operational urgency. If the gate will clear after kickoff, the venue must decide whether to reroute, open more lanes, delay entertainment sequencing, strengthen outer filtering or communicate clearly with the crowd.

BOIDS TRANSLATION

How boids becomes useful for stadium access

Boids modelling starts with simple local rules: separation, alignment and cohesion. For stadium operations, these behaviours must be translated into practical crowd movement logic.

Boids layer: individual rules reveal system-level flow

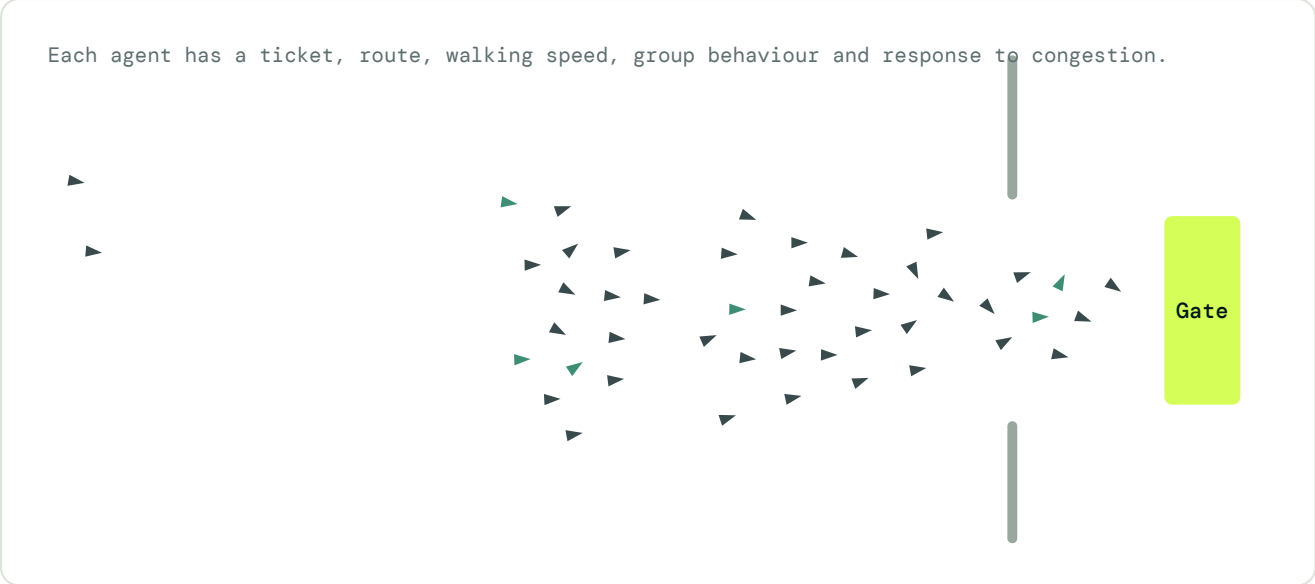


Figure 6. A boids-style model treats fans as agents moving toward gates while responding to nearby people, barriers, signage and congestion.

Boids concepts translated for stadium operations

BOIDS CONCEPT	STADIUM TRANSLATION	OPERATIONAL USE
Separation	People avoid direct collision and uncomfortable density	Estimate where queues widen, bend or push into roads
Alignment	People follow the visible flow of nearby fans	Predict herd movement toward wrong gates or short-looking queues
Cohesion	Groups try to stay together	Model family, fan club and group arrival behaviour
Destination seeking	People move toward assigned gate or visible entrance	Test whether signage and ticket instructions are clear enough
Rerouting	People change path when congestion becomes obvious	Prevent uncontrolled movement from one overloaded gate to another

BOIDS FORMULA

The movement layer should be explainable

A useful simulation must be explainable to operations teams. The velocity formula below is a simplified way of saying that each fan's next movement is influenced by nearby people, group behaviour, gate destination and rerouting signals.

$$v_i^{new} = w_s S_i + w_a A_i + w_c C_i + w_d D_i + w_r R_i$$

Simplified agent movement update. Separation, alignment, cohesion, destination-seeking and rerouting forces combine into the next movement direction.

The weights should not be treated as universal constants. They should be calibrated by venue type and event context. A derby match, family event, concert, international fixture and VIP-heavy event will all produce different movement behaviour.

Before the event

Use the model to test gate assignment, approach routes, barricade placement, signage points and steward positions.

During the event

Use actual scan rates, queue estimates and density observations to update risk assumptions and recommend interventions.

SPATIAL RISK

Risk is spatial, not only numerical

A gate with 2,000 people waiting may be manageable in a wide, well-barricaded forecourt. The same number can be unsafe if it is compressed into a narrow approach road, stair base, wall edge or merge point. That is why the model must show spatial pressure.

$$D_{zone} = \frac{N_{zone}}{A_{zone}}$$

Density in a zone. The same number of people can be safe or unsafe depending on the area and movement conditions.

Spatial pressure map: the danger is where density forms

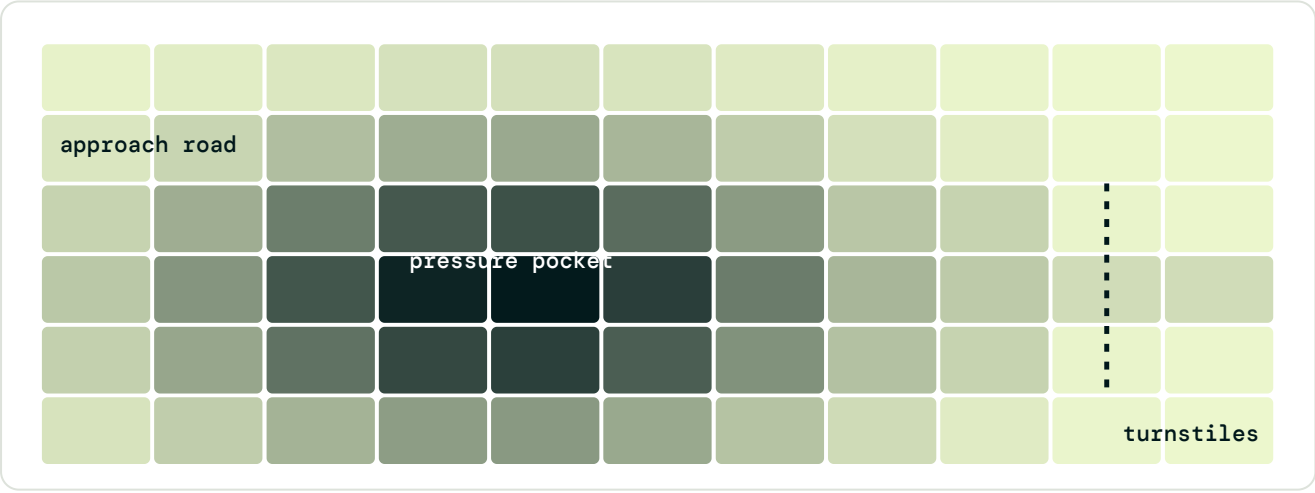


Figure 7. A spatial pressure map helps the control room see where density is forming, not just how many people have scanned.

Boids-based modelling adds value here because it shows how people redistribute when one route becomes dense. If the simulation predicts that an overloaded gate spills into a road or blocks an adjacent gate, the operations team can change the geometry before event day.

DELAYED ADMISSIONS

What actually slows the lane

Delayed admission is a full-chain problem. It can start before the fan reaches the scanner and continue after the ticket has been accepted. The access-control team should therefore measure causes of delay in the whole gate environment, not only at the validation API.

Where delayed admissions can come from

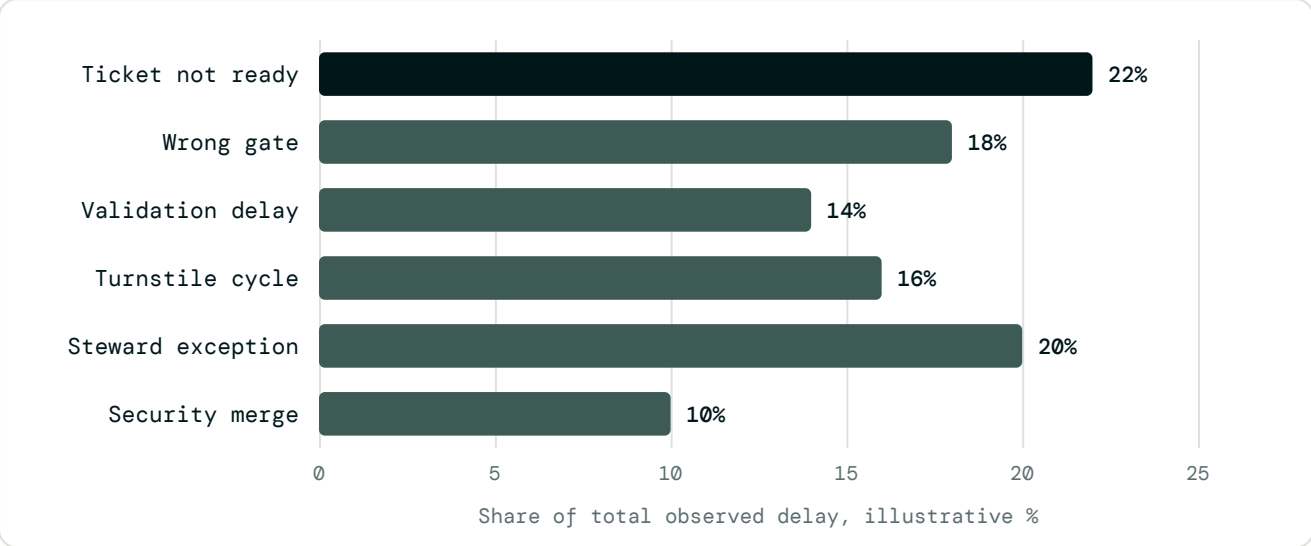


Figure 8. Illustrative breakdown of delayed admission causes. Actual values should come from gate logs, exception tags and steward debriefs.

Delay sources and countermeasures

DELAY SOURCE	WHAT IT LOOKS LIKE	COUNTERMEASURE
Ticket not ready	Fan reaches scanner then searches WhatsApp, email or gallery	Pre-gate reminders, holding screens, ticket-help desks upstream
Wrong gate	Valid ticket presented at wrong category or zone	Gate printed on ticket, WhatsApp reminder, outer steward filtering
Validation delay	Scanner waits for response or retries	Edge validation, offline cache, redundant networks
Turnstile hesitation	Accepted ticket but user delays passage	Feedback screens, better lane design, steward cueing
Exception handling	Dispute, override, accreditation issue or duplicate ticket	Remove exceptions from main lane immediately

TURNSTILE BEHAVIOUR

The turnstile is both a validator and a traffic signal

The moment after a valid ticket is presented is critical. The system must unlock fast enough, stay open long enough for safe passage, reset quickly, and give clear feedback if the ticket is not accepted. The turnstile should not be treated as a passive barrier. It is a traffic signal for the fan and the steward.

Open time

Too short creates repeated attempts, confusion and assisted passage. Too long slows the lane and increases tailgating risk.

Feedback screen

Clear messages such as accepted, wrong gate, already used, invalid event or seek steward reduce argument time and keep the lane moving.

Alarm logic

Alarms should distinguish forced passage, tailgating, wrong direction and repeated failed attempts. Not every issue needs the same steward response.

Accessibility

PWD, VIP, family and assisted-entry routes need different cycle assumptions and should not be judged only against general admission speed.

$$\tau_g = t_p + t_v + t_o + t_m + t_e$$

Cycle time must include the physical action of the turnstile and the human passage, not only server validation time.

LIVE CONTROL MODEL

What the command centre should see

A good access dashboard should answer four questions: where is pressure forming, why is it forming, what happens if nothing changes, and what intervention is most likely to help?

$$R_{g,t} = \alpha P_{g,t} + \beta(1 - S_{g,t}) + \gamma W_{g,t} + \delta D_{g,t} + \varepsilon E_{g,t}$$

Composite risk score. The weights should be calibrated by venue experience, but the structure forces the dashboard to combine pressure, success rate, wrong-gate behaviour, density and exception load.

Dashboard signals and the decisions they support

DASHBOARD SIGNAL	MEANING	DECISION IT SUPPORTS
Pressure ratio	Demand versus effective capacity	Open lanes, reroute, add handhelds
Time-to-clear	How long the backlog will last	Escalate if clearance goes beyond kickoff
Wrong-gate rate	Gate assignment or communication problem	Push messages, move stewards upstream
Manual override rate	Rules, devices or ticket types causing exceptions	Create exception desk, adjust validation rules
Density zone	Spatial crowd pressure	Change barricades, open spillover areas

The dashboard should not only say what happened. It should indicate what will happen next if operations do nothing.

INTERVENTION PLAYBOOK

How to act before pressure becomes unsafe

Interventions should be pre-agreed before the event. During pressure, teams do not have time to debate authority, permissions or messaging. The model should recommend actions from a prepared playbook.

The intervention playbook

SIGNAL	LIKELY CAUSE	RECOMMENDED ACTION	OWNER
High arrivals, high success	Pure volume surge	Open reserve lanes, deploy handheld pre-validation	Access control lead
High wrong-gate rate	Communication or signage failure	Push gate reminders, move stewards to outer ring	Ticketing + stewarding
Low readiness factor	Fans reaching scanners without tickets open	Create ticket preparation area and broadcast instructions	Fan support team
High exception rate	Duplicate, comp, accreditation or rule issue	Move exceptions to separate desk away from queue	Ticketing supervisor
Density pocket	Bad geometry or blocked route	Adjust barriers, widen queue, remove crossing flow	Security commander

Before kickoff

Interventions should be judged by whether they reduce pressure within the next 5 to 10 minutes.

After kickoff

The priority shifts from speed alone to safe clearance, accurate reporting and controlled late entry.

AFRICAN STADIUM REALITIES

The model must respect local event behaviour

Many imported stadium models assume highly predictable transport patterns, early arrival, strong compliance with signage and uniform digital readiness. African stadium operations need a model that reflects how fans actually arrive and behave.

Late clustered arrival

Traffic, workday schedules, public transport, pre-match social activity and confidence that gates will remain open can concentrate arrivals very close to kickoff.

Mixed ticket formats

Fans may present WhatsApp tickets, PDFs, screenshots, printed tickets, accreditation badges, wallet passes or box-office receipts.

Protocol pressure

VIP, VVIP, sponsors, media, teams, officials and complimentary-ticket holders can create competing flows if zoning is weak.

Mobile money effects

Payment confirmation, ticket retrieval, phone battery and customer-support issues can create queues outside the official access-control line.

The lesson is not that local conditions make modelling impossible. The lesson is that the model must be calibrated locally. Every event should update the assumptions for arrival curve, wrong-gate probability, readiness factor, exception rate, lane cycle time and steward response time.

IMPLEMENTATION ROADMAP

How a stadium can adopt this practically

A stadium does not need a full simulation platform on day one. It can build toward predictive access control in stages.

Six stages toward predictive access control

STAGE	BUILD	OUTPUT	MINIMUM DATA NEEDED
1. Measure	Gate-by-gate scan logs, rejection reasons, device uptime, active lane count	Truthful post-event report	Access logs + device logs
2. Forecast	Arrival curve, gate allocation, peak demand estimate	Pre-event access plan	Ticket sales + historic arrivals
3. Stress test	Late-arrival scenario, queue growth, lane requirement	Capacity risk assessment	Throughput assumptions
4. Simulate	Boids agents, route geometry, barricades, rerouting	Spatial pressure map	Venue map + behavioural assumptions
5. Control	Live pressure dashboard and intervention playbook	Real-time operational decisions	Live scans + steward observations
6. Learn	Compare forecast against actuals	Improved next-event model	Post-event review

DATA REQUIREMENTS

What should be captured at every event

The quality of the model depends on the quality of the event data. The goal is not to collect data for reporting decoration. The goal is to make the next event safer and easier to manage.

Ticketing data

Sales by gate, category, channel, issue time, ticket type, complimentary allocation, accreditation group and ticket delivery method.

Access-control data

Scan time, gate, lane, device, outcome, rejection reason, retry count, admission confirmation, override status and offline validation status.

Hardware data

Turnstile open time, reset time, forced passage alarms, tailgating alarms, device uptime, battery, network latency and scanner health.

Operations data

Active lanes, steward count, barricade layout, weather, kickoff time, road closures, announcements and intervention timestamps.

IMPORTANT

Timestamps are the backbone. Without reliable time logs, the stadium cannot reconstruct arrival pressure, lane performance or the effect of interventions.

KPIS

The access-control metrics that matter most

The best post-event report should help the venue understand whether the access plan worked, where pressure formed, why it formed and what should change next time.

Access-control KPIs

KPI	DEFINITION	WHY IT MATTERS
Peak arrival rate	Highest arrivals per minute by gate	Shows the true stress point
Effective throughput	Clean admissions per minute adjusted for success and uptime	Shows real capacity
Scan-to-admit success	Accepted admissions divided by total attempts	Shows quality of flow
Wrong-gate rate	Wrong-gate attempts divided by total attempts	Shows communication and zoning quality
Average cycle time	Presentation to lane ready for next person	Shows speed of the access lane
Time-to-clear	Estimated clearance time for existing queue	Supports live decisions
Intervention response time	Time from signal to action	Shows command-centre effectiveness

LIMITATIONS

What the model should not pretend to know

A credible model must state its limits. Boids and formulas can help the team reason better, but they do not replace trained stewards, safety officers, police, emergency planning, venue design or professional judgement.

Behaviour is probabilistic

People do not move like machines. Crowd mood, weather, transport disruption, team rivalry, misinformation and panic can change behaviour quickly.

Density estimates need validation

Camera-based or manual density estimates should be checked against field reports, not treated as perfect truth.

Ethics and privacy matter

Any camera, AI or analytics layer should be designed around safety, minimisation, access control and clear governance.

Simulation is only useful if acted upon

A model that predicts pressure but has no intervention authority is only a report. The command structure must be ready to act.

The point of the framework is not to automate crowd safety. It is to give trained operators better early warning, better explanations and better options.

CLOSING POSITION

From ticket validation to access intelligence

The future of stadium access control is not simply faster scanning. It is an integrated access-intelligence layer that combines ticketing, turnstiles, handheld scanners, gate operations, crowd modelling and post-event learning.

Boids gives stadium teams a way to think about how individuals become a crowd. Throughput formulas give teams a way to understand whether gates can cope. Live access data shows what is happening now. The learning loop turns every event into better planning for the next one.

- **The main leaf**

  Do not wait for a crowd to become visibly difficult before managing it. Use ticketing data, admission data and movement modelling to see pressure before it becomes unsafe.

Predict
ARRIVAL PRESSURE

Measure
SCAN-TO-ADMIT REALITY

Learn
FROM EVERY EVENT

References

This paper combines Loft Technologies operational learning with established literature and guidance on boids, sports ground safety, stadium security and pedestrian movement.

SOURCES USED TO GROUND THE FRAMEWORK

1 Reynolds, C. W. (1987). Flocks, Herds, and Schools: A Distributed Behavioral Model. ACM SIGGRAPH Computer Graphics, 21(4), 25–34. DOI: 10.1145/37402.37406. dl.acm.org ↗

2 Sports Grounds Safety Authority. Guide to Safety at Sports Grounds, Sixth Edition, also known as the Green Guide. sgsa.org.uk ↗

3 FIFA. Stadium Safety and Security Regulations. digitalhub.fifa.com ↗

4 Confederation of African Football. CAF Safety and Security Regulations. cafonline.com ↗

5 Confederation of African Football. CAF Safety and Security Officers Handbook. cafonline.com ↗

6 Fruin, J. J. (1971). Pedestrian Planning and Design. Metropolitan Association of Urban Designers and Environmental Planners.

7 Still, G. K. (2014). Introduction to Crowd Science. CRC Press.

8 UEFA. UEFA Safety and Security Regulations. documents.uefa.com ↗