Pneuma Theorem, Inc.

Access Geometry Assessment
Demonstration

Rivergate Yard

The Built-Form Row of a Nine-Cell Assessment

Report PTI–AGA–2026–003

30 July 2026

Document Information

Report identifier	PTI–AGA–2026–003
Document type	Access Geometry Assessment — Demonstration
Subject	Rivergate Yard (fictional) — built-form row of a three-axis nine-cell assessment, standing configuration, pre-event
Framework	Three-axis nine-cell — closed-form, no simulation
Constants	Local frozen copy, aligned to framework v5 (2026-07-26)
Issue date	30 July 2026
Revision	1.0
Classification	Open distribution — may be circulated in full

Author and Contact

Author	Yixuan Tao
Affiliation	PneumaTheorem, Inc., Delaware, United States
ORCID	0009-0005-8136-3886
Email	kevin@pneumatheorem.org
Web	pneumatheorem.org
Framework Zenodo	10.5281/zenodo.21261698

Notice

This document demonstrates the PneumaTheorem, Inc. Access Geometry Assessment: its structure, its assessment criteria and its reporting conventions, applied to a fictional venue (Rivergate Yard) which corresponds to no actual place. The venue, its geometry and its event are illustrative; in a commercial engagement they are supplied by the client or jointly verified, and the distinction between a supplied input and a verified one is recorded in the report. The framework constants are not illustrative — they are declared reference values or measurements, and each is identified as such where it is used.

It is not a certification, an approval, or a regulatory determination, and it does not assert any certifying status. Nothing in it is issued by, or on behalf of, any body empowered to approve a configuration for use. Whether a configuration may lawfully be occupied is a question for whoever is empowered to answer it; what follows is an input to that question and not an answer to it.

The framework establishes what load a geometry permits. It does not establish what load injures a person. Every level compared against here is supplied from outside it, is named where it is used, and is the reader's to accept or replace.

Demonstration Summary

Rivergate Yard satisfies the criteria that apply to it. Egress of 30,000 people completes in 7.95 minutes against a permitted eight; occupancy stands at 2.0 persons per square metre, far below any limit; five gates are distributed on four sides. Nothing in this report disputes any of that.

This assessment asks a question those criteria do not contain: whether the occupiable route admits a connected run of uninterrupted contact, and how long that run is.

It asks that question of the built form alone. What load a geometry permits to accumulate, and what dimension removes that permission, are both properties of the drawing: neither depends on what a venue has installed or on who is on duty. Whether the condition would be seen while the event runs, and what could be done about it on the day, are the other two rows of this framework's matrix; §3 states what each of them asks. The three are separately necessary for the hazard, so what follows is not a determination about Rivergate Yard. It is the one of the three that can be read before anyone is admitted.

The argument runs in a fixed order.

1. **The approach route is a hundred metres of connected space, walled on both flanks, with nothing dividing it.** At the spacing between successive load-bearing contacts that is 265 contributors in depth. The route is closed ahead by the gate line and fed from the street behind: a run held at both ends.
2. **The walls are what put it on the accumulating side, and not anything the crowd does.** Accumulation requires that each contributor deliver more than a receiving body absorbs. An open region distributes its occupants among several directions of travel, and at the divided alignment that follows, the per-contributor exertion falls below what is absorbed — nothing accumulates, at any length. A run walled on both flanks admits one direction, the division does not occur, and the same people exerting the same amount are on the other side of that threshold.
3. **So the yard is not the finding and the approach is.** The standing area holds four times what the route can and does not accumulate, because it is open. Occupancy does not distinguish the two; the boundaries do.
4. **Unpartitioned, the guaranteed load in the route is 19,080 N** against an adopted level of 1100 N. The figure is computed at the design end of the alignment factor — the end that makes a guarantee harder to claim, and not a prediction of what occurs — and it is what the built form permits before any arrangement is made. The level is supplied from outside this framework, is named where it is used, and is the reader's to replace; the arithmetic is unchanged in form if they do.
5. **And the same route carries 8.0 m of the venue's exit width.** Four of the five gates open directly onto the yard. The fifth is this route, used in the other direction, and it is the one width whose approach admits a single direction.

The finding therefore has a consequence for capacity, and it is stated as a choice.

Either the route is occupied in groups no longer than 6.8 m, separated by clear intervals of 12 m — on this route, 5 intervals giving groups of 6.7 m, by whatever means the venue finds workable —

or 38.0 m of exit width is counted rather than 46.0 m, which reduces the permitted capacity by about 17 per cent, or 5,248 people. At the population sold this is not a reduction in headroom but a shortfall of 5,072 people, and the eight-minute calculation is not satisfied.

Why compliance did not find it. The applicable criteria are expressed in width, rate, density and time. **None of them is how many directions a region admits.** Increasing the population by a fifth moves this venue from compliant to non-compliant and changes not one figure in this report; removing the walls from the approach closes this report's finding and changes not one figure in the compliance calculation (§4).

A third option is geometric, and §5 sets it out. A queue forms upstream of whatever controls its rate. Moving the control point from the head of the ramp to the street end of the underpass moves the queue with it, into the open forecourt — where a region admits several directions of travel and the mechanism does not operate at any length. The route then carries people rather than holding them, and partitioning it becomes a provision against a downstream stoppage rather than a standing requirement. Relocating a gate line has consequences for ticketing, search, weather protection and the public highway which this report has not assessed and is not competent to assess.

On the grouping, this report supplies a dimension and not a method. Any particular scheme — batched admission, holding points, managed release — achieves the grouping under some conditions and fails under others, and a scheme carried out at seven parts in ten is not the scheme. What is stated is a maximum group length, an interval, and the arithmetic connecting both to the declared inputs. How the dimension is met is for the venue, which knows its own constraints; a criterion that can be checked is worth more than a method that must be trusted.

This report contains no claim that Rivergate Yard is safe or unsafe. The framework establishes what load a geometry permits; it does not establish what load injures. Rivergate Yard is fictional and corresponds to no actual place.

Contents

1. Scope of Analysis

2. Assessment Matrix

Cell 1 — Hardware × Space

The question this cell answers. How many people these spaces can physically hold.

Not whether that is safe — that is Cell 4. Not whether the crowd will reach that state — that is Cell 7. Not who manages it — that is Cell 3. This cell does one thing: it measures the occupiable space and multiplies by the geometric packing limit.

L0 — Basis Criterion.

$$N_{\max} \;=\; \rho_{\max} \times A, \qquad \rho_{\max} = 7 \text{ persons/m}^2$$

$\rho_{\max}$ is the geometric packing limit — the greatest areal density attainable by upright, coplanar bodies (Tao, 2026). It assumes nothing about what anyone does, and is therefore the one quantity in this matrix that rests on no behavioural premise.

Why 7 and not 4 or 5. The lower figures in common use are *safety thresholds*. The 7 used here is a *physical limit*. They are not the same kind of number and are not interchangeable:

- 7 answers "how many fit" — beyond it the population is no longer upright and coplanar but already compressed, angled, or off its footing.
- 4 to 5 answers "where does danger begin" — which is the work of Cell 4.

This report offers no proof that densities above 7 are safe, and makes no claim that densities below it are safe.

Node or link — each space is one, never both. A space contributes either its *area* (as a holding node) or its *width* (as a flow link). Counting the same space both ways is double counting. A link contributes a width and a rate, which are treated on the time axis (Cell 7); it carries no capacity here.

L0 — Venue side (supply) Rivergate Yard is a former railway goods yard in mixed event use. The configuration assessed is Open-air standing, single stage, with a stage at the north end and admission from three of the five gates.

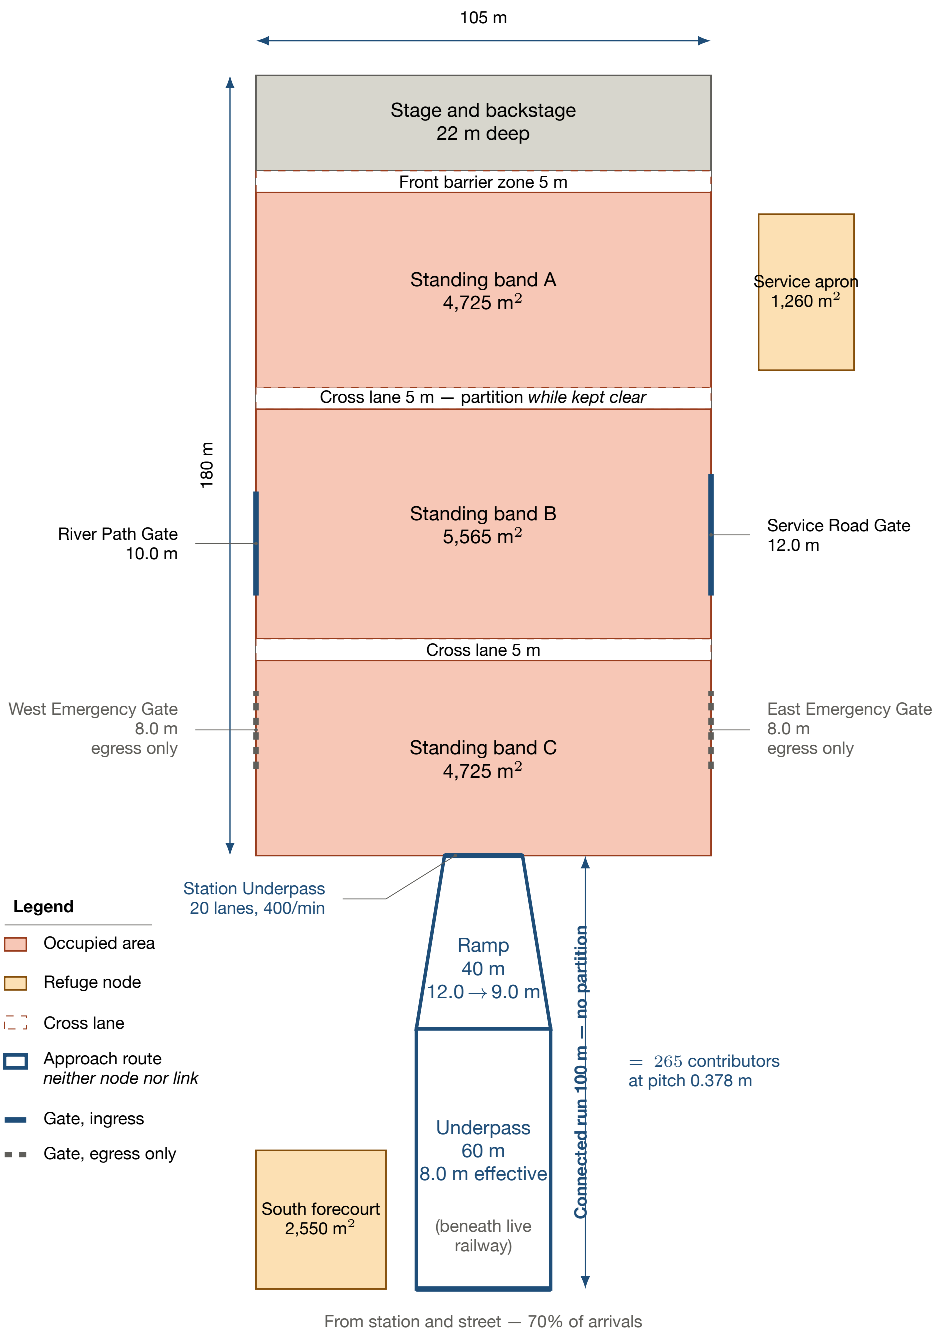

This drawing defines the scope of the assessment. The segments shown are the segments analysed; anything not drawn is outside it. Rivergate Yard is fictional and corresponds to no actual place.

Figure 1: Rivergate Yard — standing configuration site plan. **This drawing defines the scope of the present report.** The segments shown are the segments analysed; anything not drawn is outside the scope of this assessment. Where the actual space differs from this drawing, or is used in a manner this drawing does not cover, the conclusions of this report do not apply to that situation.

Segment	Role	ℓ (m)	w (m)	A (m^2)	N_{max}
Occupied areas — standing configuration					
Standing band A (north)	occupied	45	105	4,725	33,075
Standing band B (centre)	occupied	53	105	5,565	38,955
Standing band C (south)	occupied	45	105	4,725	33,075
Approach route A — see the note below					
Underpass	*see note*	60	8.0	480	—
Ramp	*see note*	40	10.5	420	—
Refuge nodes — checked per node, never summed					
South forecourt	node	85	30	2,550	17,850
Service apron	node	70	18	1,260	8,820
Gates — width only; no capacity in this cell					
Station Underpass	link	—	8.0	—	—
Service Road Gate	link	—	12.0	—	—
River Path Gate	link	—	10.0	—	—
East Emergency Gate (egress only)	link	—	8.0	—	—
West Emergency Gate (egress only)	link	—	8.0	—	—

The two cross lanes through the standing area (5 m each, 1,050 m^2 in total) are excluded from the occupied figures above. They are partitions *while kept clear*; whether they are kept clear is an L2 question.

The approach route does not fit the classification Route A is designed as a link. It carries 70% of arrivals — 21,000 people — from the station to the gate line, and the drawings give it as a width, which is how a link is specified.

But it has an area. The underpass and ramp together occupy 900 m^2, and at the packing limit that area holds 6,300 people. A link with a capacity is a node, and the classification that keeps this matrix consistent has nothing to say about which it is.

Which it is depends on whether it is clearing. While the gate line admits at least as fast as arrivals reach the ramp, the route carries flow and its area is irrelevant. Where arrivals exceed admission, the difference accumulates within the route, and the area becomes the operative quantity. The route is a link at one arrival rate and a node at another, and nothing in its construction changes at the crossing.

The accounting therefore fails exactly where it matters. A node–link classification made on the drawing is a statement about intended use. The transition it cannot represent is the one that occurs when the intended use stops being achievable, which is the condition under which the route matters at all. This report accordingly declines to classify Route A in the table above, and treats it in both capacities: as a link in Cell 2, and as an occupied region in Cell 4.

No capacity figure is entered for it here. Entering one would assert the node reading; leaving the row out would assert the link reading. Both would be a determination this cell is not entitled to make.

L0 — Event side (demand) Standing configuration: 15,015 m^2 of net standing area at the planning density of 2.0 persons/m^2, giving a population of 30,030, sold at 30,000.

Of these, 70% — 21,000 people — are expected to arrive via Route A, the station approach. That expectation is an event-side input and is not independently verified; the basis for it is stated in Cell 2.

L0 — Determination **The occupied area is not the constraint.** The three standing bands hold 30,030 at the planning density against a combined physical capacity of over one hundred thousand at the packing limit. The occupied area is generously sized for the population, and nothing in this cell suggests otherwise.

The refuge nodes are read individually. The south forecourt holds 17,850 and the service apron 8,820, both at the packing limit rather than at any density at which people could reasonably be held. Neither is sized against the whole population, and they are not added: a figure formed by adding them would conceal the case in which one is full and the other empty.

The approach route is the finding. Its 900 m^2 hold 6,300 people at the packing limit — a fifth of the event population, in a space the drawings describe by its width. Whether that capacity is reached is Cell 7; what load it would impose is Cell 4. This cell records only that the capacity exists and that the drawings do not show it.

What this cell does not determine. Whether these capacities suffice depends on how much the gates carry away within the permitted period, which is Cell 7, and on how the population is allocated between routes, which is Cell 3. This cell supplies the static term only.

L1 — Measures The only variable in this cell is A, so the measures are two:

- **Increase occupiable area** — bring external space into scope, remove installations consuming effective width, or restore an area currently allocated to another use.
- **Reduce simultaneous occupation** — admission control, phased arrival, reduced ticket allocation.

Partitioning is not in this cell. A physical partition changes the length of a coherent path rather than an area; its spacing criterion comes from Cell 4.

L2 — Execution and traceability This layer records where as-built departs from design.

Known risk: effective width. w is the quantity most liable to drift between drawing and site, and it enters A linearly. The underpass is tabulated at 8.0 m effective against 9.0 m structural, the difference being a drainage channel, lighting conduit and maintenance clearance.

Temporary fencing, concession units, stored materials and works hoardings all consume effective width further and are frequently absent from the drawings.

Known risk: which gates are open. The drawings show 5 gates. During admission only 3 are open, since East Emergency Gate and West Emergency Gate are egress-only. A reader taking the gate count from the drawing will overstate the connectivity available during the period this report is concerned with.

Status for this assessment. All widths tabulated above are drawing values and **have not been verified on site**. This item is recorded as *no data*, not as a pass.

Where no site verification capability exists, this layer is marked "not verified on site" rather than left blank; a blank reads as a pass. The relevant detection capability is Cell 2; where that cell fails, this layer cannot be determined and the two move together.

Cell 4 — Hardware × Force

The question this cell answers. What load the built form permits to accumulate on one person, if the route fills.

Not whether it will fill — that is Cell 7. Not whether anyone would know — that is Cell 5. This cell takes the geometry of Cell 1 as given and computes what it allows.

L0 — Basis **The relation.** Where contact is continuous, the load borne by one individual accumulates along the connected run reaching them:

$$F \;=\; C \cdot F_{\text{den}} \cdot \Phi(L), \qquad F_{\text{den}} = 480\ \text{N},$$

where F_{den} is the declared exertion scale, C is the per-contributor exertion expressed on that scale, and $\Phi(L)$ sums the contributions along a run of depth L (Tao, 2026). At the design baseline $\Phi(L) = L$.

Accumulation is not automatic. It occurs only where each contributor delivers more than a receiving body absorbs without passing on:

$$\beta f \;>\; \delta, \qquad f = C \cdot F_{\text{den}}, \qquad \delta \approx 30\ \text{N}.$$

Where this fails, each person bears one neighbour's exertion at any depth and for any duration, and no partitioning is required. This is the first question asked of the scene, because where it fails the rest does not arise.

The residual enters the criterion and not the bounds. The value above is used to decide whether a run accumulates, where a value is unavoidable and the adverse end of the measured range is therefore taken. Every bound this cell reports is computed at $\delta = 0$, which is the conservative choice for a bound and the uninformative one for a criterion. The division is set out in §4, which also states what a properly established residual would be worth in metres.

The level compared against is supplied from outside. This framework establishes what load a geometry permits; it does not establish what load injures. The values available fall into two classes answering different questions:

- *Acute structural failure* — whether the ribcage breaks. Low thousands of newtons, from biomechanical models (Kroll et al., 2017).
- *Sustained compression* — whether respiration is defeated by a load held for minutes. About a thousand newtons, from retrospective estimates (Hopkins et al., 1993).

This assessment adopts the sustained class, at 1100 N, because the exposure at issue is a queue held for tens of minutes rather than an impact. It is the strictest of the levels available, and the arrangement this report arrives at is correspondingly the least permissive. Had the structural class been adopted, the permitted run would be some two to four times longer — and the arrangement of §4, which brings the borne load to 1,055 N, would sit within the structural levels by that same factor of two to four. Against the level actually adopted the margin is some four per cent. Adopting the stricter level therefore costs occupiable length rather than closing off the answer. That is the correct arithmetic for a different question — whether the geometry could break a ribcage — which this report does not raise.

The provenance of the adopted value is not strong, and the chain is worth stating. It reaches the crowd-safety literature through Hopkins et al. (1993), who relay it rather than report it; Madea (2020) records the figures as descending from work referenced from the United Kingdom Home Office, and that primary work has not been located. A reader preferring a different level substitutes it and the arithmetic below is unchanged in form.

L0 — Venue side (supply) The built form supplies three things to the relation above.

Depth. The connected run of Route A is 100 m with no partition along it. At the chain pitch of 0.378 m — the spacing between successive load-bearing contacts at the packing limit — this is

$$L = \left\lceil \frac{100}{0.378} \right\rceil = 265 \text{ contributors.}$$

Rigid boundaries at both ends. The gate line closes the run ahead: it is a structure, and the load bearing on it does not relieve through it. Behind, the underpass is fed from the street, so the upstream end is a source rather than an outlet. A run held at both ends is the configuration the relation describes.

Alignment. The route is walled on both flanks along its whole length and narrows from 12.0 m to 9.0 m at the gate line. Lateral movement is unavailable and the geometry orients exertion along the route axis.

Which value of the alignment factor, and where it comes from. The factor is the mean projection of the contributions onto the route axis. Two things set it, and only one of them has anything to do with what the crowd intends.

How many directions the region admits. An open region distributes its occupants among several directions of travel; those heading along a given axis contribute to it and the others do not. A run walled on both flanks admits one direction, so no such division takes place. That difference alone is a multiple of about three, and it follows from the walls — not from a target, an instruction, or anyone's intention.

How wide the spread is within that direction. Within a single direction the contributions are spread over a sector, and the mean projection over the framework's reference sector is about 0.83 (Tao, 2026). A common visible target narrows the spread and carries the factor toward unity.

Route A has both features. It is bounded on both flanks along its whole length, which supplies the first, and it terminates in an admission point, which supplies the second. A design guarantee takes the adverse end of an admissible range, because overstating the load makes the guarantee harder to satisfy and therefore harder to claim wrongly. This assessment accordingly takes $C_h = 1$.

And the standing area has neither. The yard is open on its long axis, is crossed by lanes, and is entered and left in several directions. Its alignment factor is the divided figure, about 0.28, and the determination below returns to what follows from that.

A retrospective analysis would take the floor instead. Arguing that a load *was* reached calls for the lowest defensible alignment, and obtains a longer permitted run from the same arithmetic. The two are different questions, and the difference between them is not a disagreement.

L0 — Event side (demand) **The event supplies the target.** The gate line is a common visible objective for everyone in the route, which is what makes the alignment factor above applicable at all. This is not a behavioural assumption about crowd psychology; it is a statement that the population is trying to reach one place and can see where it is.

It is also not avoidable. An admission point is a common target by construction, and no operating measure removes it.

The event supplies the duration. Admitting 21,000 people at 400 per minute takes 53 minutes. Individuals in the route are therefore standing for tens of minutes, which bears on the sustained fraction below.

The sustain factor, and a route that does not depend on it. The vertical factor is taken at the declared reference $C_v = 0.15$, and with the alignment at unity,

$$C \;=\; C_h\, C_v \;=\; 0.15, \qquad f \;=\; C \cdot F_{\text{den}} \;=\; 72\ \text{N}.$$

That reference rests on a proposition about sustainable muscular exertion (Rohmert, 1960) which the later literature does not support: the original determination treated a hold exceeding fifteen minutes as of unlimited duration (Hagberg, 1981), and fatigue is observed well below the stated fraction (Sjøgaard et al., 1988). A reader is entitled to ask what becomes of the figure.

Nothing, because the same value is available without any claim about sustained exertion. A body inclined forward by about six degrees and supported in front of it delivers, at this alignment, a horizontal force of about 72 N — and its source is weight rather than muscular effort. Six degrees is not a posture an observer would remark on in a wedged queue, and it is not one anyone has to intend. The fatigue that lowers the sustainable fraction is the same fatigue that produces the inclination, so the two routes to the figure fail in opposite directions. Over the tens of minutes this route holds people, the second is the one that applies.

L0 — Determination **First: does it accumulate?** With $f = 72$ N against $\delta \approx 30$ N, the growth condition is satisfied by a factor of about two and a half. The route is on the accumulating side.

And the standing area is not. At the divided alignment factor of an open region the per-contributor exertion is about 20 N, which is *below* the residual. Nothing accumulates in the yard, at any depth and for any duration, and no partition is required there.

And what would move it. That figure uses the framework's three-family bound on alignment, which is the low end of what an open region admits. Crossing the residual would require the per-contributor exertion to exceed 30 N, which is an alignment factor above about 0.42 — a yard whose occupants were oriented as a single direction family rather than divided among several. The determination holds while the yard is entered and left in several directions, which the drawing shows and which §4 tests.

So what separates the two is the boundaries. The same people, exerting the same amount, are on opposite sides of the growth condition according to whether the region they stand in admits one direction or several. Occupancy does not distinguish them — the yard holds four times as many people as the route can, and the yard is the one that does not accumulate.

Which is what this report's finding rests on. It is not that the approach is long. It is that the approach is walled; that walling supplies the alignment on which the growth condition turns; and that the length then determines how much accumulates once it does.

Second: how much? Unpartitioned, at the values above:

Declared C	Basis	$F_{\text{guar}} = CF_{\text{den}}L$
0.075	C_h at its floor — the forensic reading	9,540 N
0.15	**C_h at unity — the design reading**	**19,080 N**
0.30	C_h at unity, with a lean of about twelve degrees	38,160 N

The middle row is the one this assessment operates on. Every entry exceeds the adopted level by a wide margin, and the design value exceeds it by more than seventeen times. The route as built does not admit a bound below the level at any of these values.

And an unpartitioned reading is not the assessment. The figures above are what the built form permits before any arrangement is made. They are the starting point of the assessment rather than its outcome: what determines the outcome is whether an arrangement brings the borne load within the adopted level, and §6 states the three states that follow.

Third: what would. Inverting the relation gives the partition spacing at which the guarantee meets the adopted level:

$$s_{\max} = p \left\lfloor \frac{F_*}{C\, F_{\text{den}}} \right\rfloor .$$

Declared C	Depth admitted	Spacing
0.075	30	11.3 m
0.15	**15**	**5.6 m**
0.30	7	2.6 m

This is the bound the cell delivers. It is not a verdict and not a recommendation: it is the depth, and hence the spacing, at which the guaranteed load meets the adopted level at the design value of the exertion. *How* that bound is met is a separate question — §4 and §5 set out the arrangements available and why a rigid partition is not the one this report favours. The spacing is inversely proportional to the declared C, exactly and with no interaction, so a reader who prefers a different declaration reads the answer off the same table — and a reader who prefers the forensic reading of the alignment obtains 11.3 m, which this assessment does not adopt and states for comparison.

A check against a second route through the framework. The coherent length at which the adopted level is reached under the design value is 6.05 m. The framework's treatment of the reversal problem, working from the same declared quantities but by a different route, obtains six metres for the same configuration (Tao, 2026). The agreement is arithmetic rather than evidential — both use the framework's declarations — but it confirms that the two are computing the same thing.

What this cell does not determine. Whether the route will fill is Cell 7. Whether anyone would know is Cell 5. Whether the adopted level is the right one is a question for the body that sets the applicable standard, and this report does not answer it.

L1 — Measures

- **Bound the connected run** at the spacing tabulated above. This is the only class of measure that removes the condition rather than managing it. The bound can be met by occupying the route in separated groups (§4) or by an installed partition; an installed partition is effective only if it interrupts load-bearing contact — a structure that narrows the route without interrupting contact concentrates the chain rather than breaking it.
- **Reduce the connected run** by relocating the gate line into the route, so that admission occurs partway along rather than at its head. This shortens L directly.
- **Reduce alignment** — not available here. The target is the admission point, and it cannot be removed while admission occurs.

Reducing occupancy is not a measure in this cell. The relation contains no headcount: two hundred people at a given partition spacing and two hundred thousand at the same spacing produce the same guarantee. What is bounded is the length of a connected run, and admission control does not change it.

L2 — Execution and traceability **Known risk: a partition is only a partition while it is clear.** The two cross lanes through the standing area are partitions on the drawing. Whether they are kept clear during the event determines whether they are partitions in fact, and a lane occupied by standing spectators is not one. The same applies to any partition installed in Route A.

Known risk: temporary structures that narrow without interrupting. Fencing erected to guide queueing narrows the route and does not break contact. It looks like a partition on a plan and is the opposite of one in the relation above.

Known risk: the declared C is not measured. There is no instrument for it (Cell 5). The tables above are accordingly a sensitivity rather than a prediction, and a report using a single column of them as a result has overstated what is known.

Status for this assessment. No partition exists in Route A. This is a design condition and is recorded as such, not as an execution failure.

Cell 7 — Hardware × Time

The question this cell answers. Whether the route reaches the state Cell 4 computes, and how long that takes.

Cell 4 establishes what the geometry permits if the route fills. This cell asks whether it does.

L0 — Basis **Accumulation requires an excess sustained over time.** Occupancy of the route rises at the difference between the rate at which people reach it and the rate at which the gate line admits them. Where that difference is negative or zero the route carries flow and nothing accumulates; where it is positive the occupancy climbs at that rate.

The state that matters begins before the space is full. The packing limit of 7 persons/m^2 is the greatest density upright bodies can reach. Continuous contact — the condition under which the relations of Cell 4 apply — begins at some *lower* density, because contact becomes continuous before further compression becomes impossible.

That density is not derived. The framework states the criterion for the transition in terms of the floor area a recovery step requires, and does not evaluate it (Tao, 2026). Its absence has a specific consequence here, stated below.

What the framework declines to predict, and what it does not. The rate at which a crowd loads a geometry is not predicted, because it would require a behavioural quantity the framework does not possess. That refusal does not extend to a rate fixed by a timetable. Arrivals at this site are delivered by scheduled services, and the number a service delivers is not a matter of crowd behaviour. Where the arrival rate is scheduled it can be computed; where it depends on when people choose to travel it cannot, and both apply here in different proportions.

L0 — Venue side (supply) The built form supplies two terms.

Capacity. Route A holds 6,300 people at the packing limit (Cell 1). This is the endpoint of filling, not the point at which the state of Cell 4 begins.

Drain. The gate line admits at 400 per minute across 20 validated lanes. This is the only outflow from the route. It is fixed by the number of lanes and their individual rate, both of which are design quantities.

Neither term is adjustable during the event beyond the reserve lanes noted in Cell 3. The capacity is built and the drain is installed.

L0 — Event side (demand) **On average there is no accumulation.** 21,000 people arriving over a three-hour admission window is 117 per minute against a drain of 400. The average case has a wide margin.

Averages are not the operative figure. Arrival at a station-served site is delivered in bursts and is weighted toward the last hour before the principal performance. The table below states what follows from a sustained excess, without asserting that any of these rates occurs:

Arrival (per min)	Excess (per min)	To packing limit (min)	After 15 min (persons/m^2)	After 30 min (persons/m^2)
480	80	79	1.3	2.7
560	160	39	2.7	5.3
800	400	16	6.7	—

The arrival profile is an event-side estimate. Cell 2 establishes that it is not checkable during the event, so the row that applies is not knowable at the time it would matter.

L0 — Determination **The computable figure is an upper bound, and it is optimistic.** Given an excess rate, the time to reach the packing limit follows by division. But the relations of Cell 4 apply from the point at which contact becomes continuous, which is earlier and by an amount this framework does not determine. Every figure in the third column above is therefore longer than the time actually available, and by how much is not stated because it is not known.

Two uncertainties compound. The excess rate depends on an arrival profile that is an estimate and cannot be checked while the event runs (Cell 2). The onset density is not derived. A time-to-onset would require both, and neither is available.

Verdict: not determined. This cell cannot state whether or when the route reaches the state of Cell 4. That is recorded as *no data* and not as a pass, and it is not a defect of this assessment: the quantity is not available from any source known to this assessment.

What follows from an undeterminable time axis. The three axes are each necessary for the hazard, so any one of them failing is a route to safety. Two of them can be closed by determination:

- *Space* — closed by bounding the connected run (Cell 4). Requires no rate and no arrival profile.
- *Force* — would be closed by failing the growth condition. Not available here: the route has a common target by construction (Cell 4).
- *Time* — would be closed by showing the state is not reached. Not available here, for the reasons above.

So the space axis is not one option among three. It is the only one of the three that can be closed at this site, and it can be closed without resolving anything on this axis. That is the argument for treating the question at the design stage: a configuration made geometrically incapable of accumulation does not require the time axis to be determined at all.

This is the framework's general position and it is worth stating in the site's own terms. A demonstration that something cannot occur is available here. A prediction of when it would occur is not, and would not be defensible if it were offered.

L1 — Measures

- **Raise the drain.** Additional validated lanes increase 400 and reduce the excess directly. This is the one term the venue controls, and Cell 3 notes that holding some in reserve is what makes it a lever at all.
- **Flatten the arrival profile.** Staging services reduces the peak. Cell 3 records that this lever is held by the transport operator.
- **Close the space axis instead.** Partitioning at the spacing of Cell 4 makes this cell's determination unnecessary. It is listed here because the honest answer to an undeterminable time axis is not a better estimate.

L2 — Execution and traceability **Known risk: the drain is assumed constant.** 400 per minute presupposes twenty lanes running at rate. A raised search level, a lane fault, or a ticketing system fault reduces it, and the excess is the *difference*, so a reduction in the drain has the same effect as an equal increase in arrivals. A twenty-five per cent reduction in lane rate turns an arrival of 400 per minute — which produces no accumulation — into an excess of 100.

Known risk: the profile is taken from elsewhere. The arrival distribution is drawn from comparable events at other sites. Nothing in the record establishes that this event's audience travels the same way, and Cell 2 establishes that a departure from the assumption would not be detected.

Status for this assessment. The onset density is not derived and the arrival profile is not verifiable. Both are recorded as *no data*. The upper bound in the table above is reported as an upper bound and is not a time available.

3. The Two Rows Not Assessed Here

The framework this report applies arranges an assessment as a matrix. The three columns are the axes on which a compressive hazard depends: the *space* a configuration provides, the *force* it permits to accumulate, and the *time* in which that state is reached. The three rows are the kinds of thing that can act on each axis: the *built form*, which is fixed on the day; the *instrumentation*, which determines whether a condition can be seen while the event runs; and the *management*, which determines what can be done about it once it is seen.

This report is the built-form row. Cells 1, 4 and 7 above are its three entries. The other six are named here, with what each of them asks, so that the boundary of this document is visible rather than implied.

3.1 Why the built-form row can be read on its own

It is the only row that can be answered from a drawing. The quantities it uses — lengths, widths, areas, the pitch of load-bearing contact, the alignment a boundary admits — are all present in the plan before anyone is admitted, and none of them depends on what equipment is installed or how the event is staffed. That is why an assessment of this row can be produced without a site visit, without access to a control room, and without knowledge of an operator's procedures.

The other two rows cannot be answered that way. What a venue can see depends on what it has installed and where; what a venue can do depends on who is on duty, what they are authorised to decide, and how long a decision takes to reach the ground. Those are facts about an organisation, not about a geometry, and they are supplied rather than derived.

3.2 What the instrumentation row asks

- **Cell 2 — instrumentation × space.** Can the quantity that governs this determination be measured while the event runs? Not whether the spaces are large enough, which is Cell 1, but whether the occupied extent of a route is observable at the time it matters.
- **Cell 5 — instrumentation × force.** Can the load itself be measured on a live crowd? The framework declares a per-contributor exertion rather than measuring it, and this cell is where that gap is stated rather than concealed.
- **Cell 8 — instrumentation × time.** Can a trend be seen while there is still time to act on it? A reading that arrives after the state has formed answers a different question from one that arrives before.

3.3 What the management row asks

- **Cell 3 — management × space.** Knowing the geometry, can anything be done about it, and does that action do what is wanted? This is where a measure is tested against its own side effects rather than against its intent.

- **Cell 6 — management × force.** Knowing the load, what can be done about it? The force axis admits fewer levers than it appears to, and this cell sets out which of them act on the quantity and which only appear to.
- **Cell 9 — management × time.** Once the state has formed, what can be done, and how fast? A measure that is correct but slower than the process it addresses is not a measure against that process.

3.4 What the full matrix adds, and what it does not

The three axes are separately necessary. A compressive hazard requires all three: a geometry that permits the load, a state that is reached, and an interval in which it is held. Any one of them failing is a route to a determination, which is why a finding on one row is not a finding about a venue.

Within a column, the weakest entry governs. Whether a condition can be detected is not improved by the built form being well instrumented elsewhere; whether it can be acted on is not improved by it being visible. A complete assessment compares the three entries of each column and reports the weakest, and that comparison cannot be made from one row.

What this report therefore establishes is narrower than a venue assessment and is not a part of one. It answers a question that is complete in itself: what load this built form permits to accumulate, and what dimension removes that permission. Both are properties of the drawing. Neither depends on what is installed or who is on duty, and neither is altered by the answers to the six cells named above.

Nothing in this section states an answer for any of the six. Their answers depend on the venue's own systems and procedures, which are not inputs to this report.

4. Sensitivity: What Each Criterion Responds To

The matrix above returns a finding that the applicable compliance checks do not. That could mean the framework is a stricter version of the same test. It is not, and the difference is demonstrable rather than asserted: the two respond to different perturbations, and the response of each to the other's governing variable is nil.

Perturbing the population

The population is the quantity a venue adjusts when a capacity question arises, and it is the first remedy proposed when an approach is congested. The table below varies it and reports both criteria.

	Compliance			*This framework*		
Population	Egress (min)	Verdict	Density (per m^2)	L	F_{guar} (N)	$s_{\max}$ (m)
24,000	6.36	pass	1.60	265	19,080	5.6
30,000	7.95	pass	2.0	265	19,080	5.6
33,000	8.75	**fail**	2.20	265	19,080	5.6
36,000	9.54	**fail**	2.40	265	19,080	5.6
45,000	11.93	**fail**	3.00	265	19,080	5.6

A twenty per cent increase moves the venue from compliant to non-compliant on egress. Not one figure on the right of the rule changes.

The reason is visible in the relations. The two quantities this framework reports are

$$F_{\text{guar}} = C \cdot F_{\text{den}} \cdot \left\lceil \frac{\ell}{p} \right\rceil, \qquad s_{\max} = p \left\lfloor \frac{F_*}{C\, F_{\text{den}}} \right\rfloor,$$

whose terms are a declared exertion, a declared scale, the length of a connected run, and the pitch. None of them is a number of people. What is bounded is the length of a run, and a run of given length imposes the same load whether two hundred people or two hundred thousand stand behind it.

The remedy that does nothing

The row for 24,000 is the one worth dwelling on. Reducing the population by a fifth relieves the egress calculation comfortably — and leaves the guaranteed load and the required partition spacing exactly as they were.

Admission control, phased entry and reduced ticket allocation all act on the same variable. Against this mechanism they are inert, and a venue that adopts them in response to a congested approach has incurred the cost of the measure without changing the condition it was taken for.

This is not a criticism of those measures. They act on the quantities the compliance criteria are expressed in, which are real quantities that matter for other reasons. It is a statement about which hazard responds to which lever.

And what occupancy does govern

The foregoing invites an obvious question. If the load does not respond to how many people are present, why is occupancy the quantity the whole field works in?

Because it governs something else, and that something matters. A fall is a common event. Most end with the person standing up again. What decides whether one becomes something else is whether recovery remains available — whether there is floor space to rise into, and whether the load from behind relents long enough for rising to be possible.

That is what occupancy determines. Not whether someone loses their footing — that is a matter of what reaches them and what they can withstand, neither of which is a density. Occupancy determines whether the space a recovery requires still exists (Tao, 2026).

Question	Governed by
Does a load accumulate along a run?	the boundaries, via alignment
How much does it reach?	the length of the run
Can someone who falls get up?	**occupancy**

So the density criteria in use are not measuring nothing. They are measuring the third row, which is a real hazard with a real remedy, and they have been doing so for decades. Nothing in this report displaces them.

What this report adds is the first two rows. They are not visible to a density criterion, they are not made visible by measuring density more carefully, and at Rivergate Yard the first of them is where the finding lies.

The two are complementary rather than competing, and a venue satisfying one has not thereby addressed the other.

Perturbing the geometry

The converse test is the more informative of the two, and it has two parts because this framework can be satisfied in two different ways.

The first way bounds the load. Suppose the approach were partitioned so that the longest run of uninterrupted contact were 5.6 m rather than 100 m, with nothing else altered — same population, same gates, same exits.

The second way prevents the accumulation altogether. Suppose instead that the queue formed in the open forecourt rather than in the walled route. An open region distributes its occupants among several directions of travel, so the alignment factor is divided among them; at the divided figure the per-contributor exertion falls below what a receiving body absorbs, and nothing accumulates at any length.

	This framework			*Compliance*
Configuration	Directions	f (N)	Verdict	Egress (min)
As designed, 100 m walled	one	72	**not closed**	7.95
Partitioned to 5.6 m	one	72	closed, load bounded	7.95
Queue in the open forecourt	several	20	closed, no growth	7.95

Neither change moves the compliance figure by a single digit, because neither the length of a connected run nor the number of directions a region admits appears anywhere in it.

And the two rows close the finding differently. Partitioning leaves the mechanism operating and bounds what it can reach. Relocating the queue stops it operating at all. The second is the stronger result and, at this site, the cheaper one — which is the subject of §5.

What the boundaries do, and why it is not a matter of degree

The third row above is worth dwelling on, because it is where this framework departs furthest from the reading it is usually compared with.

On the familiar reading, a walled approach is a place where a crowd happens to get denser, and the walls make a bad situation worse. On the present reading the walls are not aggravating: they are the condition. A crowd in an open region and the same crowd in a walled one, at the same occupancy and exerting the same amount, are on opposite sides of the threshold at which accumulation begins — and nothing about the people distinguishes them.

Which inverts a common formulation. A large space is not dangerous because it is large: in an open region the exertion is not aligned and the accumulation does not begin, whatever the extent. What a large space permits is that enough people be admitted to fill it, and whether they are is a question about admission rather than about the space. A narrow walled run is unfavourable directly, and without any question of numbers, because it supplies the alignment on which the threshold turns.

Where the two do meet: the egress calculation

The preceding says the two criteria respond to disjoint quantities. There is one place where they meet, and it is where this assessment has a consequence for capacity.

The egress calculation totals the exit widths. 46.0 m across five gates discharges 30,000 people in 7.95 minutes against a permitted eight. The calculation assumes each gate performs — not that each is equally used, but that each is available and that what passes through it keeps moving.

This framework asks something of each gate's approach. Not its width, which the calculation already has, but how many directions the approach admits:

Gate	Width	Approach	Directions
Service Road Gate	12.0 m	from the yard	several
River Path Gate	10.0 m	from the yard	several
East Emergency Gate	8.0 m	from the yard	several
West Emergency Gate	8.0 m	from the yard	several
Station Underpass	8.0 m	**through the walled route**	**one**

Four of the five open directly onto the yard. A crowd leaving through them is in an open region, the alignment factor is divided among the available directions, and by §2 nothing accumulates.

The fifth is the route this report has been concerned with, used in the other direction. Egress through Station Underpass passes down the ramp and through the underpass — 100 m, walled on both flanks, one direction. If anything downstream of it stops, that is the configuration of §2, and the guaranteed load is what Cell 4 reports.

Which is a width the framework cannot warrant on the same terms as the other four. The arithmetic of setting it aside is straightforward:

Exit width counted	Capacity at eight minutes	Time for 30,000	
46.0 m (all five)	30,176	7.95 min	pass
38.0 m (four)	24,928	9.63 min	**fail**

The difference is 5,248 people, or about 17 per cent of the permitted capacity.

What would remove the condition, and what this report does not supply

The route admits one direction because it is walled along its whole length. Two things follow, and only the second is available without building something.

Lateral relief removes the cause. Openings in the flanks restore the alternative directions, and by §2 the accumulation does not then begin at all. That is a structural alteration beneath a live railway, and this report offers no view on its feasibility or cost.

Dividing the run bounds the effect. A run occupied in separated groups carries a free edge at the head of each. An occupant at such an edge has nobody to deliver force to, so they are not a contributor; and the same body is then available to resist what arrives from behind. The edge is constituted by the occupants themselves — nothing is installed, nothing is rigid, and nothing has to be strong.

Two bounds, and they answer different questions

The first bounds the group. A group's own accumulation must stay below the adopted level, which for this route's geometry gives a maximum group length of about 6.8 metres.

The second bounds the interval. A free edge persists only while what arrives at it stays below what its front occupant can hold without stepping, and that is reached at about 1.8 metres — so a 6.8-metre group exceeds it by a factor of about four. The front occupant will step. But each step consumes one contact spacing of the interval and no more, so an interval of 12 metres absorbs some 31 of them before it closes.

The interval therefore does not prevent the edge from being consumed. It determines how slowly.

Arrangement over 100 m	Group length	Load at a group's head	
Two intervals of 10 m	26.7 m	4,219 N	exceeds
Three intervals of 12 m	16.0 m	2,531 N	exceeds
Four intervals of 12 m	10.4 m	1,645 N	exceeds
5 intervals of 12 m	**6.7 m**	**1,055 N**	**within**

Sixty metres of interval buys forty metres of usable run, and it does so with 5 gaps rather than the seventeen partitions that a spacing of 5.6 m would require over the same length. That is a different proposition from a partition every few metres, and it is available without installing anything.

One feature of the arrangement works in its favour. Once the route is occupied in groups, the group ahead stands between the next and the gate line. The common target is no longer visible from within a group, the alignment factor falls from unity to the figure for a single direction, and the permitted group length is correspondingly longer than it would be for a continuous queue with the gate in view. The figures above already reflect this; the measure partly reinforces itself.

And the report stops there

How a venue arranges for its arriving population to occupy a route in separated groups is not addressed here, and the omission is deliberate.

Any specific scheme has an execution rate. Batched admission, holding points, staff-managed release and signalling all achieve the grouping under some conditions and fail to under others, and a scheme carried out at seven parts in ten is not the scheme. A report specifying one would be specifying an outcome it cannot observe and a venue would be complying with a method rather than with a criterion.

What is supplied instead is the dimension. A maximum group length, an interval that determines how long the arrangement survives, and the arithmetic connecting both to the declared inputs. A venue knows its own constraints — its staffing, its ticketing, its transport interface — and is in a position to meet a stated dimension by means this report is not.

The same discipline applies throughout. This framework establishes what a geometry permits and states it as a checkable quantity. It does not prescribe operations, and a criterion that can be checked is worth more than a method that must be trusted.

So the trade

The approach to Station Underpass admits one direction along its whole 100 m. Four of the five exits do not have this property; this one does.

Either the route is occupied in groups no longer than 6.8 m, separated by clear intervals of 12 m — on this route, 5 intervals giving groups of 6.7 m, by whatever means the venue finds workable — **or** 38.0 m of exit width is counted rather than 46.0 m, which reduces the permitted capacity by about 17 per cent, or 5,248 people. At the population sold this is not a reduction in headroom but a shortfall of 5,072 people, and the eight-minute calculation is not satisfied.

A serpentine is not the answer. It bounds the run and lowers the discharge rate, and for a route counted as exit width the second effect defeats the purpose of counting it. The grouping above costs no throughput, because a route occupied in groups discharges at the same average rate as one occupied continuously.

And there is no margin to absorb the alternative. The egress calculation requires 45.7 m and the venue provides 46.0 m — a margin of three tenths of a metre, under one per cent. Setting aside any one exit therefore fails the calculation, and this section has a consequence at this site that it would not have at a site designed with slack. A venue providing fifty-five metres against the same requirement would receive it as an observation rather than as a trade.

What this report does not do is choose. Whether a width whose approach admits one direction may be counted alongside four whose approaches admit several is a question for whoever sets the applicable criterion. This report supplies the distinction, the arithmetic on both sides of it, and the cost of removing it.

A third sensitivity, and this one is ours

The two tests above vary the venue. This one varies nothing about the venue at all: it varies a choice made inside this framework, and it moves the deliverable by nearly a factor of two.

The relations come in two forms. A body in a chain passes on what it receives, and it may also absorb a fixed amount irrespective of how much arrives. Writing δ for that residual, the load borne at depth L is

$$F = Lf \quad (\delta = 0), \qquad \text{or} \qquad F = Lf - (L-1)\,\delta.$$

The second is the more realistic of the two — a standing body does absorb something — and it is also the more permissive, because absorption is what stops a chain growing.

This report uses the first form for every bound it reports. The residual has been measured, but on a single subject and over a range, and a quantity of that provenance is not used to relax a requirement. It is used only where a value is unavoidable: whether a run accumulates *at all* is a comparison rather than a bound, and at $\delta = 0$ that comparison is satisfied trivially and decides nothing.

Residual assumed	Bound on a group	Configuration here	Route occupied
0 N *(adopted)*	6.8 m	5 intervals of 12 m	40 of 100 m
30 N (measured, low end)	13 m	4 intervals of 12 m	52 of 100 m
35 N (measured, mid)	16 m	3 intervals of 12 m	64 of 100 m

The consequence is stated rather than buried. Adopting the conservative form costs this route twenty-four metres of occupiable length against the mid-range measurement, and the queue that no longer fits stands further upstream. That is a real cost, borne because a figure measured once on one person should not be what a guarantee rests on.

And it prices one of the gaps. A properly established value for the residual would relax the requirement here by close to a factor of two, at this venue and at every other. It is a property of bodies rather than of this site, it is measured once, and it is not a cost this venue is being asked to bear. It is the clearest available example of what the framework's stated gaps are worth in metres.

What follows

Compliance responds to	This framework responds to
population	**how many directions a region admits**
exit width	length of a connected run
flow rate	declared exertion
occupancy density	declared scale
permitted period	chain pitch

The two lists do not intersect. Neither criterion is a coarse or fine version of the other; they are expressed in quantities that do not overlap, and a configuration can therefore satisfy one while the other has not been evaluated at all.

This is the finding of the present assessment stated in its most compact form. Rivergate Yard satisfies the applicable criteria. Those criteria are expressed in width, rate, density and time, and none of them records how many directions a region admits. The condition this report identifies is not a failure of compliance; it is a quantity compliance does not contain.

The criteria are not named here, and the finding does not depend on which they are. The figures this report checks against — a discharge rate per metre of effective width, an evacuation period, a permitted time in a place of relative safety — are of the kind common to assembly-venue guidance generally, and a reader working to a different set substitutes their own and repeats the comparison unchanged in form. What is claimed is not that some particular document omits the quantity. It is that a criterion expressed in width, rate, density and time cannot contain it, because the number of directions a region admits is not a function of any of them.

5. A Configuration That Closes the Condition

The matrix of §2 reports a condition. This section sets out what would remove it, and what would remain if it were removed.

Nothing here is a recommendation that Rivergate Yard be altered. It is the answer to the question the assessment raises, stated so that the cost of the answer is visible alongside it.

The governing feature is not the length

It is natural to read the finding as being about a hundred metres. It is not, and reading it that way leads to the wrong remedy.

The relations of Cell 4 apply where contact is continuous and a recovery step is unavailable — that is, where the floor area a person would need to step into no longer exists. Length matters only after that condition holds. A hundred metres of crowd in which people can still step sideways carries no chain at all.

What removes the lateral option here is the walls. The underpass is bounded on both flanks along its whole length, and the ramp likewise. A crowd in it has nowhere to go but forward or back, and back is occupied. The same population standing in an open forecourt retains the option, and the condition does not arise however many of them there are.

This reframes the question the drawing should be asked. Not *how long is the route*, but *where does the queue form, and does that place have edges people can leave through.*

Where a queue forms is a design decision

A queue forms upstream of whatever is controlling its rate. That is not a behavioural observation; it is what a control point does.

At Rivergate Yard the control point is the gate line at the head of the ramp. The queue therefore forms in the route — in the walled space — and the route's length becomes the run length. Moving the control point to the street end of the underpass moves the queue with it:

	As designed	With the control point moved
Gate line	head of ramp	street end of underpass
Queue forms in	Route A (walled)	south forecourt (open)
Route A is	a queueing space	a through-route
Lateral escape in the queue	none	to the street on two sides
Longest connected run	100 m	bounded by the forecourt's free edges

The route is not made shorter. It is made into something people walk through rather than stand in, and a body in transit is not a contributor to a sustained chain.

Whether the forecourt can take it

The south forecourt is 2,550 m^2. At a managed queueing density of 2.0 persons/m^2 it holds some five thousand people, which accommodates a sustained excess of a hundred and sixty per minute for half an hour — the middle row of the scenario table in Cell 7.

And when it fills, the failure mode is the right one. The queue extends into the street, and the street has free edges on both sides. Under the configuration as designed, the queue extends into the underpass, which does not.

Both configurations overflow. The question a design decides is not whether an overflow can occur but where it goes, and that question is answerable from the drawing before anyone is admitted.

What remains

The route can still become a queue. If anything downstream stalls — a fault at the gate line, an obstruction in the yard, a stoppage for any reason — people already in the route stop moving, and it reverts to a queueing space of 100 m with no partition. The guaranteed load in that condition is what Cell 4 reports, unchanged.

So moving the control point does not remove the need for a run-length measure, and it does not make it contingent either. What it changes is how much of the route is standing rather than moving in the ordinary case. The measure itself has to be in place before a stoppage, because it cannot be introduced after one: the state it prevents does not reverse when the measure is resumed (§6).

In practice the separation is therefore carried in the moving stream. A route admitted in groups keeps its gaps whether the stream is moving or stopped, and §4 states how many steps a gap of 12 m absorbs before it closes. Installing fixing points into an already-continuous crowd is not an option at the moment it would be needed.

What the relocation buys is therefore not the removal of the measure but a reduction in what it has to hold: a route people move through carries fewer of them, for less time, than one they wait in.

The resulting configuration

Measure	Acts on	Status
Gate line at the street end of the underpass	where the queue forms	standing
Serpentine in the forecourt, legs $\leq$ 5.6 m	run length in the open area	standing
Groups $\leq$ 6.8 m at intervals $\geq$ 12 m in Route A	run length if the route reverts to a queue	standing
Fixing points in Route A at 5.6 m	the same, by installation	alternative to the above
Forecourt overflow route to the street	where an overflow goes	standing

The grouping is stated in preference to the fixing points because it installs nothing and costs no throughput (§4). Both bound the run; only one of them can be maintained by admission alone.

The serpentine in the forecourt is secondary. An open area with free edges does not reach the condition of Cell 4, so the leg length there is a matter of orderly flow rather than of load. It is specified against the same criterion because doing so costs nothing and because an area's edges can be obstructed by the same temporary structures that obstruct anything else.

What this section does not claim

It does not claim that the resulting configuration is safe. This framework establishes what load a geometry permits and does not establish what load injures; a configuration in which the accumulating mechanism cannot operate is a configuration in which *that* mechanism cannot operate, and nothing follows from it about any other.

It does not claim that the change is the only one available. Relocating a gate line has consequences for ticketing, search, weather protection and the public highway that this report has not assessed and is not competent to assess.

What it does claim is narrow and checkable: under the configuration above, the connected run available to a chain is bounded by an open area's free edges rather than by a walled route's length, and the condition §2 identifies does not arise while the premises hold. Which premises they are is stated in the table; a later reader can determine whether they still do.

6. Where the Safety Rests

The preceding sections establish what the geometry permits and what would change it. This one asks a different question, and it is the one that distinguishes two venues that both satisfy a criterion: *how much has to go right.*

Two kinds of safety, equivalent in verdict and not otherwise

A configuration can fail to produce the mechanism of §2 in two ways.

Geometrically. The region admits several directions of travel, the alignment factor is divided among them, and the per-contributor exertion falls below what is absorbed. Nothing accumulates at any length, for any duration, and under any occupancy the exits permit. No measure is required, so no measure can fail.

Operationally. The region admits one direction and the mechanism can operate, but a measure prevents the run from reaching a length at which the load matters. The configuration is safe while the measure holds.

Both satisfy the same criterion. They differ in that the second has an execution rate and the first does not.

Which elements of this venue rest on which

Element	Directions	Safety rests on	Can a lapse reach it?
Standing area	several	geometry	no
Service Road Gate, River Path Gate, East Emergency Gate, West Emergency Gate	several	geometry	no
South forecourt	several	geometry	no
Station Underpass, approach	**one**	**a measure**	**yes**

Every element of this venue but one is safe by its shape. A steward shortage, a missed briefing, a barrier moved to relieve a pinch point, a schedule that runs late — none of these reaches the mechanism in any of them, because the mechanism does not operate there at all.

The exception is the approach to Station Underpass. Its safety, under the grouping arrangement of §4, is the arrangement.

And a lapse there is not recoverable

It might be supposed that a measure executed imperfectly yields proportionally less of the protection it was taken for — that seven parts in ten of the grouping gives seven parts in ten of the benefit. It does not, and the reason is set out in the framework's treatment of how such a configuration forms.

Forming a compacted region requires work, supplied by advance from behind. Maintaining it requires none: once the space between bodies has been closed it is held by geometry, not by anything continuing to push. A measure that lapses for a period does not merely fail to protect during that period — what forms during it persists after the measure is restored.

So the execution has to be right from the outset rather than corrected on observa-

tion. There is no state of the configuration from which a resumed measure returns it to where it was. This is the sense in which an operationally safe element differs from a geometrically safe one by more than its execution rate.

What this implies for the record

The distinction determines what a venue's documentation has to contain in order to be checkable afterwards.

Where safety is geometric, the drawing is the record. It can be verified at any time, by anyone, against the built form. Nothing about the day of the event enters.

Where safety rests on a measure, the drawing is not enough. What has to be recorded is the measure as a dimension, and whether it was met. A duty expressed as managing the approach appropriately cannot be checked afterwards, because nothing about it is falsifiable. A duty expressed as groups no longer than 6.8 m at intervals no shorter than 12 m can be: it either was or was not.

This is the practical consequence of the discipline of §4, where the report supplies a dimension rather than a method. A dimension can be recorded. A method can only be described.

What sustaining a measure requires, and what it does not

Where safety rests on a measure, something has to establish that the measure is being met. That is a narrower requirement than it is often taken to be, and the narrowness is worth stating because the wider version is unobtainable.

What is required is an observation of the dimension. The arrangement of §4 is stated as a length and an interval. Whether it holds is a question about lengths, answerable by looking, and a record of it is a record of lengths at times.

What is not required, and not presently available, is a rate. The density of a region over time *is* computable: occupancy rises at the difference between arrival and discharge, and both are schedulable quantities. What is not available is the density at which contact becomes continuous, and without it the computed curve cannot be told where it crosses into the state this report is concerned with.

So a density instrument is not useless here; it is uncalibrated. A reading of occupancy is a real measurement of a real quantity, and it locates the region on the path from free movement to continuous contact. What is missing is the mark on that path. A density threshold adopted without it is a number compared against a convention, and the instrument will report faithfully against a figure that has not been derived.

This is worth separating from the finding of §4, which is that the *load* does not respond to occupancy. Both hold. Occupancy is not what the mechanism responds to, and it is nonetheless the quantity that says which phase a region is in — once the mark is placed.

And a rate for the load itself is a separate gap. How fast the borne load rises once contact is continuous depends on the per-contributor exertion, which this framework declares rather than measures (§3). A response time or a permitted inflow rate would require both that figure and the mark above, and this report offers neither.

So the record that matters is short. Whether the groups were of the stated length, whether the intervals were maintained, and at what times. A record of activity — staff deployed, briefings held, barriers positioned — records that something was done. It does not record that the dimension was met, and only the second is what the safety rests on.

Outcome

The foregoing supports a classification, and it is worth stating plainly because the distinction it draws is the one a reader has to act on.

State	What it means	What safety rests on
Closed by geometry	The configuration keeps the load below the adopted level with no measure required.	the built form
Closable by a stated measure	It does not, but a measure of stated dimension does.	execution of that measure
Not closable within scope	No measure within this report's scope brings it below the adopted level.	—

Rivergate Yard is closable by a stated measure. Every element but one is closed by geometry. The approach to Station Underpass is not, and two things close it: lateral relief in its flanks, which is structural and returns the route to the first state; or the grouping arrangement of §4, which leaves it in the second. Adopting neither leaves the approach in the third, and §4 states what that costs in counted exit width.

What this classification is, and is not

It is a statement about a computation. Given the level the reader adopts and the geometry supplied, the load either stays below that level or it does not. Nothing in the classification depends on a view about whether that level is the right one; a reader adopting a different level obtains a different classification from the same arithmetic, and the report has no preference between them.

It is not an authorisation, and it is not a certification. Nothing here permits a configuration to be occupied or withholds permission. Whether a venue may lawfully operate is a question for whoever is empowered to answer it, and this classification is an input to that question in the same way as any other technical finding.

And the third state is a finding, not a refusal to assess. Where no measure within scope closes the gap, that is the assessment's result and it is as informative as the other two — more so, arguably, since it is the one that cannot be resolved by spending money on barriers.

What each missing quantity would close, and where its cost falls

The assessment above leaves several determinations open, and a reader is entitled to know what would close each and who would bear the cost of doing so. The framework states its gaps rather than working around them, and the gaps have a structure worth setting out.

Quantity	Obtained by	Would enable	Cost falls on
Area a recovery step requires	controlled measurement; one-off and general	the density at which a region enters continuous contact	**research, once**
Per-contributor exertion for a scene class	contact-force measurement under controlled loading; one-off per class	replacing the declared value with a measured one	**research, once**
Force at which a standing person must step	controlled measurement; one-off and general	the interval bound on a firmer footing than one subject	**research, once**
Net inflow at a boundary	one counting line	whether accumulation is occurring, in sign	**the venue, per site**

Three of the four are not the venue's to pay for. They are properties of bodies rather than of this site, they are measured once, and once measured they are available to everyone. A venue is not being asked to fund them, and a report that presented them as site costs would be misrepresenting where the work lies.

The fourth is, and it is modest. A counting line at the upstream mouth of the approach is the same instrument already installed at the gate line, placed earlier. It closes three cells of the matrix that are not assessed here (§3), and it is the only measure in this report that does so.

And the choice this presents

The missing quantities are not an argument for waiting. They set out a choice, and it is the choice a venue is actually making whether or not it is stated.

One option is to obtain the information. Instrument the approach, record what it reports, and act on it. This is a real option, it will become a better one as the research quantities above are measured, and there is nothing in this framework that argues against it. Environmental sensing and video analysis are advancing, and a determination that is not observable today may be observable in a few years.

The other is to configure so that the information is not needed. A region in which the mechanism cannot operate requires no reading, no threshold, and no response. It is safe in the sense this framework can state, at any occupancy the exits permit, whether or not anyone is watching.

The two differ in a respect that outlasts the instrumentation. An instrumented approach is safe while the instrument is read and acted on; a configured one is safe because of its shape. The first has an execution rate and the second does not, and §6 sets out why a lapse in the first is not recovered by resuming it.

One caution about undertakings

Where a permission rests on an undertaking to observe something, the undertaking becomes part of what makes the configuration acceptable. The observation is then not an optional refinement but the thing the permission was granted against.

An undertaking of that kind that is not performed leaves the configuration in a worse position than one that was never given, because the alternative measures that might have been required instead were not required — precisely because the undertaking was given.

This report makes no observation about any particular case. It notes the structure because it bears directly on the distinction of §6: an undertaking to observe is an operational measure, it carries an execution rate, and what it displaces may have carried none.

What this section does not do

It does not advise on liability, insurance, or the allocation of responsibility. It states which elements of a configuration are safe by their shape and which are safe by something someone does, and it states that distinction in the same terms as the rest of the report — as a property of the configuration, checkable from the drawing and the operating plan.

What follows from that distinction is for whoever needs it. An insurer pricing a venue, an authority setting a condition, an operator deciding where to place supervision, and a venue comparing two possible layouts are all asking the same question in different words: how much of this rests on someone doing something correctly. This report answers it for one configuration and offers no view on what should be done with the answer.

A configuration's tolerance for operational error is a property of its shape, and it is not recorded by any pass or fail. Two venues meeting the same criterion may differ entirely in how much of their compliance depends on execution. At Rivergate Yard, every element but one is safe by its geometry; the exception is the approach to Station Underpass, whose safety under the grouping arrangement *is* the arrangement — and which, once a lapse has allowed the configuration to form, is not restored by resuming it.

References

Mats Hagberg. Muscular endurance and surface electromyogram in isometric and dynamic exercise. *Journal of Applied Physiology*, 51(1):1–7, 1981. Records that the original determination treated a hold exceeding fifteen minutes as of unlimited duration.

I. H. G. Hopkins, S. J. Pountney, P. Hayes, and M. A. Sheppard. Crowd pressure monitoring. In R. A. Smith and J. F. Dickie, editors, *Engineering for Crowd Safety*, pages 389–398. Elsevier, Amsterdam, 1993.

Mark W. Kroll, G. Keith Still, Tom S. Neuman, Michael A. Graham, and Larry V. Griffin. Acute forces required for fatal compression asphyxia: A biomechanical model and historical comparisons. *Medicine, Science and the Law*, 57(2):61–68, 2017. doi: 10.1177/0025802417695711.

Burkhard Madea, editor. *Asphyxiation, Suffocation, and Neck Pressure Deaths*. CRC Press, Boca Raton, FL, 1st edition, 2020. ISBN 9781498759038. doi: 10.1201/9780429188947.

Walter Rohmert. Ermittlung von Erholungspausen für statische Arbeit des Menschen. *Internationale Zeitschrift für angewandte Physiologie einschließlich Arbeitsphysiologie*, 18:123–164, 1960. doi: 10.1007/BF00698869.

Gisela Sjøgaard, Guy Savard, and Carsten Juel. Muscle blood flow during isometric activity and its relation to muscle fatigue. *European Journal of Applied Physiology and Occupational Physiology*, 57(3):327–335, 1988. doi: 10.1007/BF00635992.

Yixuan Tao. A physical quantification framework for crowd compression risk: The p-index system and its contact-mechanical foundation. Zenodo, 2026. URL `https://doi.org/10.5281/zenodo.21261698`. Preprint. Concept DOI (all versions).

7. Applying This to a Real Venue

Rivergate Yard is fictional. This section states what the same assessment looks like when the venue is not, what has to be supplied for it, and what each stage costs.

What changes when the venue is real

The method does not change; the inputs do. Every figure in this report was derived from a geometry and a declared set of framework constants. For a real venue the constants are the same and the geometry is yours, which means the determination may differ — a route that admits several directions of travel produces no finding at all, and that outcome is as common as the one demonstrated here.

What this report cannot tell you is which of those your venue is. That is the point of running it on your own drawing rather than reading it on someone else's.

What has to be supplied

The built-form row is assessable from drawings alone. What is needed is the occupiable extent of each approach and holding area, the positions and effective widths of the entry and exit lines, and the position of anything that narrows or divides a route — including temporary fencing, which on the evidence of the retrospective literature is frequently the element that matters and frequently the element not shown on the plan.

The analysis is performed on what is supplied. Where the material is insufficient to determine a quantity, the report states that the quantity was not determined and lists what would be needed to determine it. It does not substitute an assumption and carry on. This is stated here rather than discovered later, because a determination based on an assumed geometry is worth less than an honest gap.

The stages

Stage	What it covers	Fee (USD)
This report	The built-form row, on a fictional venue.	—
Full matrix demonstration	All nine cells, same fictional venue. Establishes what the complete assessment contains, so that you can judge what to supply.	1,000
Access geometry assessment	The built-form row, on your venue. Delivered with a quotation for the full assessment and a list of any further material required.	2,000
Full assessment	All nine cells, one approach and one event.	quoted
Venue programme	All approaches, assessed by zone.	quoted
Simulation	Agent-level reconstruction where the geometry warrants it.	quoted

The first three are fixed because they can be priced without seeing your site. Two of them are documents that already exist, and the third is read from drawings, so the work is known before it begins.

The rest are quoted, and the quotation is delivered with the access geometry assessment. What a full assessment costs depends on how many approaches a venue has, how they interact, and what condition the supplied material is in — none of which can be judged honestly at a distance. The third stage is where that judgement is made, and its fee covers making it.

The demonstration fee exists for a reason worth stating. The assessment is performed on material you supply, and what to supply is not obvious until you have seen what the complete matrix asks. The fee buys the document that makes that judgement possible, and 500 of it — half — is credited against any later stage commissioned by the same client, for any venue. The other half is the price of the document itself.

Neither the credit nor a quotation carries an expiry. A quotation issued for a venue holds for that venue on unchanged conditions, whenever you return to it. There is no advantage to us in a decision made quickly, and a deadline would only produce a hurried refusal. If the framework is revised substantially enough to change what a quotation was based on, the quotation is reissued and the difference is met by us, not by you.

Neither the demonstration nor the assessment is refundable. What each buys is a document, and a document that has been read cannot be returned. The demonstration exists precisely so that what an assessment contains is known before it is commissioned rather than after — which is why a refund is neither offered nor needed: nothing about either deliverable is discovered too late to act on it. This is stated here and again at the point of payment.

If the determination is the third state

§6 sets out three states, and the third — *not closable within scope* — is the one a venue is right to be wary of before commissioning anything. It produces a document that says something uncomfortable and cannot be unsaid.

There is a priced route out of it, and it is stated before you commission rather than after. Where a configuration is not closable by a measure within this report's scope, what remains is engineering: relief in the flanks of a run, a relocated control point, a changed boundary. Analysis of what would move a configuration from the third state to the first is quoted on the same basis as the assessment itself, and the quotation is given with the finding rather than in a later conversation.

What is not available, at any price, is the withdrawal of a determination already made. A report that could be bought back would be worth nothing to the people you would want to show it to, and that includes your insurer and your regulator. The value of the document to you depends on its not being that kind of document.

What this service does not do

It does not allocate responsibility. This framework is not qualified to, and no report produced under it should be used to. The distinction between what a geometry permits and who ought to have done what is set out in §1 and it is not a formality.

It does not give legal or insurance advice. What follows from a determination — for a licence application, a policy, or a contract — is for advisers qualified in those fields.

The simulation engine is not licensed with the reports. Where a case warrants agent-level reconstruction, the interface is the output data rather than the model. This is stated at the outset so that it does not surface unexpectedly at the point it matters.

Contact

Yixuan Tao — PneumaTheorem, Inc.
kevin@pneumatheorem.org
https://www.pneumatheorem.org

A drawing and a description of the event, sent to the address above, are enough to begin. If the approach in question admits several directions of travel, we will say so and there will be nothing further to bill for.