

# The Floor That Peaks at Midnight: Hour-of-Day Demand Concentration in a Large-Format Transit-Adjacent Retail Location

Indivd Analytics — Research

September 17, 2026 · Version 1.1 (Public)

## Abstract

Conventional retail operations assume a daytime demand curve. This paper describes a location where that assumption fails. The location is one large-format, transit-adjacent retail site whose footfall is recorded continuously by entrance-line counters; it was selected for this study precisely because its traffic pattern departs furthest from the conventional curve, and the paper should be read as a rigorous description of that case, not a survey of the whole estate. Across the 27 complete local days in August 2026 (canonical per-hour totals — the standard, audited visitation metric — summed over 17 functional zones), the six busiest clock hours carry 41.7% of all daily traffic, and the single busiest clock hour is 00:00 local time — not a trading hour — with a window mean of approximately 4,516 location-visits per hour, nearly  $1.9\times$  the location’s overall hourly mean. The conventional daytime block (10:00–17:00) contributes only 29.2% of the day’s movement: scarcely more than the pre-dawn and late-night cluster alone (00:00, 05:00, 06:00 and 23:00 jointly carry 28.2%). The implication for this format is structural: staffing, restock and queue capacity built around a standard retail curve would under-serve the actual demand window. All figures are observed, aggregated at format level and anonymized; the single-location scope is stated honestly as the paper’s principal limitation.

## 1 Problem and motivation

Most retail scheduling folklore treats 10:00–18:00 as “the day”: it is the window that staffing rosters, replenishment cycles and marketing calendars are built around. A natural empirical question follows: *when does demand actually arrive* at a retail format that does not live on a high-street clock — specifically a large-format location whose traffic is anchored to a flight schedule rather than to shop hours?

The question matters for three audiences. For operators, the peak-to-daytime gap is a direct staffing and queueing decision. For users of the analytics platform, it demonstrates that the canonical hourly visitation signal is informative enough to re-order a business’s mental model of its own day. And for anyone studying retail demand curves, it is a testable, fully anonymized description of a curve that looks nothing like the textbook one.

## 2 Data

The analysis uses the canonical per-hour visitation total (`location_visits`) recorded by entrance-line counting at one large-format transit-adjacent retail location, summed across 17 functional zones [1]. Four recording and selection conventions matter for honesty:

- **How the location was selected.** This location was chosen for study *because* its hourly profile departs furthest from the conventional daytime curve (its single busiest clock hour

Table 1: The six busiest clock hours (local time), by window-mean hourly volume and share of daily traffic, computed from the canonical hourly series.

| Clock hour  | Mean location-visits/hr | Share of daily traffic |
|-------------|-------------------------|------------------------|
| 00:00–01:00 | 4515.7                  | 7.8%                   |
| 06:00–07:00 | 4448.1                  | 7.7%                   |
| 23:00–00:00 | 4002.1                  | 6.9%                   |
| 17:00–18:00 | 3983.8                  | 6.9%                   |
| 18:00–19:00 | 3830.3                  | 6.6%                   |
| 05:00–06:00 | 3416.4                  | 5.9%                   |

is 00:00). The paper is therefore a case study of a format where the standard assumption breaks, not an unbiased sample of the estate — read the headline numbers as describing this location, not retail in general.

- **Canonical metric.** Figures are computed on the canonical per-hour total; earlier figures published elsewhere used a zone-line sum and are about  $2.19\times$  larger [1].
- **Window and completeness.** The frozen window is August 2026 (local days). Two calendar days are partial (one with 21 of 24 hours, one with 3); both are excluded, leaving 27 complete days, August 2–28. Including partial days would bias the overnight hours, because the truncated days are missing exactly the late-night hours this paper is about.
- **Anonymity.** Data are aggregated to hourly totals before analysis; no individual, store, brand or regional information is used, and the location is described only at format level.

### 3 Method

The method is deliberately simple so that every number can be re-derived from the raw hourly series:

1. **Hour-of-day profile.** For each clock hour  $h \in \{0, \dots, 23\}$  (local time), total visits across the window:  $S_h = \sum_d v_{d,h}$ , and the window mean of the clock hour  $\mu_h = S_h/27$ , where  $v_{d,h}$  is the canonical hourly total.
2. **Top- $k$  concentration.** Rank the clock hours by their window totals; the top- $k$  share is  $C_k = (\sum_{h \in \text{top } k} S_h) / \sum_h S_h$ . We report  $k = 6$ .
3. **Daytime block share.** The daytime block is fixed at clock hours 10:00–17:00 — eight clock hours; the folkloric “retail day” invoked above runs an hour longer, and every daytime share quoted here is for this eight-hour block. Its share is  $\sum_{h=10}^{17} S_h / \sum_h S_h$ .
4. **Peak-hour level.** The busiest hour’s level is quoted as the window mean of that clock hour,  $\mu_h$  — the mean of the hour, not a single day’s maximum.

Every figure in this paper — and in Figure 1 and Table 1 — is recomputed from the raw hourly series by a single deterministic analysis that accompanies the paper.

### 4 Results

Three findings describe the curve.

**The peak is at midnight.** The single busiest clock hour is 00:00 local time, with a window mean of approximately 4,516 location-visits per hour — nearly  $1.9\times$  the location’s overall hourly mean of roughly 2,418 (Table 1). The runner-up hours are early morning (06:00) and the late-evening cluster (23:00, 17:00, 18:00, 05:00).




Figure 1: Mean hourly location-visits by local clock hour over the 27 complete days of the window. Purple bars are the six busiest clock hours (Table 1); the shaded band marks the conventional daytime block 10:00–17:00. The busiest clock hour is 00:00, outside every trading-hour assumption.

**Six hours carry two-fifths of the day.** The top six clock hours jointly carry 41.7% of all daily traffic. Four of the six (00:00, 23:00, 05:00, 06:00) fall outside the conventional retail day entirely.

**The daytime block is a minority of demand.** The 10:00–17:00 block — the hours most retail schedules treat as core — carries 29.2% of the day’s movement. The four pre-dawn and late-night hours 00:00, 05:00, 06:00 and 23:00 jointly carry 28.2%, nearly identical to the daytime block’s share.

## 5 Validation and quality

All reported numbers are recomputed from the raw hourly series by a single deterministic analysis that accompanies the paper. Two quality rules were applied deliberately: (i) partial days are excluded rather than imputed, because the missing hours of the two partial days concentrate exactly in the overnight window under study; (ii) the peak-hour level is quoted as the mean of the clock hour across the window, not the maximum single hour, so that the headline number cannot be an outlier event. What the analysis does not do is characterize day-to-day variability around those means (for example a dispersion band on the 00:00 hour); the headline figures are means over one month at one location and should be read as such.

### Where the evidence stands

This paper is deliberately explicit about what is and is not anchored in the wider literature. There is no strong peer-reviewed baseline for “which clock hour peaks” in transit-adjacent duty-free-style formats — no comparable published dataset exists, as far as we can establish — so the concentration figures (the 41.7% top-six share, the 00:00 peak) rest entirely on our own measurement of one location, and we state that openly rather than dressing them up as benchmarks. What the literature does support is the mechanism: scheduled air service is organized in waves that concentrate passenger traffic into specific parts of the day [3] — an observation whose roots go back decades in the scheduling literature [2] — and airport commercial areas live off that timetable-driven flow, with passenger segments whose dwell time and spending track their flight times [4]. Scheduled air service also shows strong calendar rhythms, documented at monthly granularity worldwide

[5]. The reader gets both halves: what the literature establishes (mechanism), and what only this dataset can establish (the measured midnight peak and its size).

## 6 Discussion

The finding is best read as an operating question rather than an instruction: for a format of this kind, the demand day starts where the retail day ends. Rosters, restock windows and queue capacity calibrated to 10:00–18:00 would under-serve 00:00, 05:00, 06:00 and 23:00 — which together carry as much traffic as the entire conventional daytime block — while daytime service levels calibrated to the peak would over-provision. Whether a specific location should re-phase its operations depends on its own measured curve, which is exactly what the underlying hourly data provides.

### Limitations

The scope is one large-format transit-adjacent location over one month — selected because it departs furthest from the conventional curve — so the shape, not the absolute level, is the transferable finding, and the figures describe this case rather than the estate in general. A one-month window cannot separate stable schedule rhythm from season or timetable changes; the rank ordering of runner-up hours can shift, and the paper reports the top-six *group* rather than a rigid rank list. Day-to-day variability around the reported means is not characterized, though as a basic robustness description (not a significance test) excluding any one day changes neither the peak hour nor the top-six membership. All figures are descriptive; no significance testing is implied; data are aggregated and anonymized at format level.

## References

- [1] Indivd Analytics Research. Preprocessing Standards for Entrance-Line Demand Data. Papers page, version 1.0, 2026. Methodology for the canonical `location_visits` metric, entrance-zone aggregation and the capture-window gates underlying the hourly series used here.
- [2] Schoennauer, A. W. Airline schedules and their impact upon passenger traffic. *Transportation Research*, 3(1), 1969. A historical pointer to the long-standing observation that scheduled airline timetables — not shop hours — dictate when passenger traffic arrives at airport facilities; the modern wave-structure evidence is [3].
- [3] Pyrgiotis, N., Malone, K. M., and Odoni, A. Modelling delay propagation within an airport network. *Transportation Research Part C: Emerging Technologies*, 27:60–75, 2013. Models hub operation as a sequence of arrival and departure *waves* (banks); establishes the wave-structured, schedule-driven nature of hub-airport traffic that this paper’s mechanism rests on.
- [4] Freathy, P. and O’Connell, F. J. Spending time, spending money: passenger segmentation in an international airport. *The International Review of Retail, Distribution and Consumer Research*, 22(4):397–416, 2012. Documents how airport passengers’ dwell time and spending behaviour follow their flight times — the airport-commercial-demand-follows-the-timetable evidence.
- [5] Dobruszkes, F., Decroly, J.-M., and Suau-Sánchez, P. The monthly rhythms of aviation: A global analysis of passenger air service seasonality. *Transportation Research Interdisciplinary Perspectives*, 14:100582, 2022. Worldwide evidence that scheduled air service carries strong calendar rhythms (documented at monthly granularity), the demand driver of transit-adjacent retail.