DUBLIN AIRPORT OPEX EFFICIENCY STUDY: 2027–2031

Annex B: Airport security benchmarking

Evidence summary

H008 D001 V1.0

June 2026

Executive summary – Dublin Airport (i)

This Annex provides supporting evidence for the Security element of CEPA's Dublin Airport Opex Efficiency Study: 2027–2031, commissioned by the Irish Aviation Authority (IAA).

It summarises supporting benchmarking analysis of airport passenger and hand baggage security. The study assessed four key performance areas – planning, throughput, productivity, and queues – using operational data for a busy summer month for 12 airport terminals, thus providing a snapshot of relative performance at that time[1]. Not all terminals were able to provide all data required, so complete comparisons could not be produced for every performance indicator. Data was provided by airport operators under strict confidentiality and anonymity conditions, so the terminals are not identified in the report but are labelled as A through to L. All the terminals operate under similar aviation security regimes so are broadly comparable. Dublin Airport Terminal 1 is labelled B and Dublin Airport Terminal 2 is labelled K.

Planning: Compared to the benchmarking sample, Dublin Airport's passenger forecasting accuracy ranked sixth (T2) and seventh (T1) out a sample of nine terminals for which data was available. T1 showed a small bias towards over-forecasting, that is the forecast systematically over-stated the actual demand. T2 showed a comparatively large under-forecasting bias, with actual demand systematically exceeding forecast demand. Estimation of images per passenger (IPP) is important in translating the passenger forecast into demand for hand baggage scanning. IPP can be highly variable and hence challenging to estimate. Although not strictly comparable due to some IPP data pre-processing, Dublin Airport's IPP for both terminals appears generally less variable, day-to-day, across the sample month than that experienced by the other terminals. By contrast, the intra-day, hour-to-hour IPP volatility was higher at both Dublin Airport terminals than at any other comparators.

In summary:

i. Dublin Airport's passenger forecasting was less accurate than most of the benchmarks

ii. Its IPP variability is likely more challenging to deal with than for the benchmarks

IPP, together with the throughput of the passenger- and image-screening processes, determines how well capacity is balanced between the body-scanner and CT-scanner components of the lane. The analysis indicates that at present in T1 the body scanner is likely the single binding capacity constraint, whereas in T2 capacity appears more evenly distributed across both components. Across the wider benchmark set, body scanners were also the predominant constraint. This condition is not fixed and can be mitigated through body-scanner process improvements and/or by increasing the number of body scanners to adjust the body-scanner-to-CT ratio.

[1] This summary highlights the results and draws conclusions from Dublin Airport's perspective.

Executive summary – Dublin Airport (ii)

Throughput: The analysis focused on body scanner and CT scanner throughput although some terminals, including Dublin T1, were operating a mixture of CT scanners, body scanners, conventional two-dimensional X-ray machines and archway metal detectors (AMDs) for the period which they provided data. Our throughput analysis considered both peak-hour and all-hour throughput performance for both body scanners and CT scanners:

- **Peak-hour:** For T1, peak hour body scanner throughput was the lowest and most variable (least stable) in the benchmarking sample, which comprised 11 out of the 12 participating terminals, indicating relatively poor performance. However, T1's CT scanner performance ranked as third highest in terms of throughput and joint third in terms of stability, i.e. relatively good performance within the sample. This imbalance between CT and body scanner performance may be part of the explanation for body scanners likely being the main capacity constraint in T1. T2 ranked fourth out of 11 for peak hour body scanner throughput and eighth in terms of consistency: this is relatively good performance compared to the sample. Similarly, T2's peak hour CT scanner performance ranked fifth for throughput and second for consistency, again, indicative of good performance.
- **All-hours:** As for peak hour body scanner performance, T1's all-hours body scanner performance ranks lowest in the sample for both throughput and consistency. Relative to the sample, T2's all-hours body scanner performance degraded more compared to its peak hour performance: it ranks second lowest for both throughput and stability. Our Dublin Airport dataset did not include data on all-hours CT scanner performance to compare with other airports.

Productivity: There was only limited staffing data available from the benchmarking participants so the productivity benchmarking sample is very small: only three other terminals with complete data along with Dublin's two. T2 has the lowest staff-per-lane ratio in the sample and T1 scores as middle of the five comparators. These results are likely driven by infrastructure requirements, with the other terminals having more body scanners to each CT scanner than Dublin and hence, likely needing more staff. T1's actual staffing per lane is very near to the plan. T2's actual staffing per lane is, on average, around 8% lower than the plan over the sample period.

Using passengers per hour per staff hour deployed on lanes as a productivity indicator, T1 and T2 have the joint highest planned productivity across the small sample. However, achieved productivity has a 30% shortfall from plan for T1 and a 16% shortfall for T2. For this realised productivity, T2 ranks highest and T1 ranks third in the sample.

Queues: Comparisons should be interpreted with caution because of the different ways queues are measured and reported at different airports. In the sample period, Dublin Airport met its queue target, although indicatively, its reported average queue times were higher than the comparator sample, with T1 higher than T2. Maximum queues at two other terminals were higher than Dublin's maximum queues.

Contents

1. Introduction
2. Planning
3. Throughput
4. Productivity
5. Queues

Introduction

This annex summarises benchmarking evidence used to inform the Security assessment in CEPA's Dublin Airport Opex Efficiency Study

The work was commissioned by the IAA to inform the regulatory price determination of Dublin Airport. Passenger and hand baggage security forms a large proportion of Dublin Airport's operating cost base.

We use operational data from 12 airport terminals, all operating central search facilities, to calculate and compare performance indicators spanning the planning, throughput, productivity and queue dimensions of the passenger and hand baggage security process.

Performance indicators

Planning	Throughput	Productivity	Queues

• Variance between forecast and actual passengers
  • Accuracy
  • Bias
• Images per passenger (IPP)
  • Variability
  • Intra-day volatility
• Capacity balance between passengers and hand baggage
	• Peak hour throughput
	  • CT scanner
	  • Body scanner
	• All-hours throughput
	  • CT scanner
	  • Body scanner
		• Staff per lane
		• Passengers per lane staff hour
			• Passenger waiting time distributions

Our analysis focuses on the performance of new C3 technology operations, i.e. the use of computed tomography (CT) machines for hand baggage and mm-wave body scanners for passengers. However, some of the terminals in our comparator set were still in a mixed operation with some conventional x-ray and archway metal detectors (AMD) systems when the data was collected. In this report, we do not assess the performance of these legacy systems. The data analysed covers a peak summer month although dates vary from terminal-to-terminal

All terminals operate under similar aviation security regulatory regimes and either have regulator-imposed or self-imposed queue targets. We do not investigate adherence to aviation security screening regulations, instead assuming compliance for each of the terminals.

Not all terminals were able to provide all data required, so complete comparisons could not be produced for every performance indicator.

Data has been provided under strict confidentiality and anonymity conditions, so individual terminals are not identified in this report and instead labelled as A to L.

Planning

We first look at forecast variance, i.e. how well the number of passengers presenting at security in a 15-min interval matches what was forecast

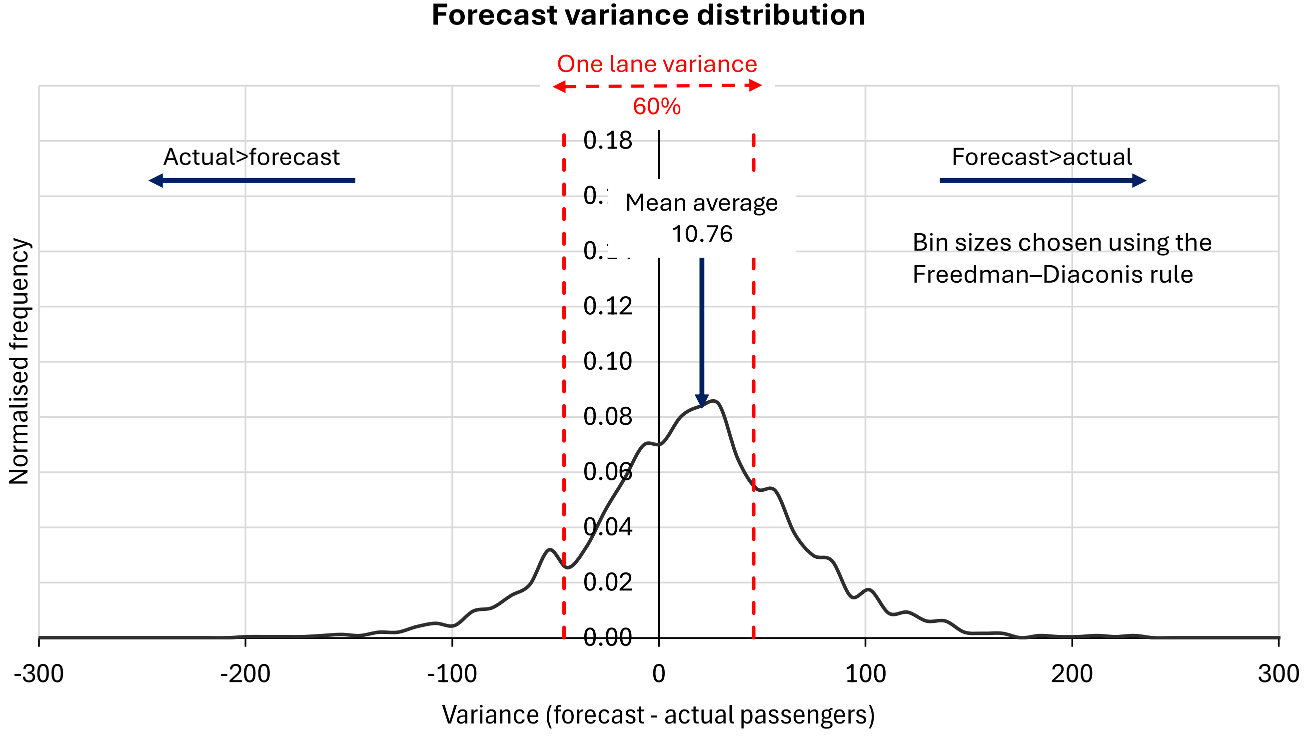

The example shows that:

- The forecast is within a one-lane threshold of the actual demand for 60% of the 15-minute time slots over the busy summer month
- The mean average of 10.76 shows that there is a tendency to over-forecast with the risk that more lanes are planned than necessary

For each terminal, we look at the distributions of forecast to actual passengers, in 15-minute intervals, over a busy summer month,

The width of the one-lane variance is derived from the 95th-percentile peak-hour unit body-scanner throughput, adjusted for the ratio of body scanners to CT scanners so that performance is normalised to a single lane. This reflects a practical definition of a ‘lane’ as a CT unit supported by one or more body scanners. The proportion of the forecast–actual variance that falls within this one-lane band represents the proportion of time the planned number of lanes is within ±one lane of the required lanes, based on the assumption that a lane can run at peak capacity for short periods of time. This tolerance band is a balance between efficiency (actual less than forecast) and queue risk (actual greater than forecast) that arises naturally from fluctuations in the passenger demand profile. The higher the proportion, the better the forecast represents reality

The mean average gives a measure of the degree to which there is a systematic bias in the forecast rather than the expected random distribution:

- A value near zero implies no systematic bias
- A positive value implies over-forecasting with a risk of inefficient operations
- A negative value implies under-forecasting with a risk of high queues

Accurate passenger demand forecasts are the basis of good security planning and underpin (but do not guarantee) subsequent performance

We measure forecast accuracy using the statistical distribution of forecast variance (i.e. forecast – actual) over the busy summer month sample period

Relative forecasting performance of terminals within the sample for which forecast data was available

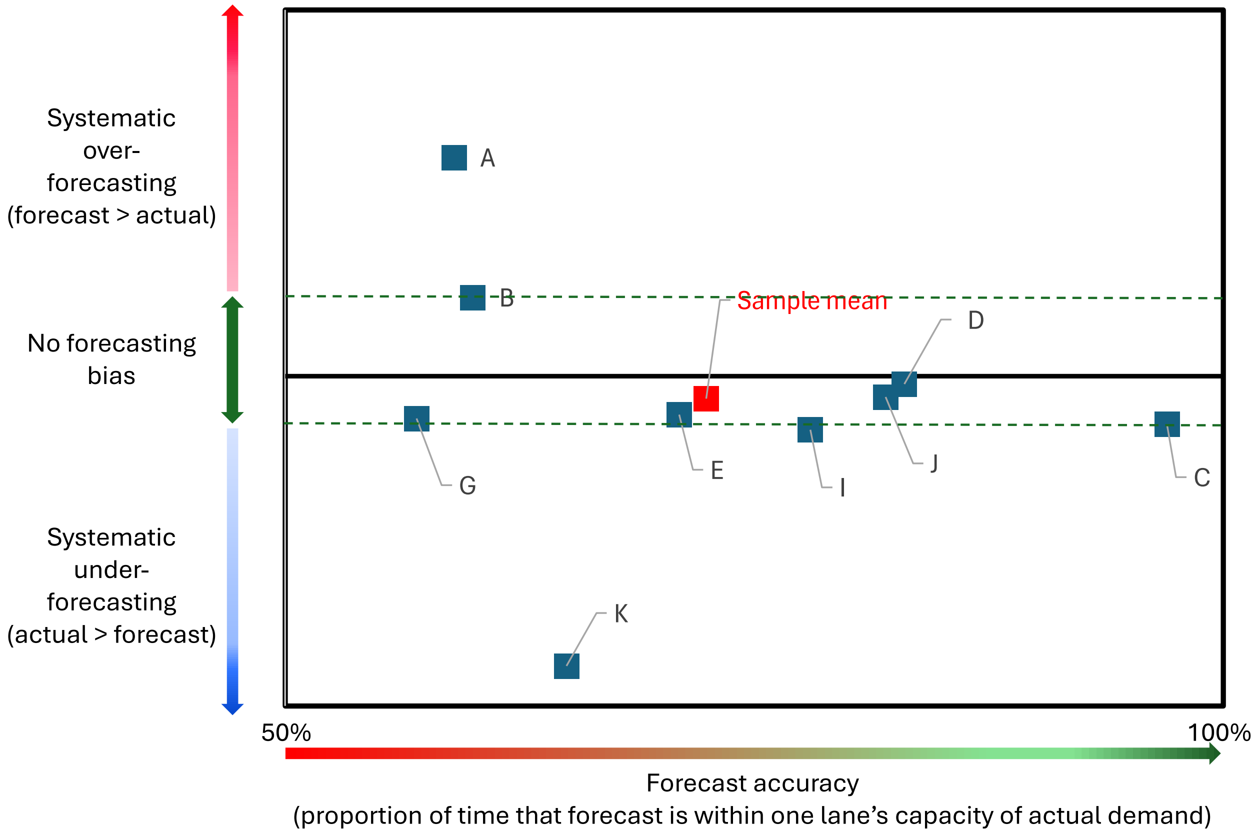

For the terminals in the sample where forecast data is available, the chart compares performance along two dimensions:

- Systematic bias, whether the forecast consistently over- or under-predicts passenger demand. Good performers are sited near the centre of the vertical axis
  - Over-forecasting indicates a tendency to open more lanes than required, potentially reducing efficiency
  - Under-forecasting indicates a tendency to open too few lanes, increasing the risk of long queues
- Forecast accuracy, measured as the proportion of the forecast–actual variance that falls within a ±one-lane capacity threshold. This reflects how well the forecast aligns with the operational tolerance for each terminal. The ±one-lane threshold represents a balance between efficiency and queue risk. Good performance lies to the right on the horizontal axis
- In the sample, C is the best performer; A, B, G and K are the worst performers

Although forecasting is inherently difficult, driven, for example, by changing passenger mix & preferences, the good performance exhibited by some terminals indicates that it should be possible for the poor performers to improve

We now consider variability and volatility in the number of images generated per passenger, IPP, which again affects the productivity of a security operation

Variability defined by the width of the IPP distribution

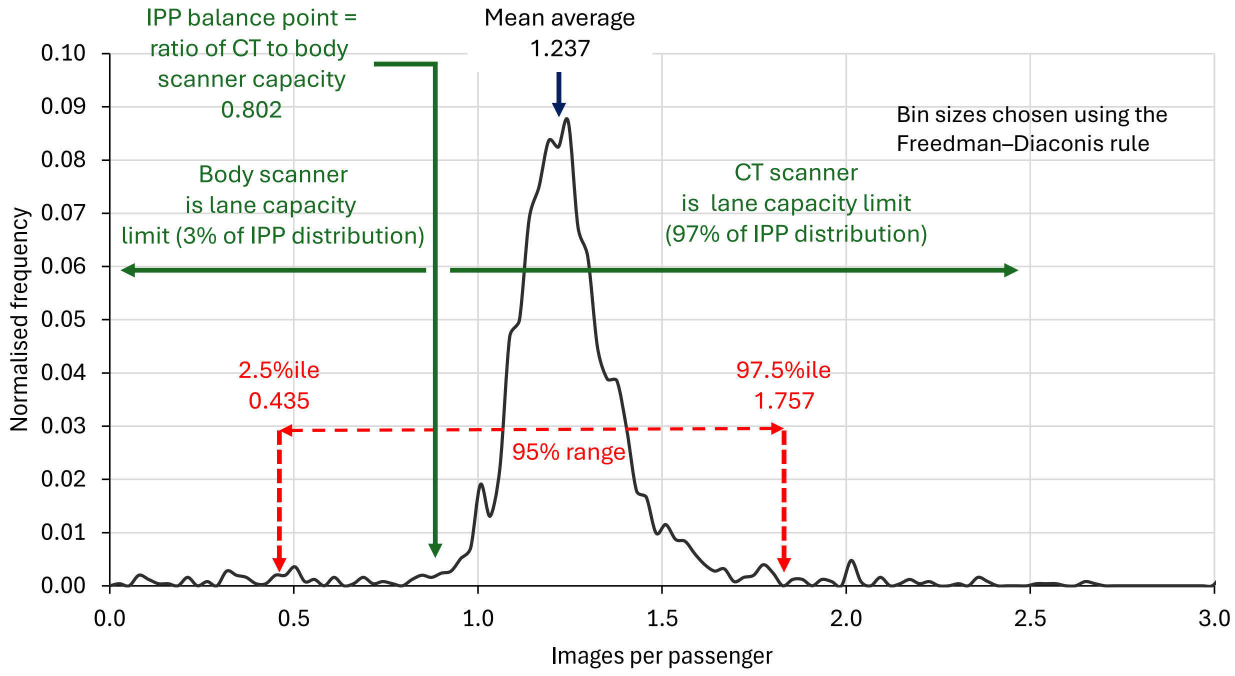

Intra-day volatility of IPP across the day

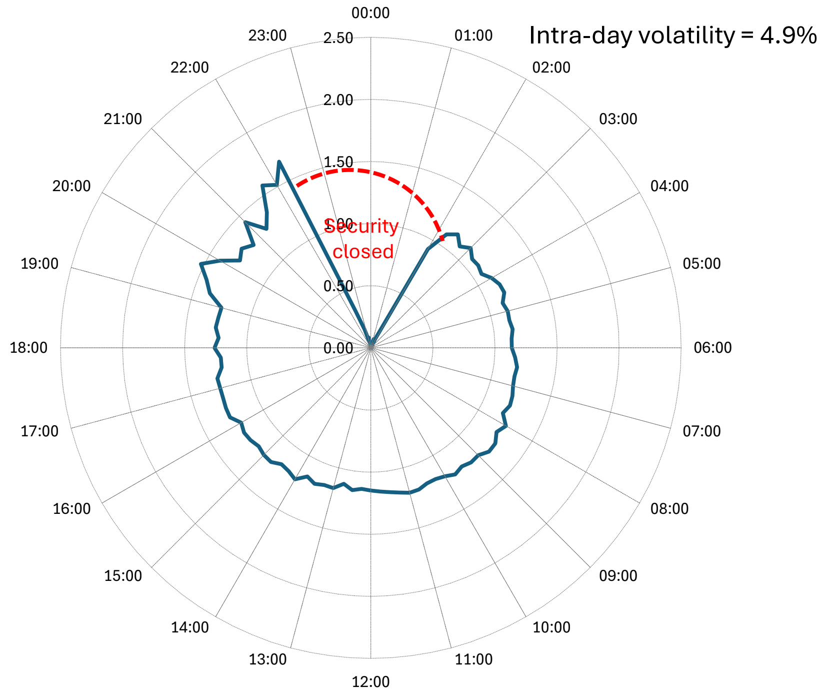

Variability is a measure of uncertainty in IPP based on the distribution of measured IPP for each 15-minute interval. As shown in the chart above, our analysis quantifies:

- The mean average and 95th percentile range of the IPP distribution.
- The IPP value where body scanner and CT scanner capacity are in balance. This balance occurs at the IPP value where the body scanner (passenger) throughput per lane matches the CT scanner (image) throughput per lane, adjusted for the ratio of body scanners to CT scanners in the search area, i.e. there may be more than one body scanner per CT scanner

The **intra-day volatility** is an indicator of the hour-to-hour variation of the IPP from across the day. It is measured as the mean of the absolute changes from one 15-minute interval to the next relative to the overall mean. This is expressed as a percentage, where a low value indicates stability and a high number indicates volatility.

Hades Consulting Limited

We observe wide IPP ranges, presenting a challenge when converting a passenger forecast into demand for hand baggage screening and security lanes

IPP is a measure that is difficult for Dublin Airport to control. It can show large natural variation: the larger the variation, the more challenge and uncertainty there is in planning a security operation

Mean average and 95th percentile ranges for IPP

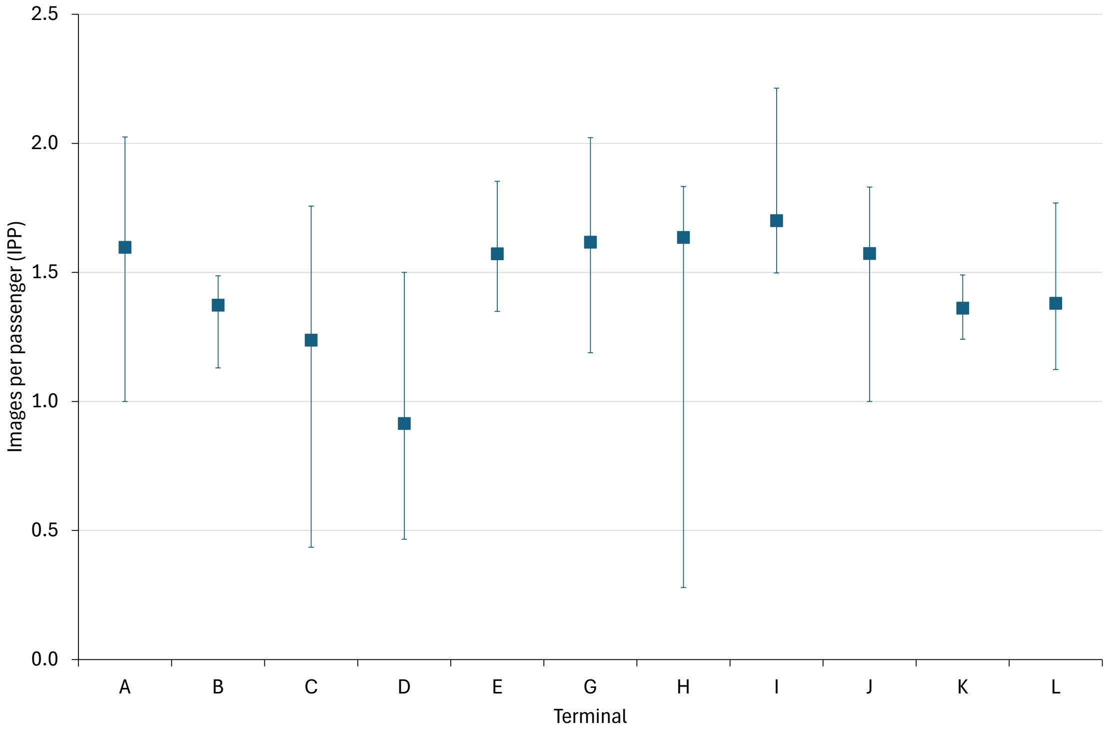

Note: the 95% ranges for terminals B and K are not representative of (shorter than) the true ranges because the data provided was pre-smoothed

For the terminals in the sample where IPP can be calculated, the chart shows two elements:

- The mean IPP across the busy summer period
- The 95th percentile range of the IPP distribution, shown as error bars. Because the IPP distributions are often asymmetric, the mean does not necessarily sit at the midpoint of the 95% bounds

When the lower bound extends further below the mean than the upper bound extends above it, this indicates a left-skewed distribution. Without adjustments to the forecasting method, this pattern suggests that the actual IPP is more likely to fall below the forecast, meaning the forecast tends to overestimate IPP. Terminal H is an example of this behaviour

Conversely, when the upper bound extends further above the mean than the lower bound extends above it, this indicates a right-skewed distribution. Without adjustments to the forecasting method, this pattern suggests that the actual IPP is more likely to be above the forecast, meaning the forecast tends to underestimate IPP. Terminal I is an example of this

We have limited data to compare actual and planned IPP, and so are unable to comment on IPP forecast accuracy.

High intra-day IPP volatility may indicate that lanes need to be adjusted frequently as image demand varies rapidly even if passenger volumes remain stable

This volatility has a greater impact when CT scanner capacity is the dominant constraint – small fluctuations in IPP may then require additional or fewer lanes increasing the risk of high queues and/or under-utilisation

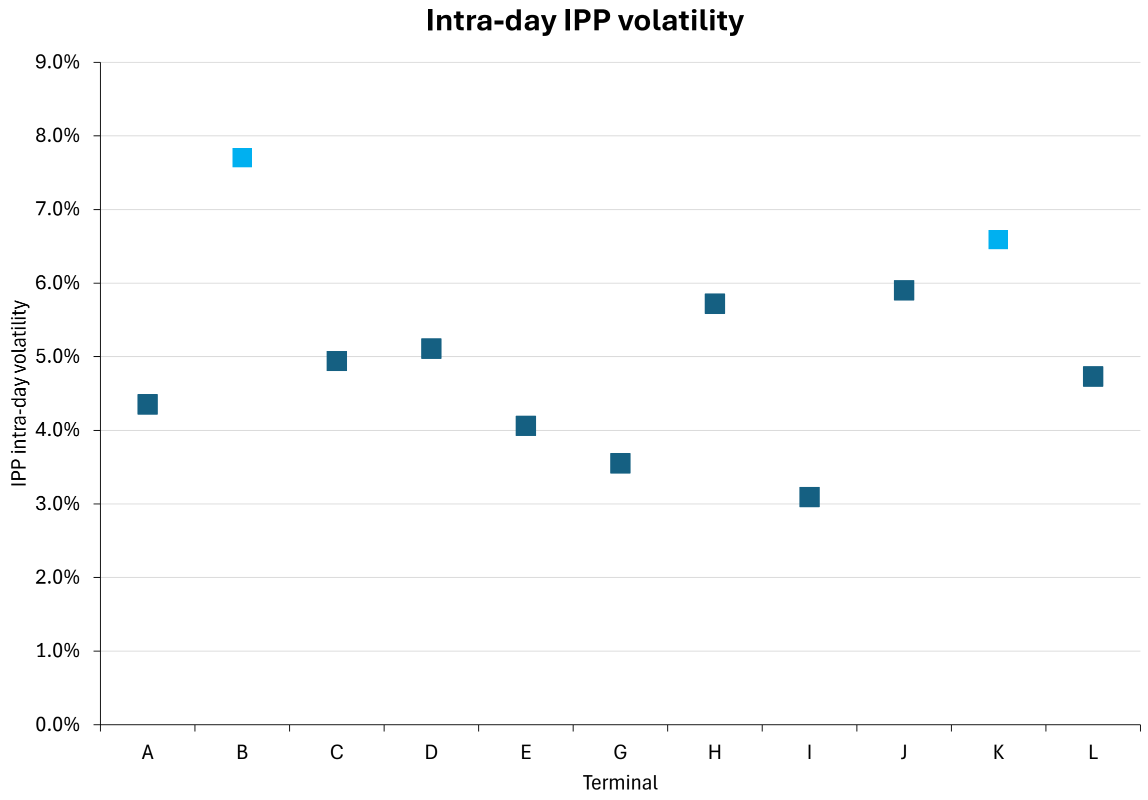

Note: the volatility for terminals B and K is an approximation and so, not directly comparable to the other terminals because the data provided was pre-smoothed

The intra-day IPP volatility for all the terminals in the sample is low – ranging from 3% to 8% of the daily average. Even at peak passenger demand (say 500 passengers in a 15-minute period) a 5% change in IPP would only result in a change in image processing demand of ±3% (or around 25 images in the 15-minute period).

It appears, therefore, that the queue and efficiency risks associated with intra-day IPP volatility are low and are manageable within current operations. And while these risks would be higher if the volatility compounded over several 15-minute periods, they would likely be ameliorated given the greater time to react to the change (if trends in IPP are monitored in real-time).

IPP for most terminals appears to be relatively stable although some terminals show a high level of instability at both operational and planning levels

IPP variability matrix

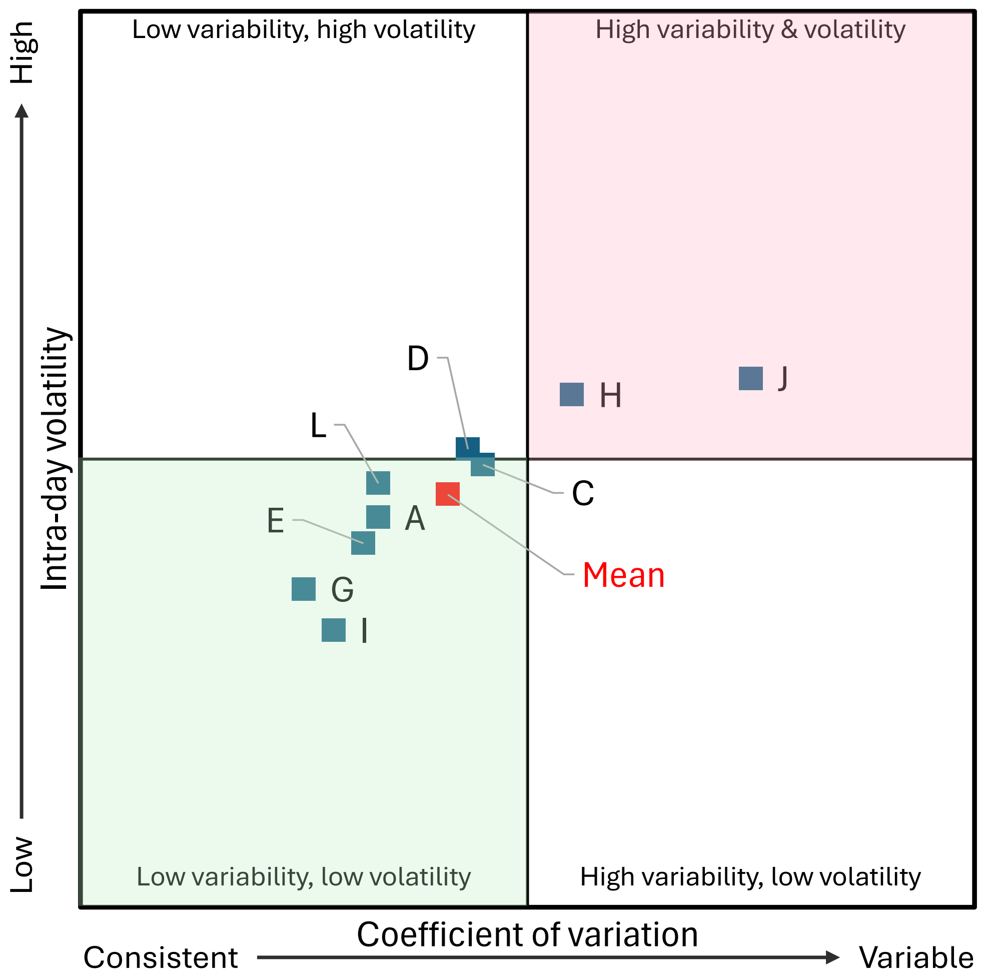

The coefficient of variation (standard deviation relative to the mean) measures the variability of the IPP over the measurement period. It provides an indicator of planning difficulty: higher variability implies greater uncertainty and is therefore harder to plan for than lower variability

Intra-day volatility measures the hour-to-hour variation in IPP. High volatility indicates potentially large short-term swings. Operationally, high intra-day volatility indicates a greater management challenge because the system must react to rapid changes over short timescales

The chart shows that most terminals operate with low IPP variability and volatility. The exceptions to this are terminals H and J and, to a lesser extent, terminals D and C – where planning and operational management is more challenging.

IPP variability means that the limiting factor can quickly switch between image and passenger capacity, depending on achieved throughput and lane configuration

For each terminal, the charts show the IPP threshold when the capacity constraint switches between CT scanner and body scanner, and the proportion of time in the sample month that each constraint applied. This is based on the ratio of CT to body scanners, achieved peak throughput and IPP distributions.

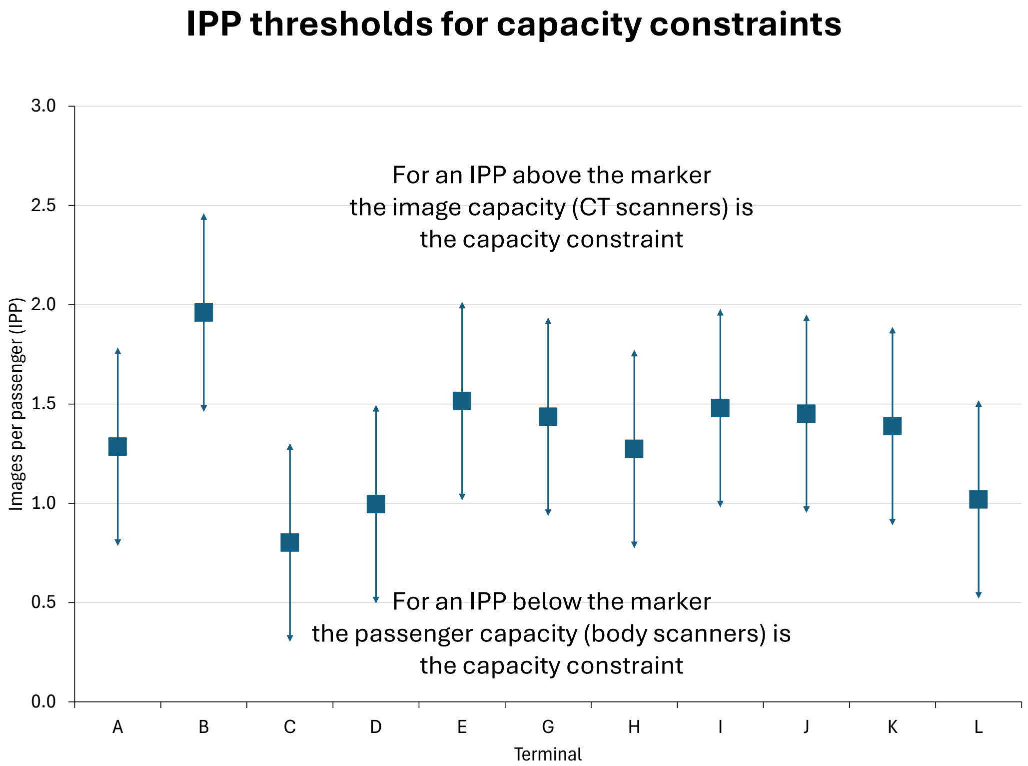

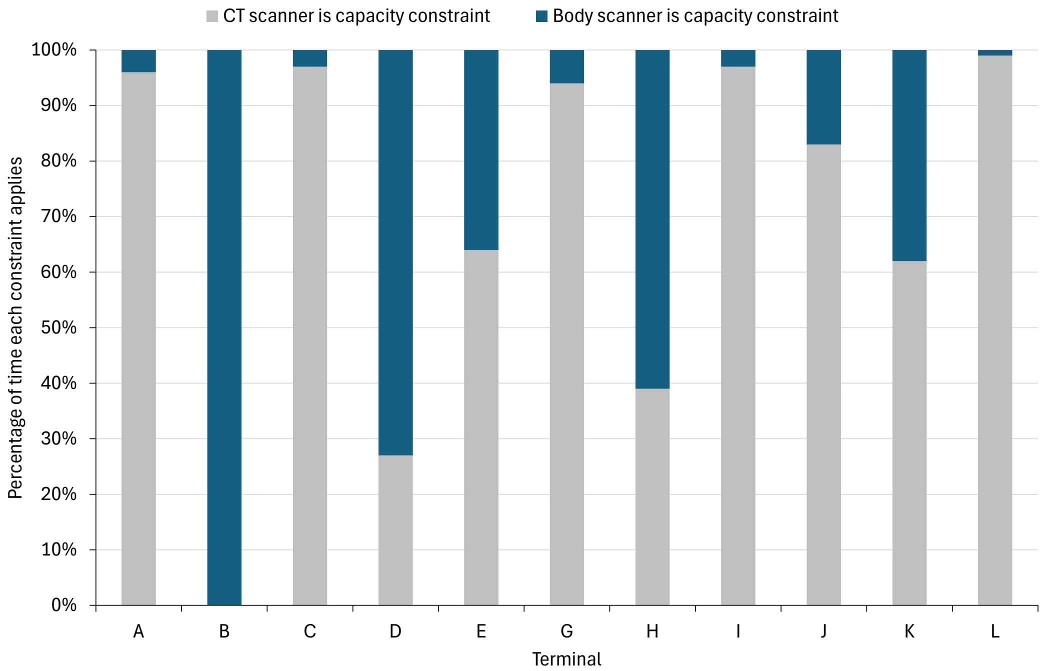

In the right-hand chart, a capacity constraint dominated by one of the scanning systems implies over-reliance on that system and under-utilisation of the other system. For example, in terminals A, C, G and L the CT scanner imposes the capacity constraint most of the time and, conversely, in terminal B the body scanner is the sole capacity constraint. It is better to have a balance between the two systems, e.g. as in terminals D, E, H and K. This can be achieved by better balance of the two types of infrastructure, e.g. CT/body scanner ratios and/or by improving the throughput of the constraining system. Note: peak throughputs were used to avoid distortions due to over-provision of capacity during off-peak periods

Throughput

In this section, we look at the distribution in peak hour throughput, for both CT and body scanners

Peak hour distributions are generally characterised by sharp narrow peaks, indicating a low-variability process with observations staying close to the mean value. Some distributions have multiple peaks likely driven by the opening and closing of lanes to match capacity and demand, and to manage queues,

Example peak hour throughput for one body scanner

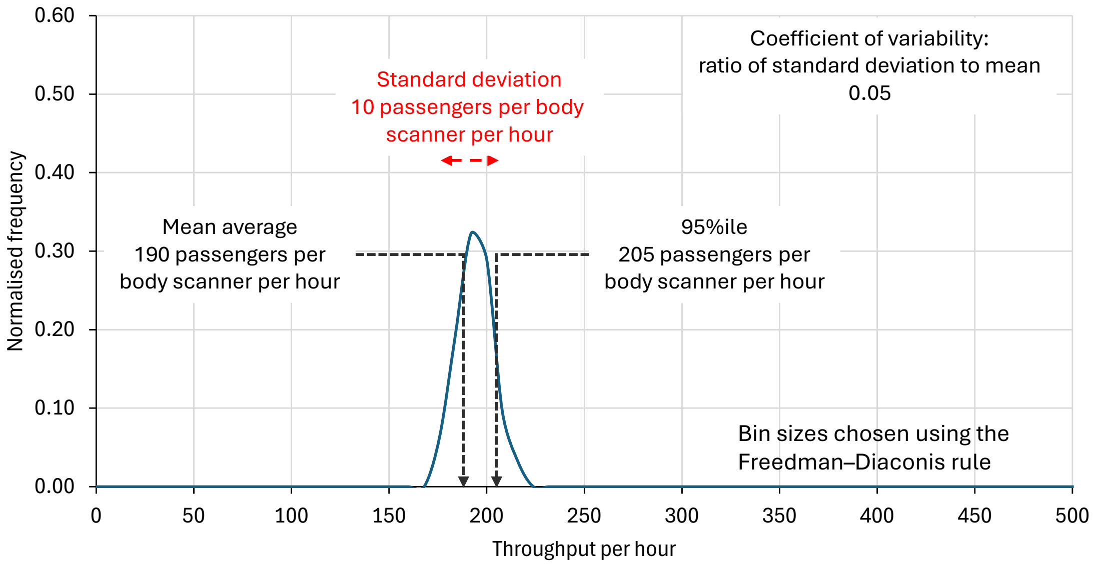

Example peak hour throughput for one CT scanner

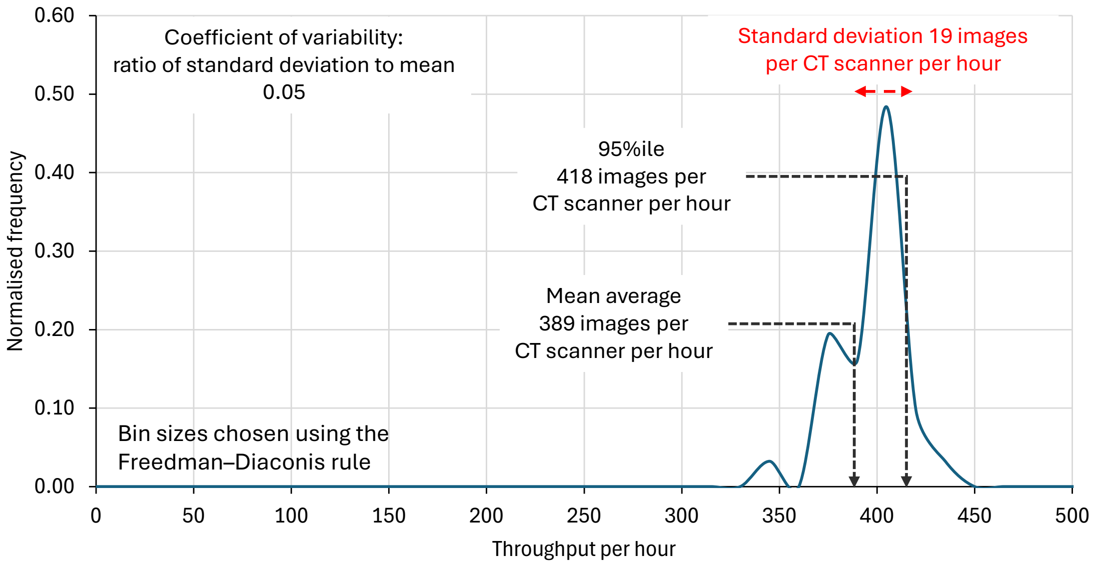

We compare throughput across terminals using four measures:

- Mean, representing the typical throughput.
- Standard deviation, representing the spread of throughput.
- 95th percentile, representing peak throughput, used as an indicator of maximum capacity for the single scanning system.
- The coefficient of variation, which is the ratio of the standard deviation to the mean. This gives a measure of the stability of the process. Typically, values range from >0.1 for low variability processes to >0.4 for unstable, highly variable processes.

HCL Hades Consulting Limited

Peak hour throughput performance for CT and body scanners, varies widely across our sample, by a factor of around two

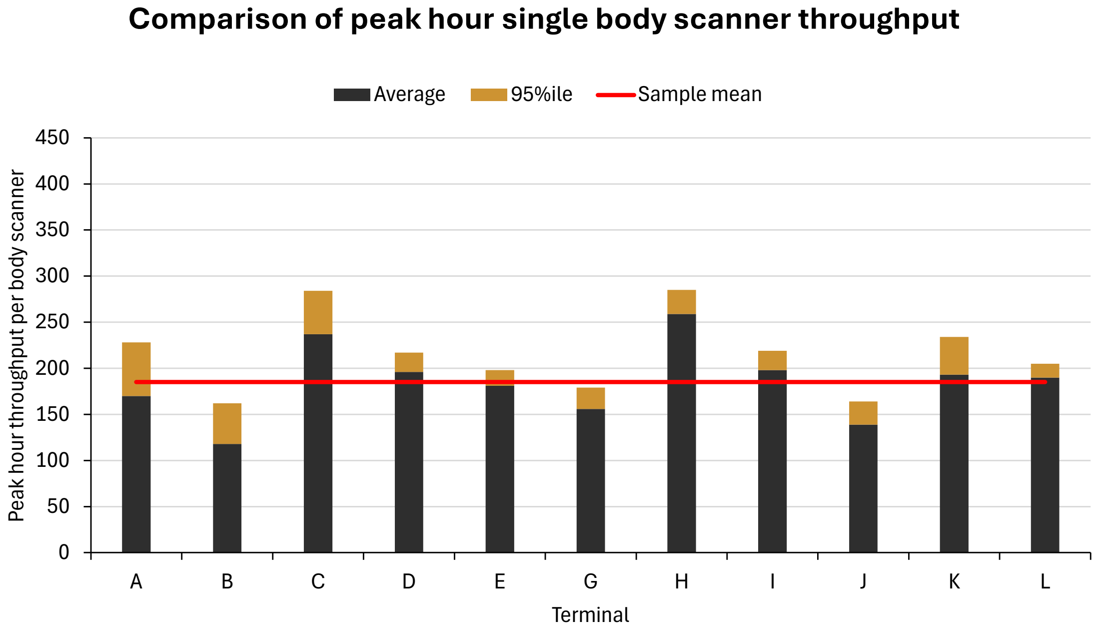

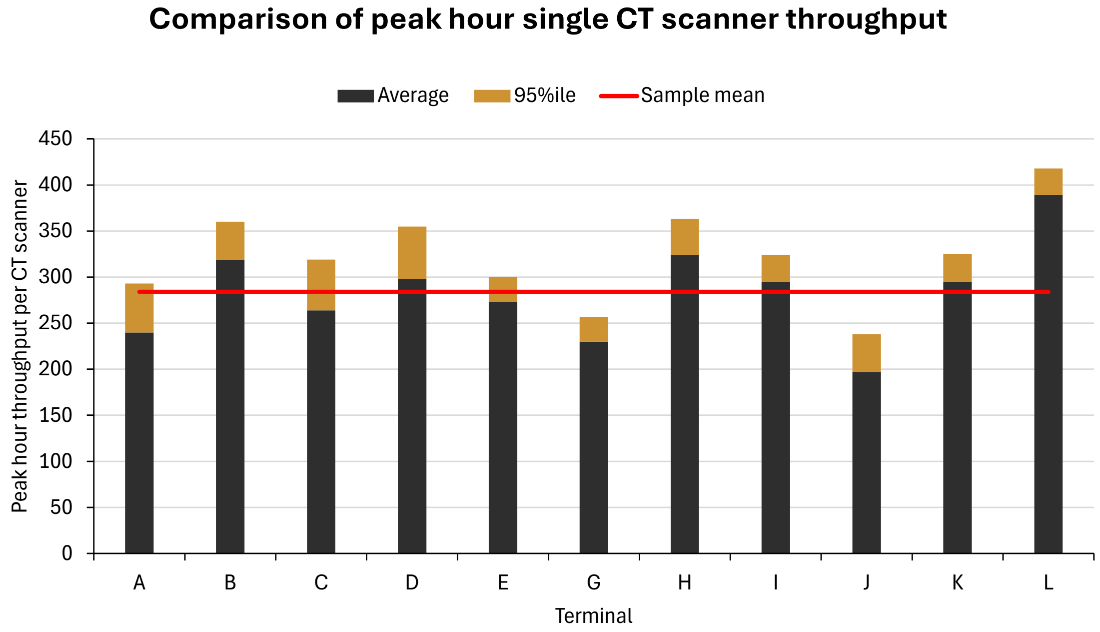

For each of the terminals, average throughput is similar to 95th percentile throughput, while the distributions of throughput are narrow and symmetric. This indicates that, as expected, at peak hours the system operates near to its maximum sustainable throughput

Although some throughput differences between terminals can be explained by infrastructure constraints such as lane length, based on our assessment of impact of lane length on throughput, the scale of variation is far greater than these factors can explain. Furthermore, the relative throughput of body scanners and CT scanners is not consistent across terminals. This implies that structural drivers are not the sole root cause of performance differences as otherwise, low body scanner throughput would be coupled with low CT scanner throughput (whereas we see the opposite in terminal B in the charts above). This suggests that most of the observed differences are operational or process-driven rather than structural.

Most terminals show high throughput, and stable CT and body scanner performance in their peak hours, as shown in the matrices below

Peak hour body scanner performance matrix

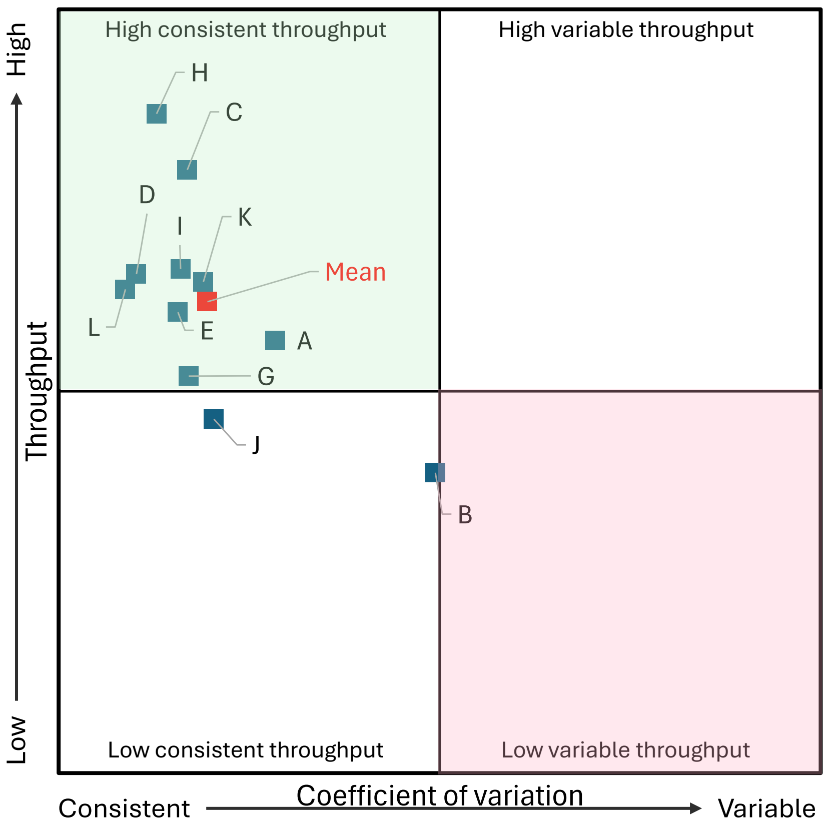

Peak hour CT scanner performance matrix

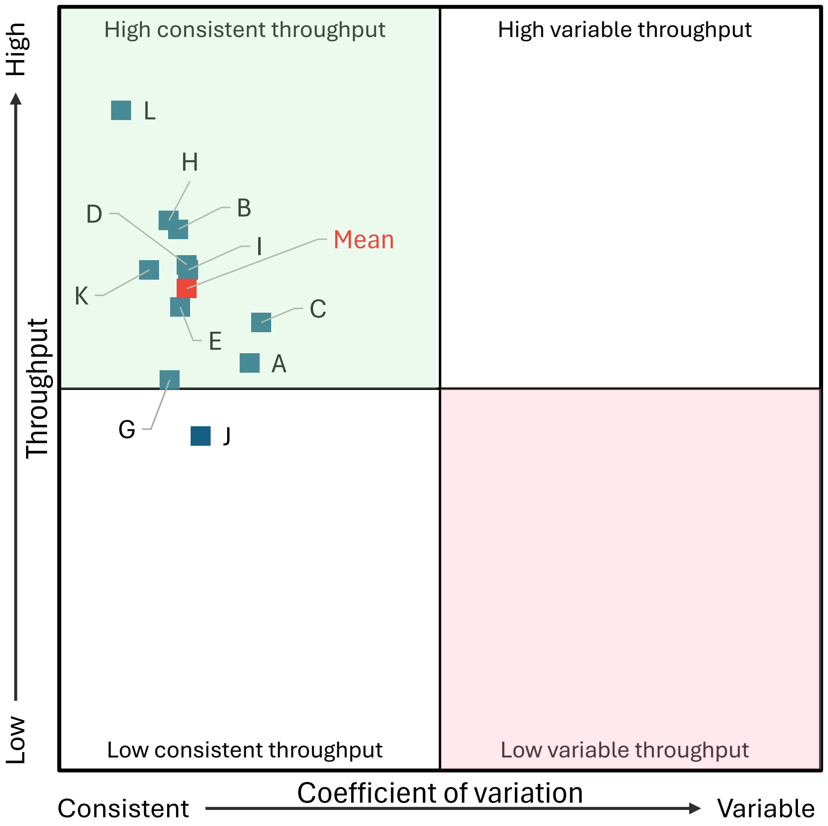

Using mean throughput and the coefficient of variation, the matrices show the relative position of each terminal for peak hour body scanner and CT scanner performance. Good performers are located in the top left quadrant, with relatively high and stable throughput.

Hades Consulting Limited

All-hours throughput is an indicator of sustained, continuous CT and body scanner performance across the day rather than just at peak hours

All-hours distributions are much broader than peak hour distributions with, as expected, lower averages and 95th percentiles than at peak hours. To illustrate these differences the charts below show the all-hours distributions for the same terminal used as an example for the peak hours distributions

Example all-hours throughput for one body scanner

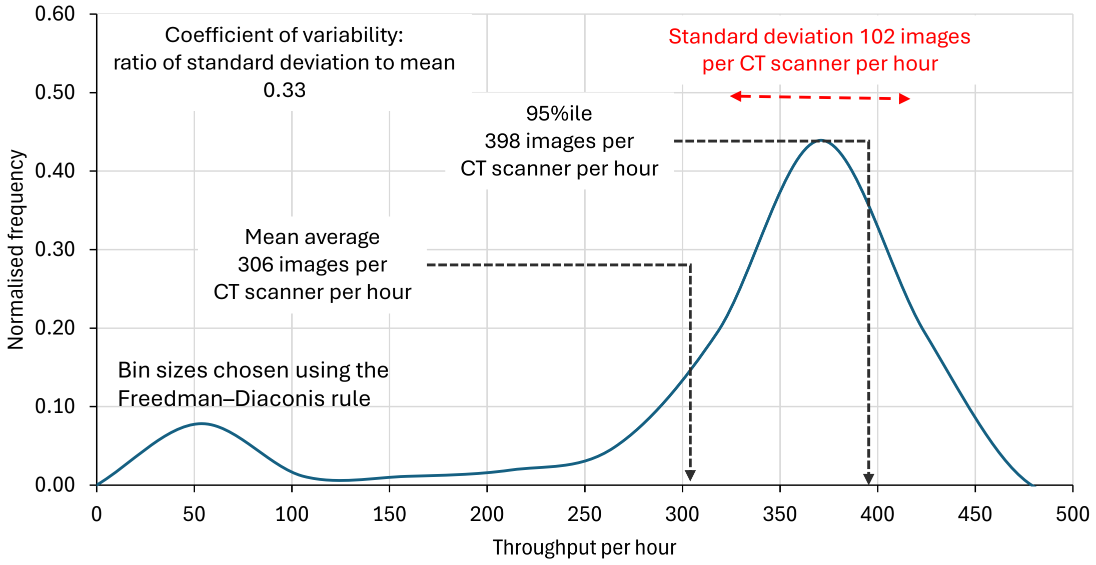

Example all-hours throughput for one CT scanner

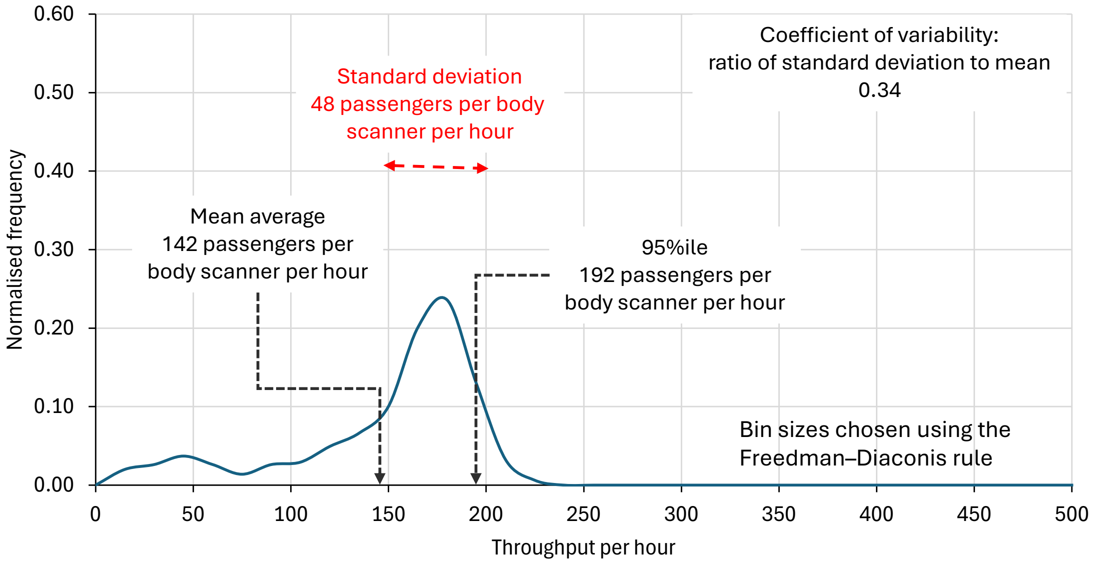

We use the same four measures to compare all-hours throughput across different terminals as we use for the peak hour comparison – mean, standard deviation, 95th percentile, and coefficient of variation.

The all-hours throughput distributions are broader than the peak-hour distributions because they combine multiple operating conditions across the day – such as mixing high and low demand. Although the all-hours distributions will never replicate the peak distributions because of random fluctuations and reaction times to changes, optimally matching capacity with demand will minimise the differences between the achieved throughputs and will deliver a more stable operation (as measured by the coefficient of variation).

There is a slightly smaller range across the sample for all-hours throughput (a factor of 1.8) than for peak hour throughput

Comparison of all-hours single body scanner throughputs

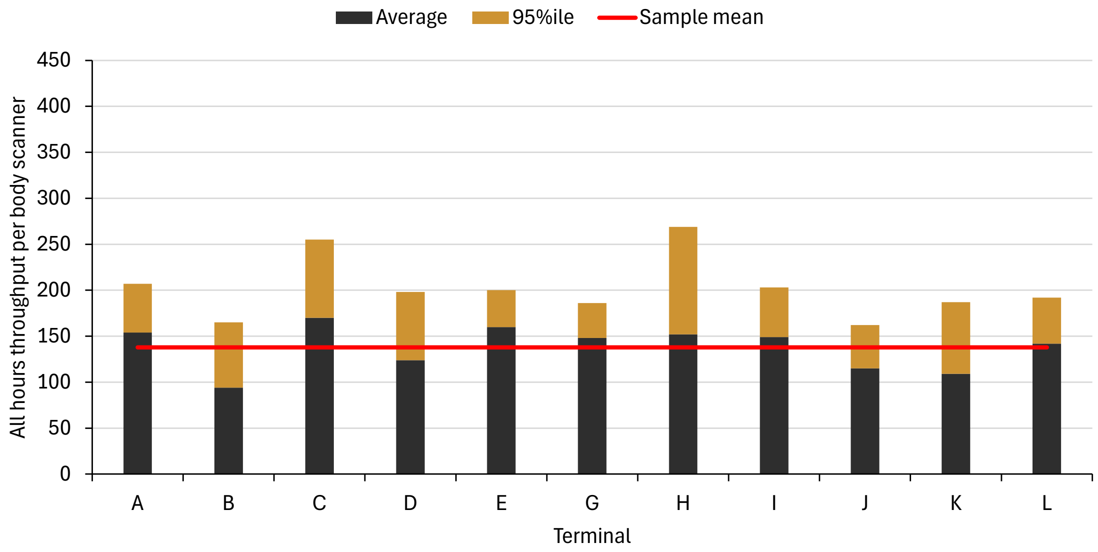

Comparison of all-hours single CT scanner throughputs

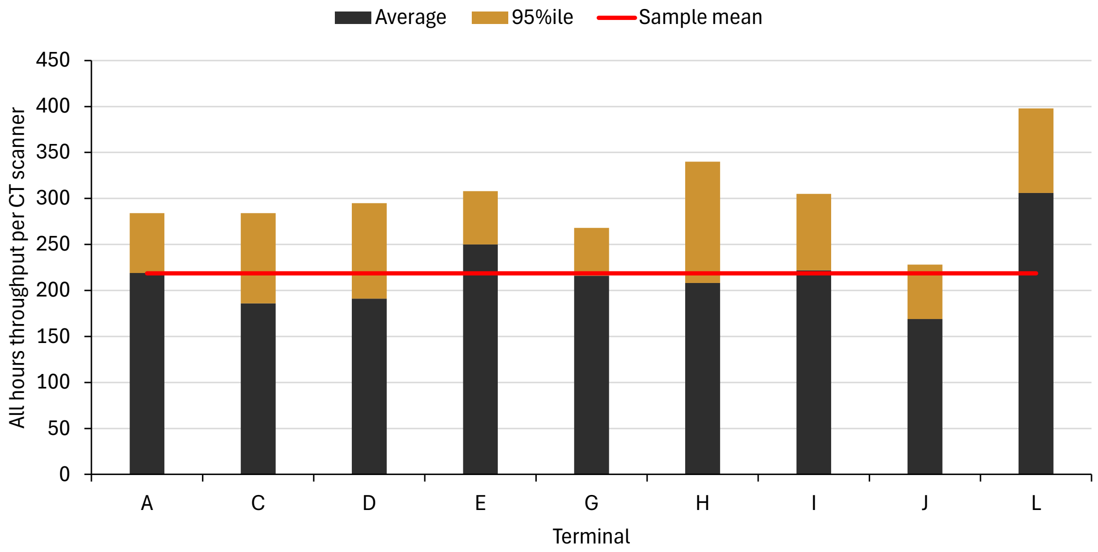

As shown in the table to the right, the relative throughput rankings of the terminals changes slightly from peak hour to all-hours. Again, this indicates that there are operational factors influencing performance alongside the structural constraints.

Low → High

Body scanner throughput terminal ranking	
All-hours	Peak hour
B	B
K	J
J	G
D	A
L	E
G	L
I	K
H	D
A	I
E	C
C	H

CT scanner throughput terminal ranking	
All-hours	Peak hour
J	J
C	G
D	A
H	C
G	E
A	I
No data	K
I	D
No data	B
E	H
L	L

HCL
Hades Consulting Limited

All-hours throughput is uniformly lower than peak hour throughput, as would be expected

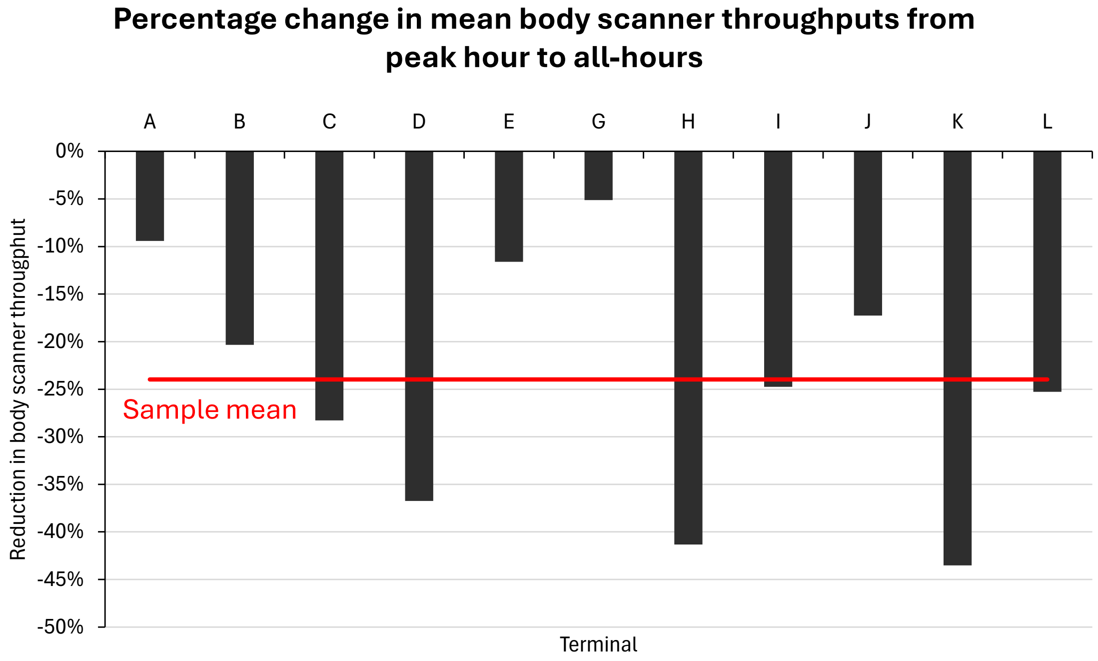

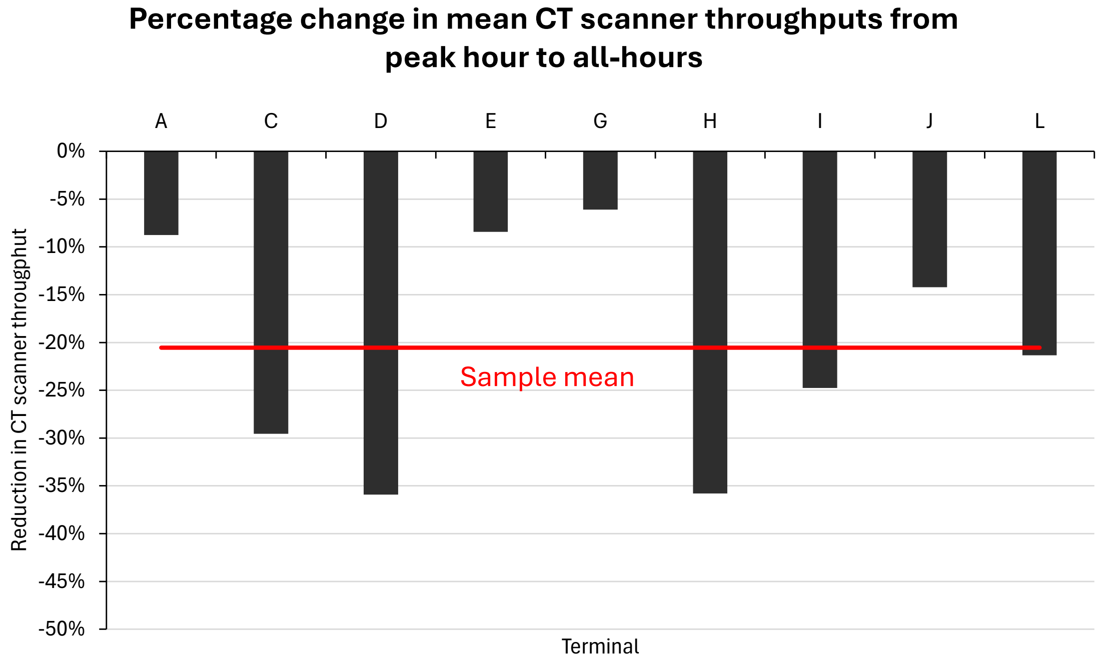

The charts show the percentage reduction in mean throughput for each terminal when comparing peak hour to all-hours. Although the samples are different, the charts suggest that reductions in body scanner throughput are larger than reductions in CT scanner throughput. Different terminals show very different reductions, from as low as 5% to nearly 45%. This is likely explained by those terminals that show low reductions:

- Operating near to capacity, i.e. at peak, for large parts of the day
- Having uniform demand across the day.

Conversely, those terminals that show large reductions may only operate at their peak for short, discrete periods and may be subject to wide variations in demand.

The performance matrices for all-hours throughput confirm lower throughput and more variability than their peak hour equivalents

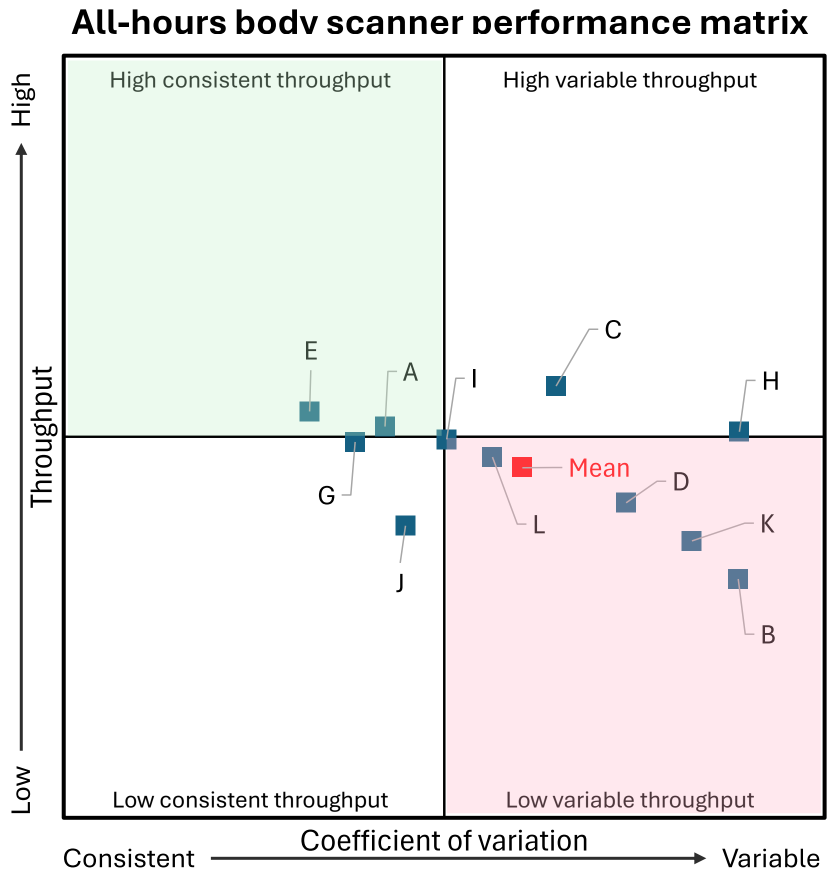

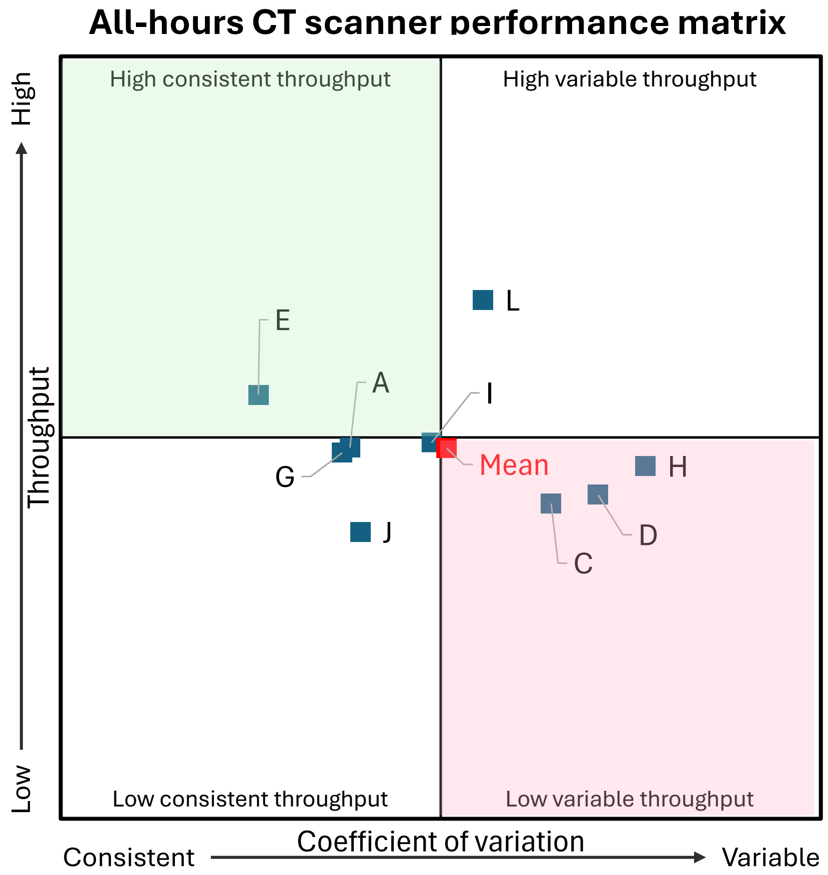

Some terminals' all-hours throughput performance is categorised in the bottom right quadrant indicating low and highly variable throughput. This is more prevalent for body scanner performance than CT scanner performance.

Productivity

We compare planned and realised staff per lane to measure the alignment between the staffing plan and actual operational deployment

Staff per lane may not be directly comparable across terminals because it can be driven by infrastructure layout, such as the number of body scanners per CT scanner and the concept of operations, e.g. the number of CT operators per scanner, etc.

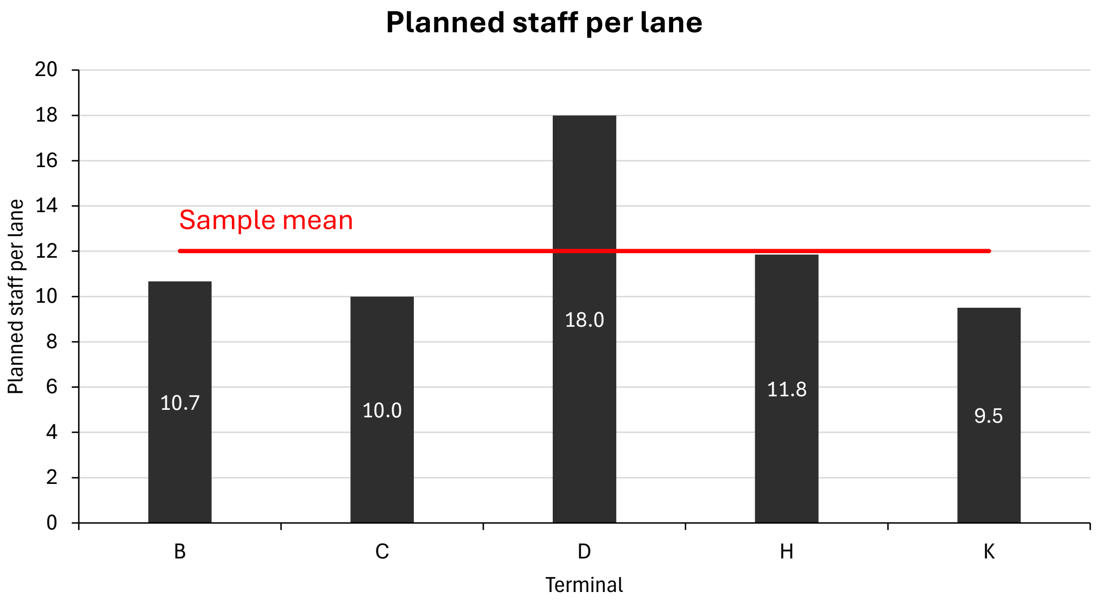

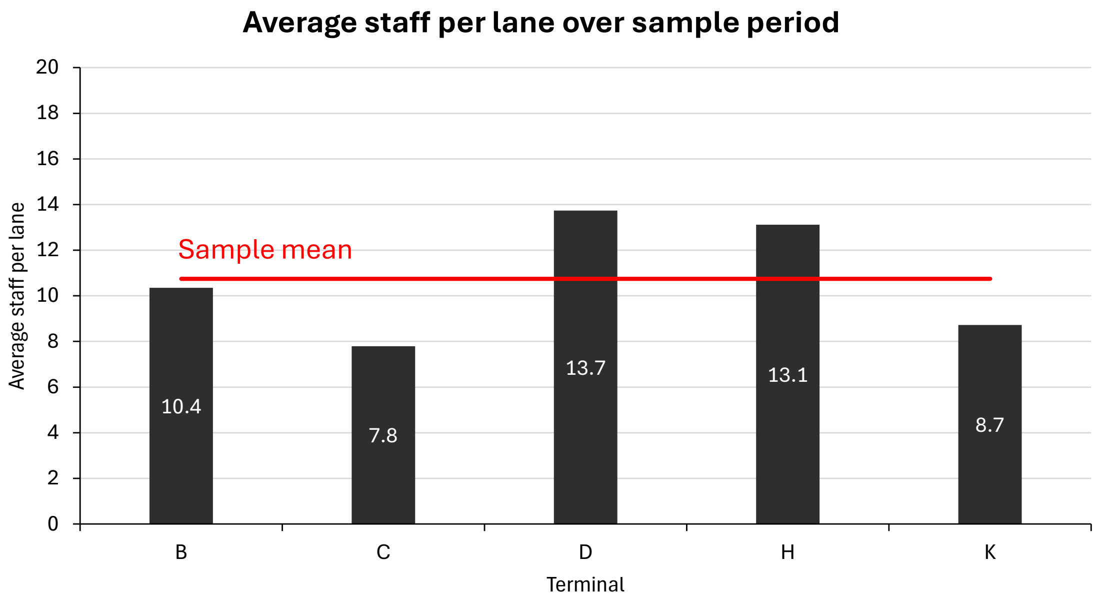

The availability of staffing data was very limited so the sample for comparison is very small.

As the charts show, actual staff per lane appears to be generally, but not always, lower than the planned staff per lane. These differences can be explained by staff absences, deployment to other duties and tactical adjustments if pressure on the system is lower than expected, e.g. opening fewer body scanners than planned. Actual staffing being greater than the plan is likely due to tactical adjustments to meet higher than expected demand and to manage queues.

Hades Consulting Limited

Passengers processed per hour per staff hour is a measure of staff productivity: it is possible to compare the plan with realised outcomes

This indicator ‘normalises’ for the infrastructure impact on staff per lane because, for example, more infrastructure and staff would be expected to process passengers more quickly

Planned pax per hour per staff hour over sample period

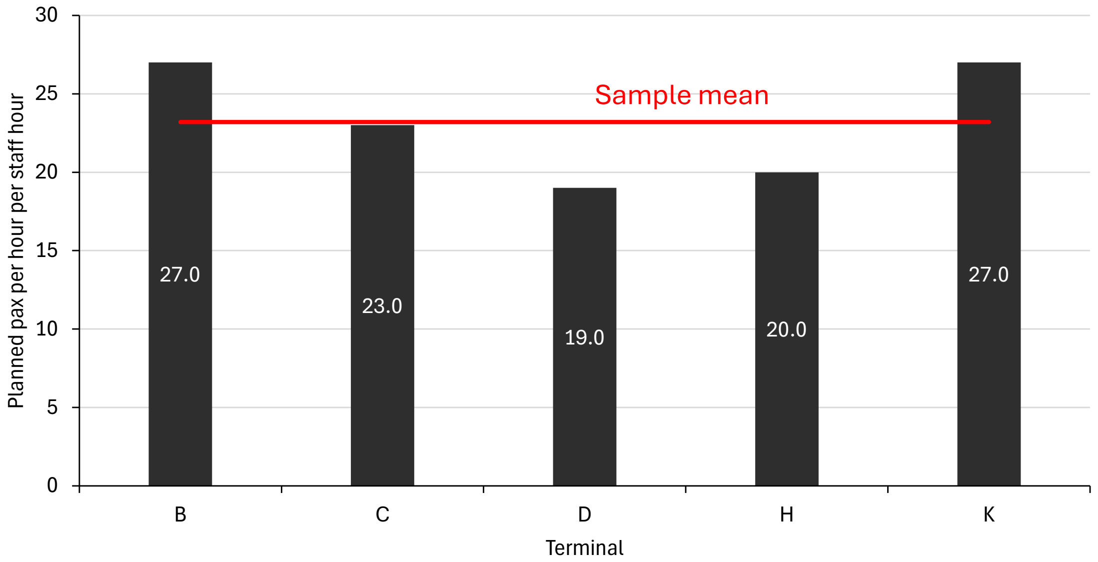

Average pax per hour per staff hour over sample period

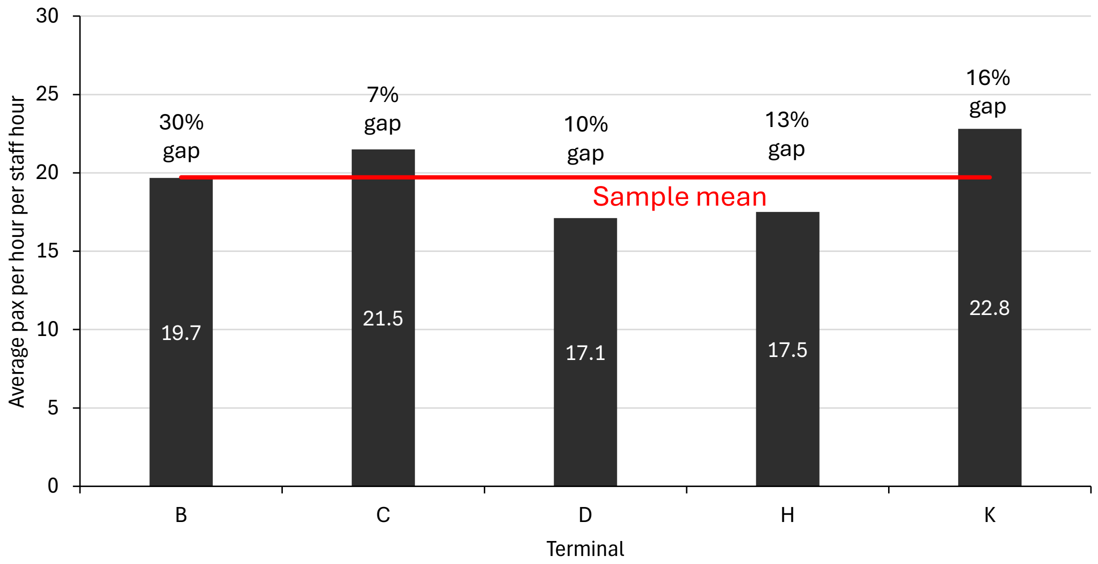

Again, the limited availability of staffing data means that the comparator sample is small.

Uniformly, the achieved productivity using this measure is lower than the planned productivity.

Although based on a very small sample, the productivity matrix suggests that productivity is generally high but variable

Productivity matrix

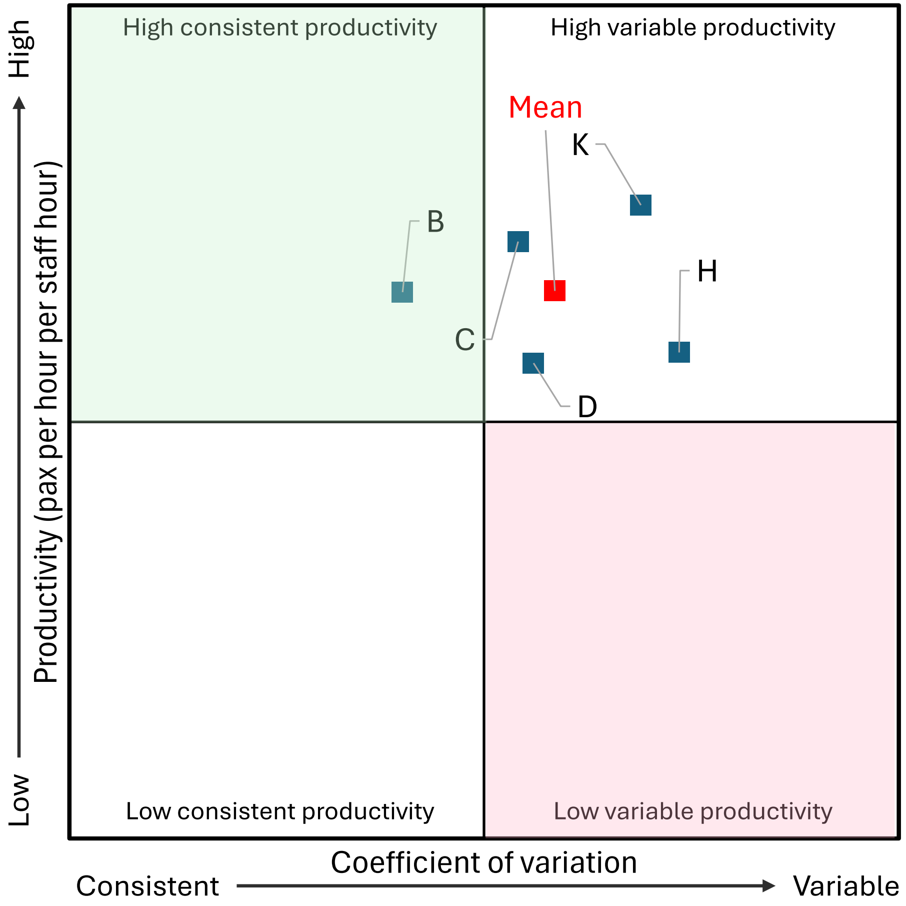

Typical pax per hour per staff hour distribution

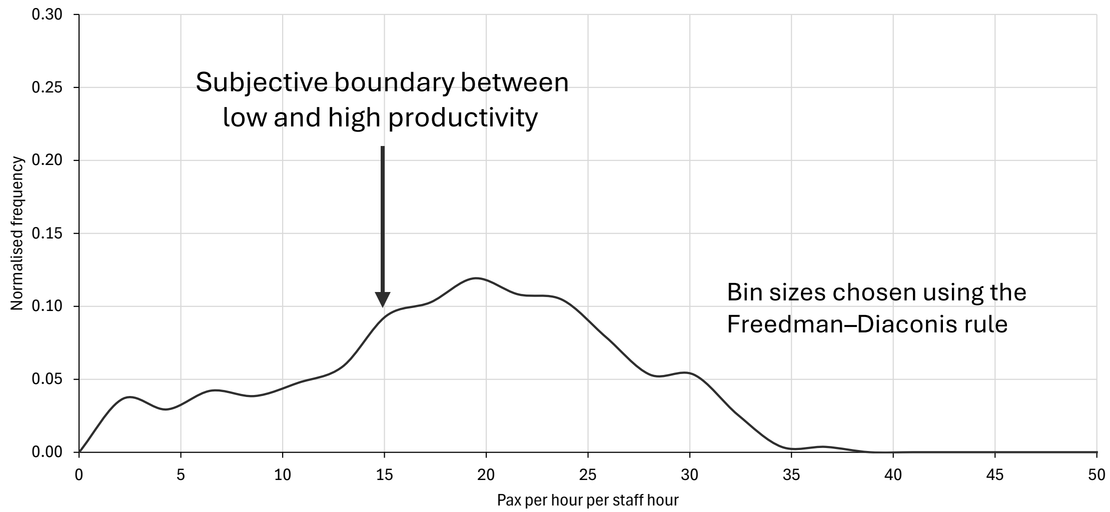

We can use mean productivity and the coefficient of variation (standard deviation divided by the mean) to compare the productivity achieved at each terminal and the stability of that productivity. This is shown in the productivity matrix. We set the boundary between low and high productivity at 15 passengers per hour per staff (based on our judgement) as illustrated in the chart above. We set the boundary on the coefficient of variation axis at 0.3 in line with standard practice.

The chart shows high productivity for all terminals but also high variability for all but terminal B.

Queues

Queue length is the most publicly visible measure of security performance: for their sample periods all the terminals performed reasonably in terms of queues

Quantitative comparison of security queues across airports is difficult because of the different ways in which waiting time is measured and processed. However, the chart below gives an indicative comparison of the queues at the different terminals.

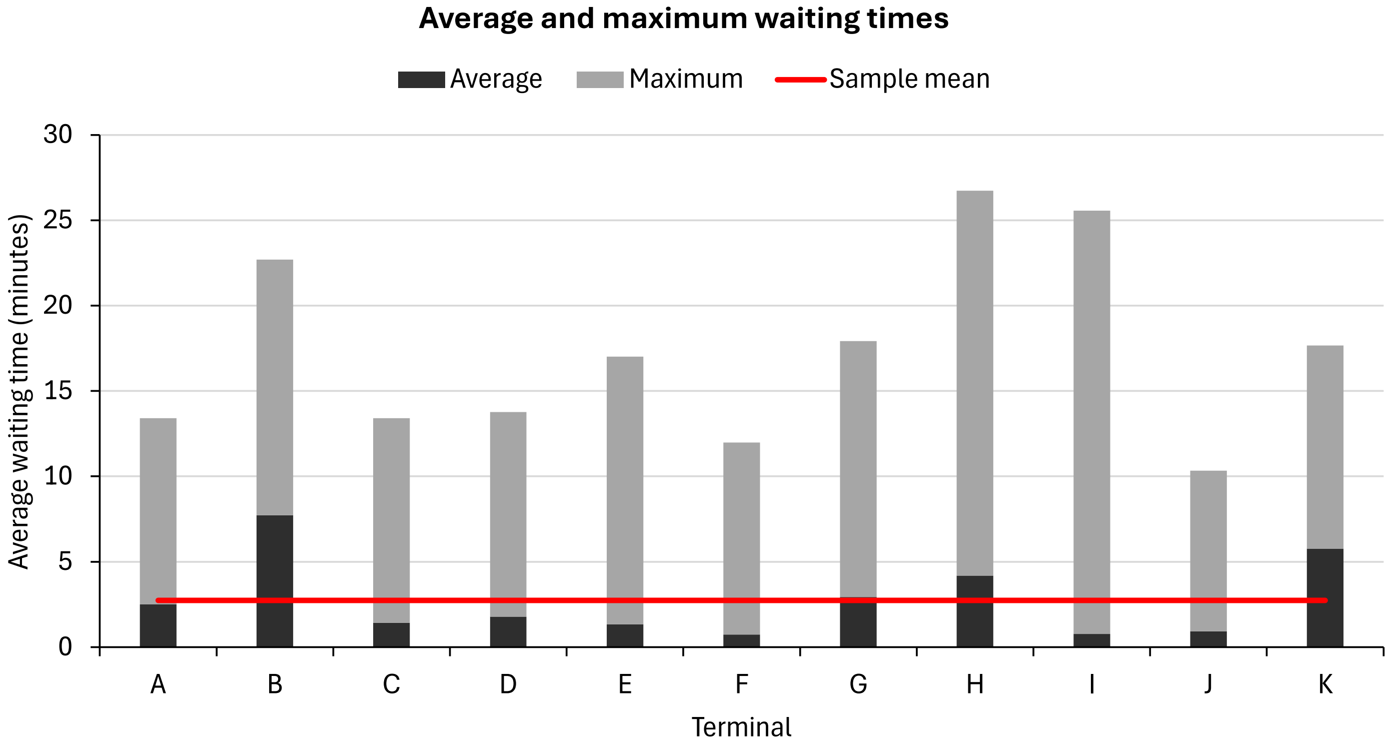

Proper statistical analysis of queue length is not possible because queues are reported in time intervals as an indicator of the waiting times experienced within that time interval, e.g. as a median or rolling average. It is, therefore, not possible to construct statistical distributions from these pre-processed queue length records. The chart above shows an approximation to the average queue length across the measurement month and the maximum queue experienced during that month.

Hades Consulting Limited

HCL
Hades Consulting Limited

Think

CEPA
economics matters