seeingmachines

Guardian Insights Report 2025-26

What 3.6 billion kilometres reveal about global trends in fleet safety



Introduction

Welcome to the 2025-26 edition of the Guardian Insights Report.

For decades, fleet safety has relied heavily on what happens after an incident: crash investigations, compliance records and other measures of past performance. These remain essential, but they cannot tell us everything about the risks drivers encounter throughout their working day. To make meaningful progress in preventing harm, we also need to understand what happens before an incident occurs.

The shift towards more proactive safety management is gaining momentum around the world. In Australia, changes to heavy vehicle safety frameworks are strengthening the focus on risk management and evidence-based decision-making. In the United Kingdom, plans to introduce advanced vehicle safety technologies reflect growing recognition of the role these systems can play in preventing crashes. Across Europe, requirements for advanced driver distraction warning systems have extended to all newly registered vehicles from July 2026. Meanwhile, in the United States, industry discussions increasingly emphasise the need to move beyond compliance alone and embed safety into everyday operational decisions. Together, these developments reflect a broader recognition that improving safety means looking beyond incidents to better understand and manage the risks that precede them.

Guardian contributes to this effort by providing insight into fatigue and distraction during real-world driving. Underpinned by more than two decades of behavioural science research, the technology combines in-cab alerts that prompt drivers to respond to emerging risks with fleet-level data that can help inform safety practices.

As the landscape for driver safety continues to evolve, so too must the way we understand and report on the data. Detection technology is advancing, regulations are changing around the world, and fleets are increasingly looking to extract meaningful insights from the information they collect.

In previous editions, we analysed fatigue and distraction region by region. This year, we have taken a different approach. As our technology and customer base have grown, we have found that some of the most valuable insights lie not in regional comparisons, which can be influenced by differences in fleet composition and system configuration, but in the patterns that emerge consistently across our global data.

This report focuses on those patterns, exploring:

- What the remarkably consistent timing of fatigue tells us about managing risk.
- How early drowsiness detection is identifying signs of fatigue before they progress.
- Why microsleep detection remains a critical safety net.
- Why repeated short glances away from the road now account for the vast majority of distraction events we detect on our Generation 3 systems.
- How the Guardian Centre turns detection into timely action when it matters most.

Ultimately, insight matters most when it informs action. By combining an understanding of driver risk with timely intervention, operational knowledge and ongoing evaluation, fleets can investigate patterns in their own operations, review existing practices and make more informed decisions about where safety initiatives may have the greatest value. This is not about treating every recorded event as a warning of an impending crash. It is about using evidence to ask better questions, identify opportunities for improvement and continually strengthen how risk is managed.

At Seeing Machines, our vision remains clear: zero transport fatalities. Achieving this requires a shared commitment across industry, regulators and operators, together with a willingness to keep learning as new evidence and capabilities emerge.

I invite you to explore this year's findings, consider what they reveal about fatigue and distraction, and reflect on how these insights can support safer operations within your own organisation.

Thank you for being part of this important work.

Dr Mike Lenné
Chief Safety Officer, Seeing Machines



About the report

Guardian, Seeing Machines' industry-leading driver fatigue and distraction technology, is underpinned by real-world driving data collected over billions of kilometres of travel. This report draws on deidentified data from Guardian-equipped fleets worldwide over the 12 months from 1 October 2025 to 30 September 2026, sourced from our Guardian Live platform. Some insights draw on subsets of this dataset. Early drowsiness figures cover Generation 3 systems from 20 April to 30 September 2026, while attention sharing figures include only Generation 3 systems with the feature enabled. All times are shown in each vehicle's local time zone.

- **62,862** vehicles
- **4,943** fleets
- **3,607,629,991** kilometres travelled
- **84,283,608** operating hours
- **5,528,994** recorded events

A global view

As Seeing Machines' technology continues to advance, so too does the data available to understand driver behaviour and fleet operations. To build on these developments and provide greater depth and context, the Guardian Insights Report continues to evolve.

This edition presents insights drawn from our global dataset rather than region-by-region results. Differences in fleet composition, system generation, configuration and event verification can have a strong influence on regional figures, so we have chosen to focus on the patterns that are consistent across markets.

Alongside this broader global perspective, this year's report also expands its analysis to include three new areas: early drowsiness detection, attention sharing and speed.

Normalising the data

To allow fair comparisons across different operating conditions, results are normalised to account for exposure. Fatigue events by time of day are measured against kilometres travelled, the basis used in previous editions, so that this year can be compared with the two before it. Measured against operating hours instead, this year's pattern is the same: the rate peaks at 4 am at about five times its 8 pm low. Early drowsiness, attention sharing and speed insights draw on newer Generation 3 data and are measured against operating hours.

Guardian event types

Guardian uses advanced algorithms to detect when a driver is fatigued or distracted. When an event is detected, the driver is alerted in-cab in real time, while event data is sent to our secure online monitoring platform, Guardian Live. High-risk events are also reviewed and classified by highly trained analysts in our 24/7 Guardian Centre.

Fatigue events

- Microsleep: A microsleep event occurs when Guardian detects that a driver has briefly and involuntarily fallen asleep while driving, typically for just a few seconds and often without the driver realising it.
- Drowsiness (event): A drowsiness event is an analyst-reviewed fatigue classification where a driver shows clear signs of unsafe drowsiness – such as heavy eyelids, slow or repeated long eye closures, or struggling to stay awake – without meeting the criteria for a microsleep event.
- Early drowsiness (state): Early drowsiness occurs when the Guardian unit detects signs that a driver is becoming tired and may be at increased risk of an accident. Unlike a drowsiness event (see above), this state is detected automatically in real time and is not reviewed or classified by an analyst.

Distraction events

- Long glance away: A long glance away event occurs when a driver takes their eyes off the road for an extended period.
- Attention sharing: An attention sharing event occurs when a driver repeatedly divides their attention between the road and another object or task, such as an in-cab device, mobile phone, or other equipment.
- Mobile device: A mobile device event occurs when a driver is detected handling, operating, or interacting with a mobile phone or similar handheld electronic device*.
- Other distraction: The driver is observed to be undertaking an activity that clearly reduces their ability to safely control the vehicle, but it does not meet the criteria for another distraction classification.

* For most of the reporting period, mobile device events were identified through analyst review of footage showing a driver handling a portable electronic device. Automatic mobile device detection was introduced for some Guardian Generation 3 fleets in April 2026 and more widely from September 2026. These system-detected events have been excluded from this report to maintain comparability with the previous year's data and will be incorporated into the report from next year.

Insight 1: Fatigue follows a predictable rhythm

The same pattern, year after year

Of all the patterns in Guardian data, the timing of fatigue is the most consistent. Across markets, fleet types and reporting periods, fatigue rises through the night and peaks in the early hours of the morning, before easing through the day to a low in the early evening, then climbing again at night.

When we overlay the last three years of global fatigue data, the daily profile is very similar. Event volumes have changed as our customer base and technology have grown, but the shape of the curve has not. This year, the fatigue rate peaked at 4 am, about five times its 8 pm low.

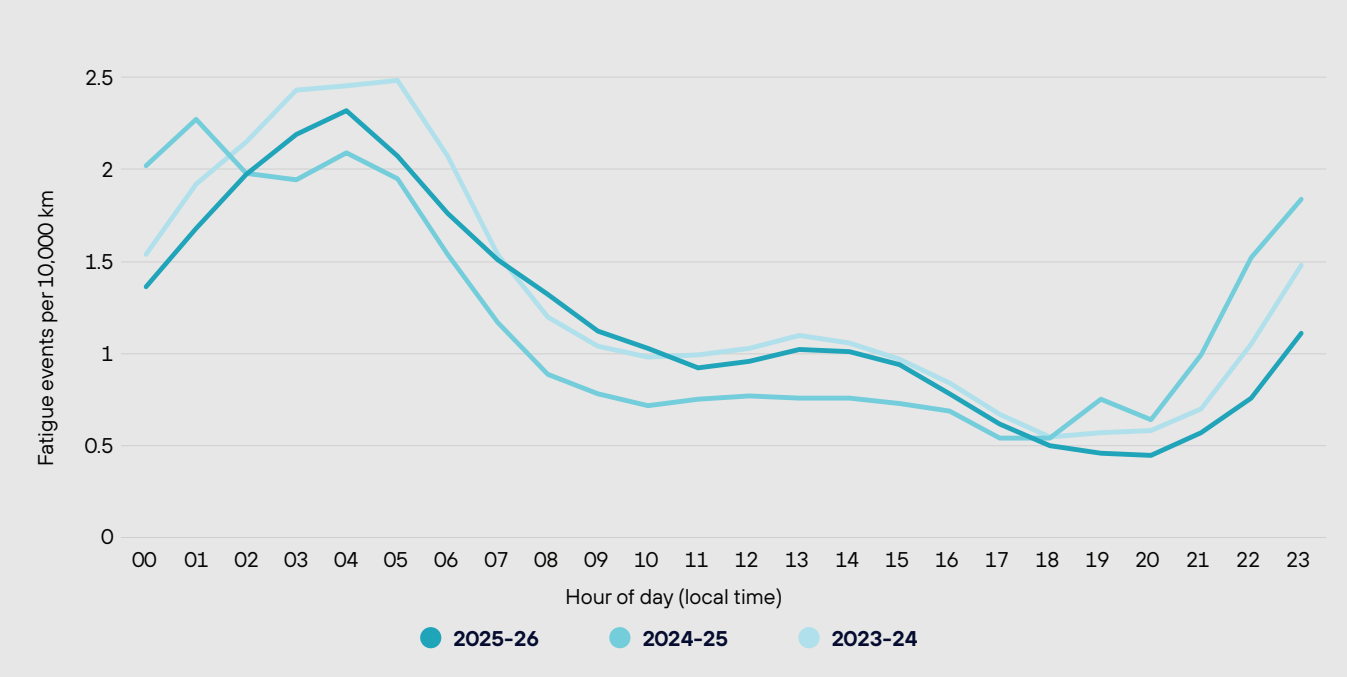


Figure 1: Fatigue events by time of day



4 am fatigue rate peaked at about **five times** its 8 pm low

37% of fatigue events occurred between **10 pm and 6 am** – a period covering just **25% of kilometres** travelled



A second measure, the same rhythm

Guardian early drowsiness detection, which measures how long drivers spend in an elevated drowsy state, shows the same rhythm independently*.

Drivers were in an elevated drowsy state for 7.5% of operating time between midnight and 5 am – around 2.5 times the 3.0% recorded between 1 pm and 5 pm.

Overall, 39% of drowsy time occurred between 10 pm and 6 am, even though this period made up only 25% of operating time.

Variation by day of week was much smaller, ranging from 4.5% of operating time on Fridays to 5.0% on Sundays. Weekends include more night-time driving, which accounts for some of that difference.

What the data suggests

The overnight peak is well known; its scale is less widely appreciated. The hours between 10 pm and 6 am accounted for 37% of fatigue events but only 25% of kilometres travelled. Early drowsiness shows almost the same concentration in those hours. Two different measures point to the same conclusion: a small part of the day carries a disproportionate share of fatigue, consistent with the circadian low in the early hours.

What this means for fleets

Because the pattern is so consistent, fleets can plan for it. Reviewing which shifts and routes run between 10 pm and 6 am – alongside their own fatigue events and, on Generation 3, drowsy time in those hours – can show where schedules, rest opportunities and fatigue management practices may make the most difference.

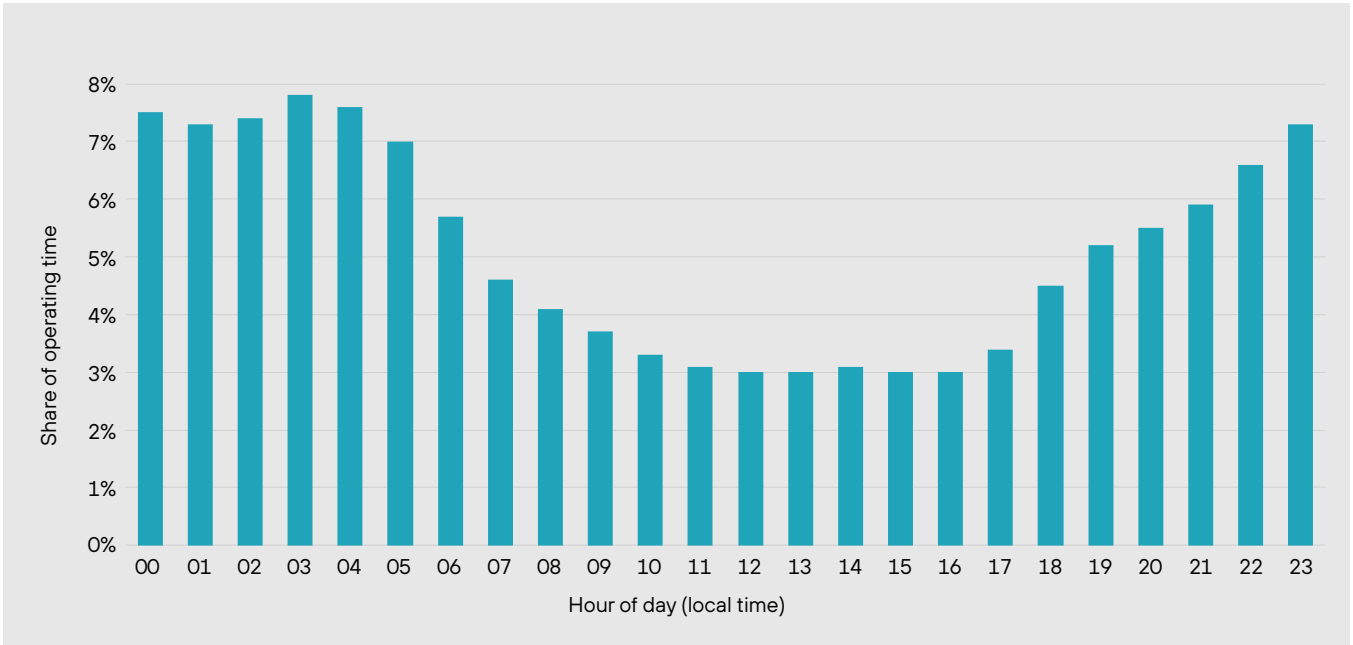


Figure 2: Share of operating time spent drowsy, by hour of day

* Findings are based on global data from Guardian Generation 3 systems between 20 April and 30 September 2026. Operating time is the total time the vehicles were operating, whether moving or stationary. Drowsy time is counted only while the early drowsiness detection algorithm was active and able to see the driver.

Insight 2: Early drowsiness detection gives drivers time to act

Identifying fatigue before microsleeps

Fatigue is a significant driver safety risk that can emerge well before a microsleep occurs. As alertness declines, drivers show earlier signs of drowsiness – creating an opportunity to recognise elevated risk before it further compromises safe driving.

Time of day is only one influence on alertness. Drowsiness also builds with time awake, known as sleep pressure, and is shaped by how well and how recently a driver has slept. Conditions in and around the vehicle also play a part. Time of day is predictable, but these other factors differ from driver to driver and from day to day. That is why drowsiness can appear in the middle of the day or early in a trip, and why measuring the driver's state directly adds to what a roster alone can tell a fleet.

Guardian extends fatigue detection beyond microsleeps by continuously assessing a driver's level of drowsiness, using principles informed by the Karolinska Sleepiness Scale (KSS), a nine-point scale widely used in sleep research and global safety regulation and policy.

Rather than relying on a single sign such as yawning, which is unreliable because people yawn for many reasons, the technology analyses a range of features in real time, including eye closure, eye closure rates and facial dynamics. When a driver exceeds the system's drowsiness alert threshold (above KSS level 7), Guardian issues an in-cab alert, giving them the chance to take corrective action, such as stopping for a rest.

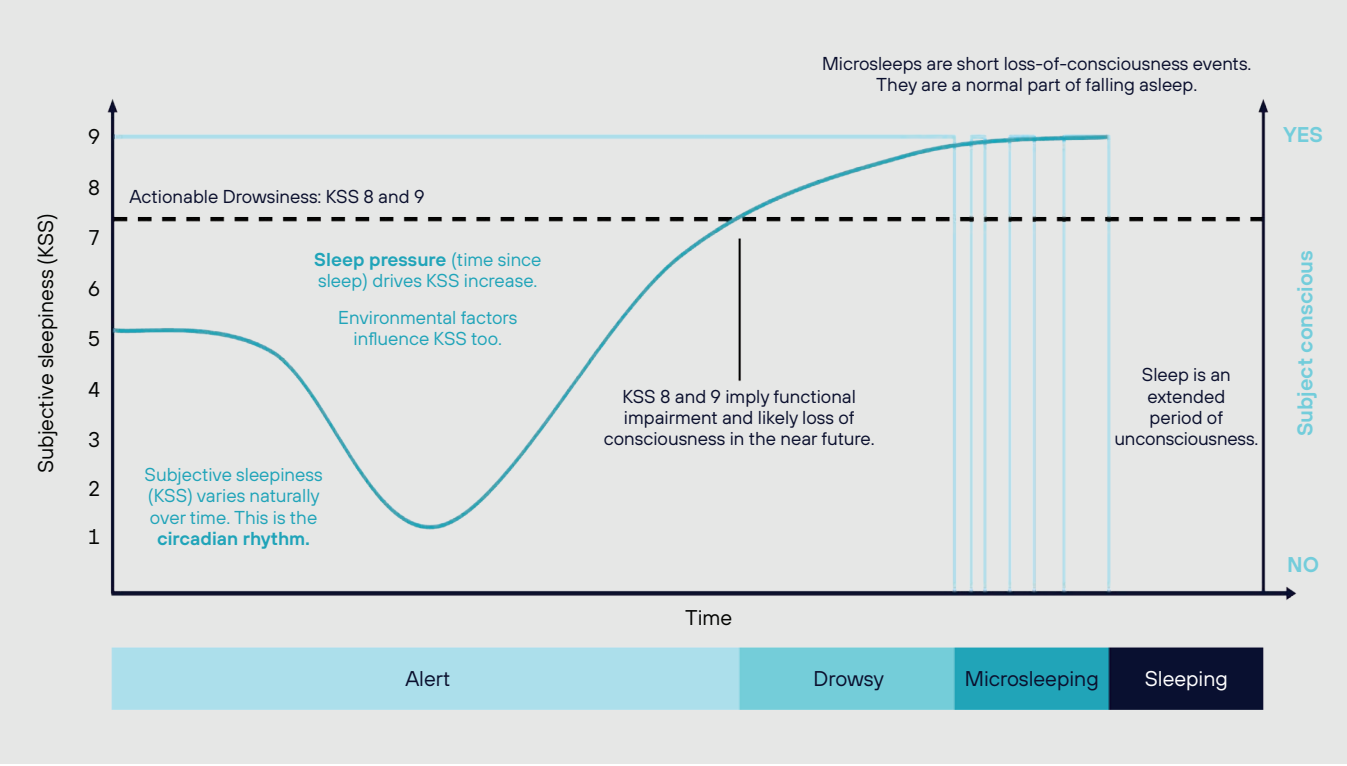


Illustration 1: How fatigue progresses, from alert to asleep *(KSS = Karolinska Sleepiness Scale)*

What the data shows

Between 20 April and 30 September 2026, drivers spent 4.8% of operating time in an elevated drowsy state. Around one in seven Generation 3 trips of 10 minutes or more (13.8%) included at least one drowsy episode.

Time behind the wheel matters. The likelihood of a drowsy episode rose steadily with journey length, from 5.0% of trips lasting 10–30 minutes to more than half (51.9%) of trips lasting eight hours or more. Among trips with a drowsy episode, the first was recorded a median of 28 minutes after the trip began.

Why this matters

Journey length is only part of the picture. The data shows that drowsiness can emerge relatively early in a trip, providing drivers and fleets with an additional source of information when assessing fatigue risk.

Early drowsiness detection also gives fleets a broader view of fatigue. Microsleeps are among the most serious fatigue events, but they are less common. In a typical month, around one in nine Generation 3 vehicles recorded a microsleep, compared with almost four in five that recorded drowsy time. Across the five-month period, 62% of vehicles recorded drowsy time without recording a single microsleep.

Because drowsiness is measured continuously, it provides a more complete picture of fatigue across vehicles, shifts and routes – including where no microsleep occurs.

This capability is already operating in commercial fleets, providing real-world insight into emerging signs of fatigue. It also reflects the direction of international driver monitoring requirements, where the detection of drowsiness is increasingly recognised as an important safety capability.

What this means for fleets

Because drowsiness is measured continuously rather than as occasional events, it gives fleets a more detailed view of where fatigue is developing. Comparing drowsy time across shifts, routes and times of day can show where risk is building before a microsleep occurs, and where changes to schedules or rest opportunities may have the most effect.

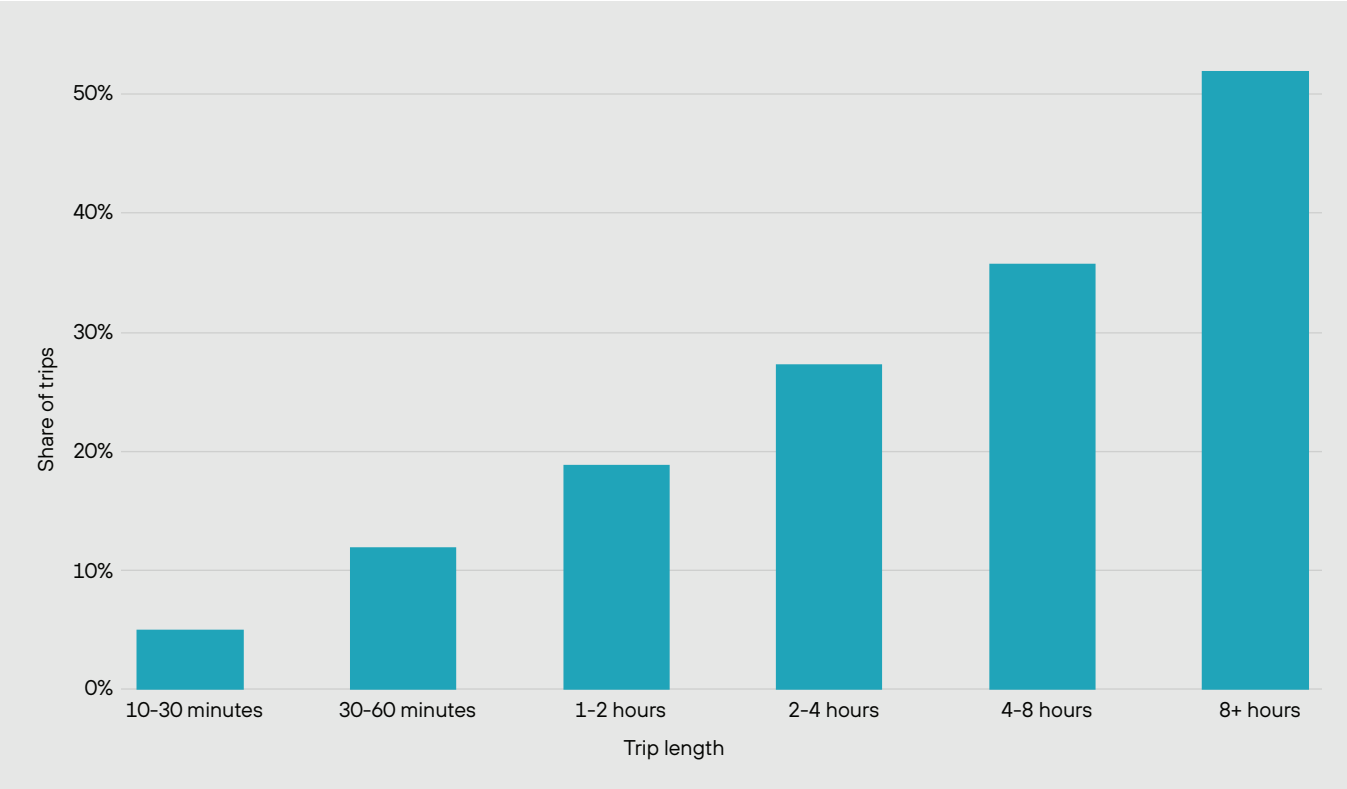


Figure 3: Trips with at least one drowsy episode, by trip length

* Findings are based on global data from Guardian Generation 3 systems between 20 April and 30 September 2026. Operating time is the total time the vehicles were operating, whether moving or stationary. Drowsy time is counted only while the early drowsiness detection algorithm was active and able to see the driver.

Insight 3: Microsleep detection remains the critical safety net

Different roles, working together

Early drowsiness detection gives drivers an opportunity to recognise increasing fatigue and respond before it progresses. However, an early warning cannot guarantee that a driver will respond appropriately or that fatigue will not escalate. Microsleep detection therefore remains an important and distinct part of fatigue protection.

A microsleep is a brief, involuntary episode of sleep and can occur without the driver recognising what has happened. At 100 km/h, a vehicle travels approximately 28 metres every second. A microsleep lasting only a few seconds can therefore mean travelling a substantial distance while the driver is not fully in control.

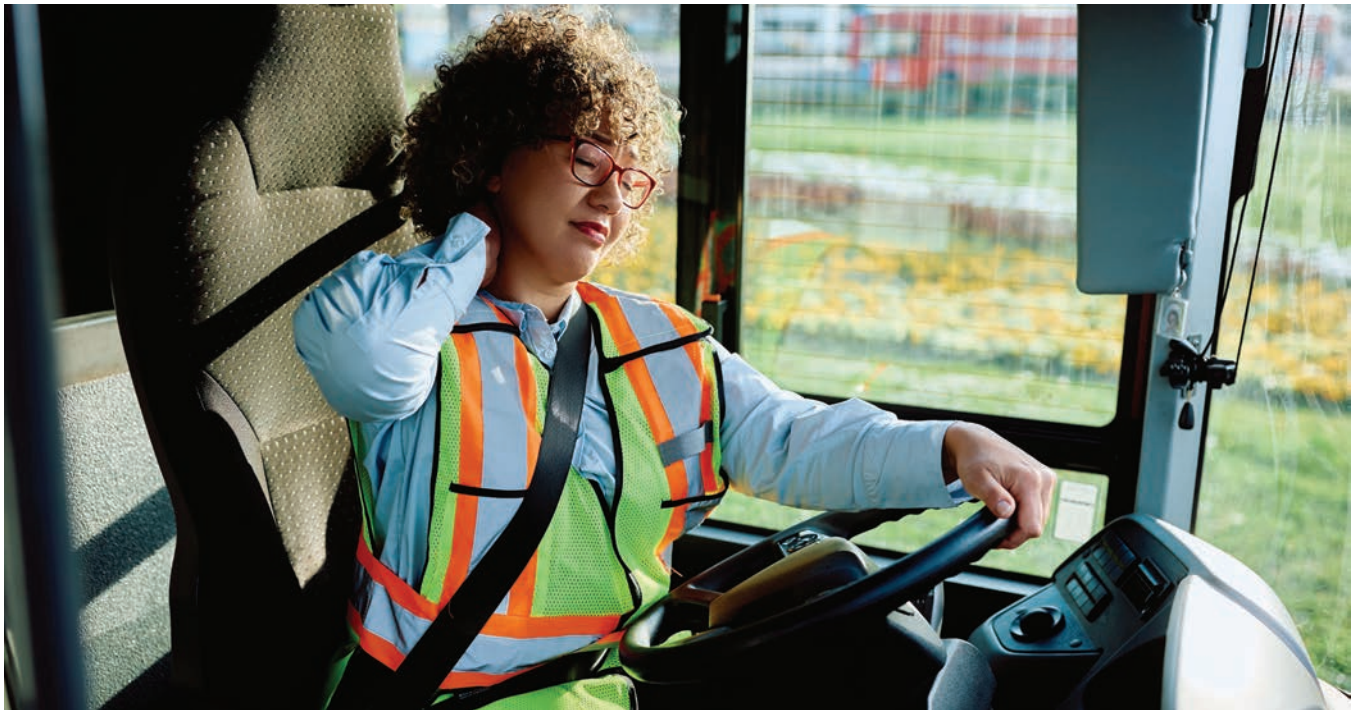
The two capabilities have complementary roles. Early drowsiness detection provides an in-cab warning when signs of an elevated drowsiness state are identified, giving the driver an opportunity to take action while they are still able to.

Microsleep detection responds to a more serious event: it triggers an immediate in-cab alert designed to wake the driver, and it supports escalation to the fleet. Microsleep events are reviewed by Guardian Centre analysts and, where a customer has a Fatigue Intervention Plan in place, a confirmed fatigue event can prompt contact with the nominated representative.

Guardian has detected microsleeps in real-world fleet operations for many years. The introduction of early drowsiness detection adds an earlier view of fatigue; it does not replace the need to identify and respond to microsleep events. The two also provide different kinds of information: microsleep events are infrequent markers of the highest risk, while drowsiness state data are continuous and show the underlying driver state.

Why this matters

Early drowsiness detection and microsleep detection are not interchangeable safety indicators. They address different stages: one helps drivers respond to emerging fatigue, the other wakes the driver and supports fleet action when a serious event is confirmed.





Higher stakes at every speed

Microsleep rates were similar between 40 and 100 km/h, at 1.0 to 1.3 events per 1,000 operating hours, and highest at 100 km/h and above, at 2.0. At these speeds a vehicle covers the most ground during a lapse, and any impact is most severe.

Lower speeds are not necessarily lower risk. A vehicle travelling at 50 km/h is often in a built-up area, among intersections, pedestrians and other road users, where there may be even less time and space to recover than on an open highway.

As regulators consider how driver monitoring should address fatigue, these results show that microsleep detection remains important at every speed – alongside earlier drowsiness detection.

What this means for fleets

A microsleep calls for a response beyond the cab. Fleets can check that a Fatigue Intervention Plan is in place, so that each confirmed fatigue event reaches someone who can act on it. Because microsleeps occur at every speed, that matters on urban and regional routes as much as on the highway. For fleets with Generation 3, early drowsiness data adds the earlier view: where fatigue is building before a microsleep occurs.

28 metres every second:
at 100 km/h, a three-second microsleep means travelling more than 80 metres without the driver fully in control

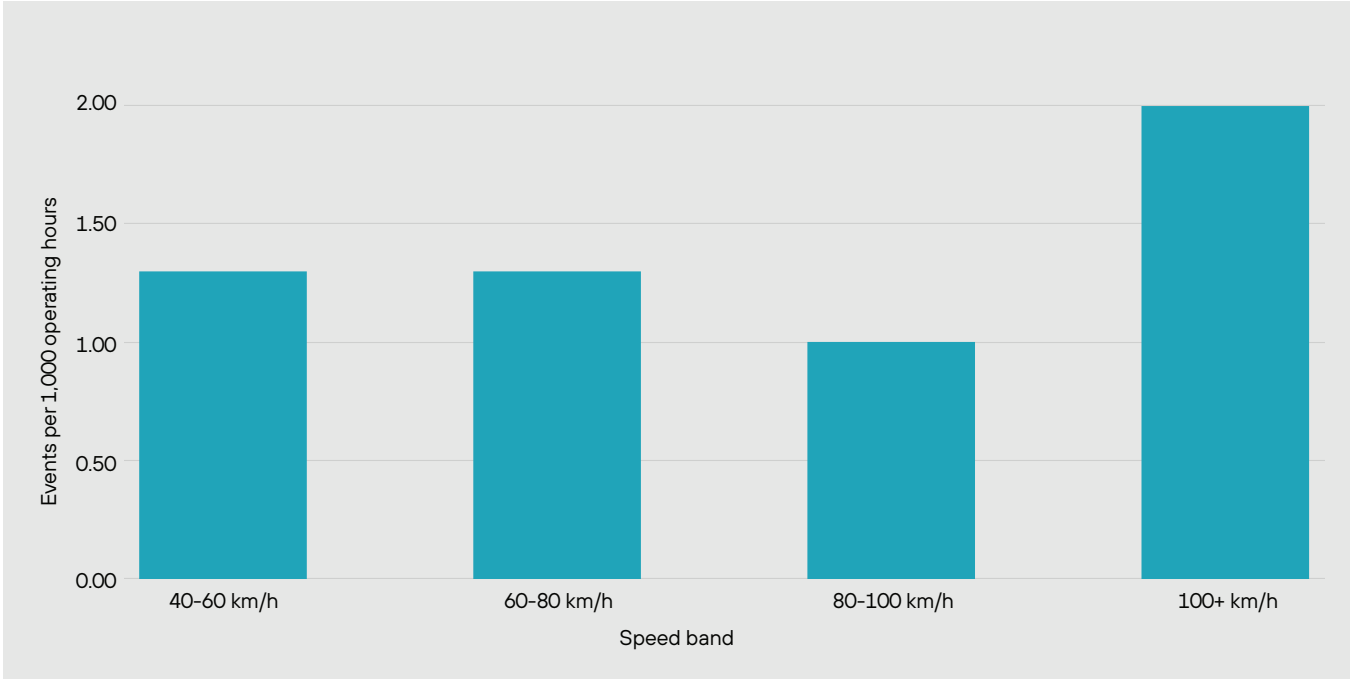


Figure 4: Microsleep events by speed band

Insight 4: Distraction is often many short glances, not one long one

When short glances add up

Driver distraction is not always obvious.

While a prolonged glance away from the road can indicate risk, repeated short glances can also compromise a driver's awareness and ability to respond to hazards. Individually, these glances may be brief, but their cumulative effect can be significant.

Guardian's attention sharing detection identifies patterns of repeated off-road glances that traditional distraction detection may miss.

Leveraging our automotive-grade eye gaze and head pose tracking, it measures the cumulative time spent looking away from the road across a series of glances. An event is triggered when that time exceeds 10 seconds, prompting an in-cab audio and visual alert.

The dominant form of distraction

Attention sharing accounted for 90% of all distraction events recorded on Generation 3 systems with the capability switched on – 1.76 million events in total*. This makes it by far the most common form of distraction captured, highlighting how frequently distraction involves repeated short glances rather than a single prolonged glance. For fleets, this provides a more complete picture of driver distraction, helping identify patterns that might otherwise go undetected and informing a more targeted approach to managing distraction risk.

Why this matters

Looking only for long glances or mobile device use may overlook a substantial category of visual distraction. Attention sharing provides a way to identify repeated short glances and explore when they occur.

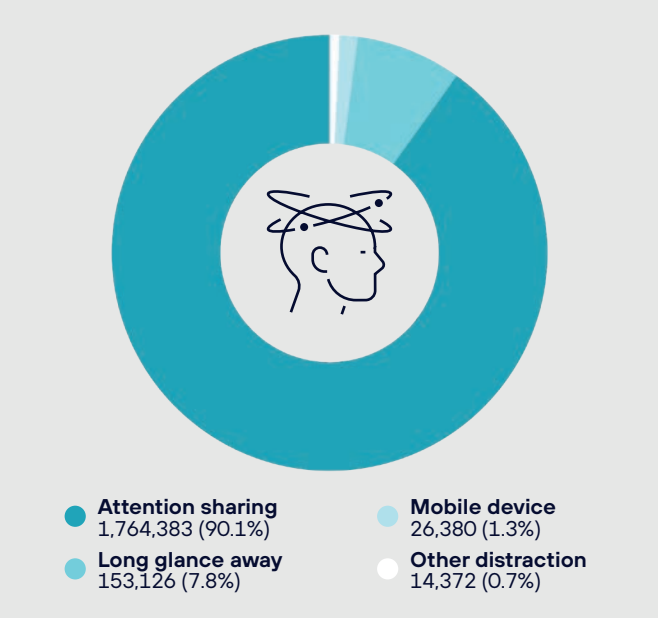


Figure 5: Share of distraction events

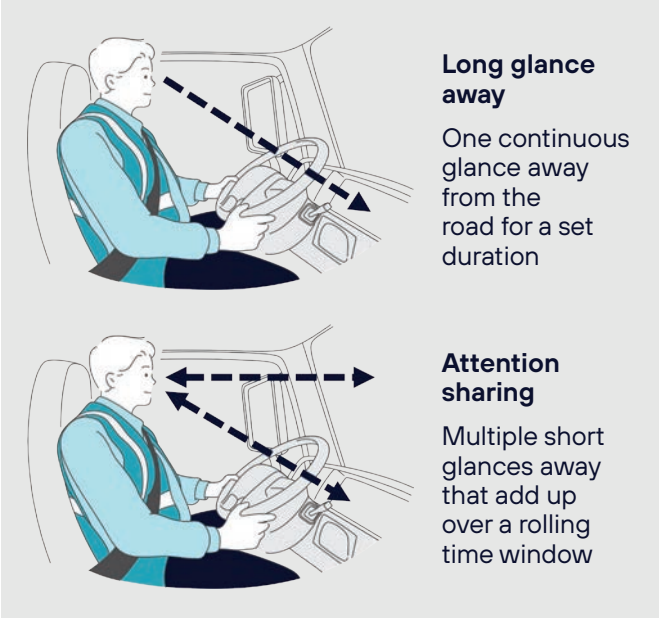


Illustration 2: The difference between distraction events



Later in the day, and at higher speeds

Once normalised by operating hours, a clear daily pattern emerges. Rates were lowest in the early morning, at 303 events per 1,000 operating hours at 5 am, rose through the morning, levelled off through the middle of the day and climbed again to a peak of 516 events at 7 pm – around 1.7 times the early-morning low.

Rates also rose with each speed band, from 587 events per 1,000 operating hours at 40–60 km/h to 951 at 100 km/h and above. At higher speeds, a vehicle covers more distance during each glance away from the road, leaving less time to respond to hazards.

What this means for fleets

Monitoring attention sharing as a distinct behaviour – particularly during evening operations and higher-speed driving – can help identify the in-cab tasks, devices or operating demands that draw drivers' eyes away from the road.

7 pm attention sharing peaked in the early evening, at around **1.7 times** the 5 am low

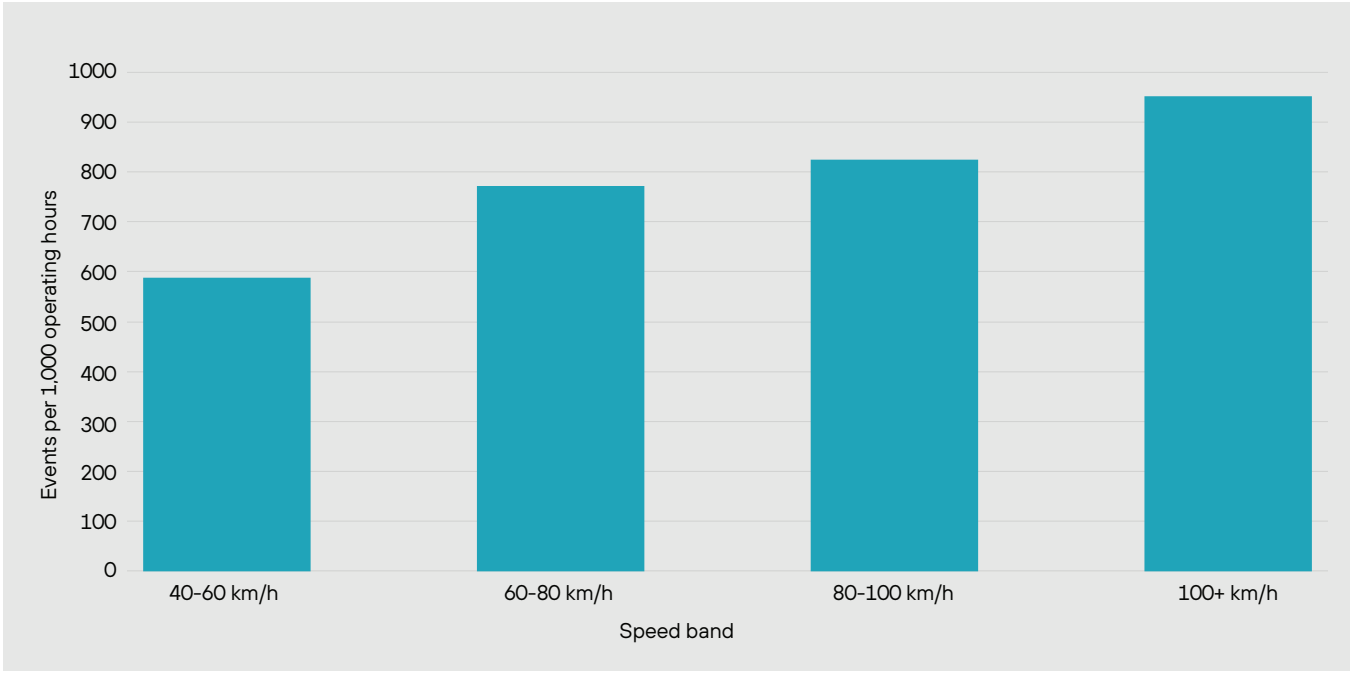


Figure 6: Attention sharing events by speed band

* Based on global data from active Guardian Generation 3 systems with attention sharing detection switched on (1 October 2025 – 30 September 2026). Attention sharing was switched on progressively during the period and not all fleets use it. Speed results below 40 km/h are not shown, as detection typically uses a minimum speed threshold of 40 km/h.

Insight 5: Human response turns detection into action

The Guardian Centre

Detection is an important first step. For high-risk fatigue events, human review and timely communication can provide additional context and help ensure that the right people are informed.

The Guardian Centre is staffed by trained analysts working 24 hours a day, seven days a week. Analysts review and classify high-risk events captured by Guardian.

Where a customer has an agreed Fatigue Intervention Plan, the team follows the plan for confirmed fatigue events, contacting nominated customer representatives by phone or email so they can respond in line with their operational responsibilities. Plans can be tailored to customer requirements, including different contacts for different days or shifts.

A call when it matters

Over the reporting year, Guardian Centre analysts made more than 200,000 calls to fleets following confirmed fatigue events, or around 560 a day. Many came outside ordinary working hours: 43% were made between 10 pm and 6 am in the vehicle's local time, and 22% on the weekend, with the busiest hour at 5 am*.

Around-the-clock review means a fatigue event at 3 am can reach the fleet while the driver is still on the road, not in the next morning's report.

What happens after the call

The data also shows what follows. When a Guardian Centre call was answered, the vehicle was about half as likely to record another confirmed fatigue event in the period that followed. Within two hours, 12.4% of answered calls were followed by a further fatigue event, compared with 22.5% of calls that went unanswered.

Why this matters

Guardian combines in-cab alerts and fleet-level information with human review and customer-specific escalation. When a fleet answers the call, a further fatigue event is much less common in the hours that follow. Other factors may also play a part, but the pattern points to the value of a timely response.

Accuracy you can act on

Analyst review also helps distinguish confirmed fatigue events from detections that do not meet the relevant classification criteria. This additional context can support fleet safety teams when reviewing event information, discussing circumstances with drivers and considering operational practices.

What this means for fleets

Around one in five calls go unanswered. Fleets can review whether their Fatigue Intervention Plan reflects their current operations, whether nominated contacts are up to date, and whether someone can take a call overnight and at weekends, when many calls are made.

Table 1: Another confirmed fatigue event on the same vehicle after the call

	Call answered	Call not answered
Within 30 minutes	7.5%	14.2%
Within 1 hour	9.6%	18.0%
Within 2 hours	12.4%	22.5%

* Call timing and follow-up figures are based on calls made by the Guardian Centre between 1 January and 30 September 2026 under a Fatigue Intervention Plan. Follow-up figures cover calls made after a vehicle's first confirmed fatigue event in at least four hours. A further event is any confirmed fatigue event on the same vehicle after the call.

Putting the insights to work

The findings in this report provide a broad view of global fatigue and distraction trends. Your fleet's operations and safety practices are unique. Comparing your Guardian data with these broader findings can help identify patterns worth exploring, inform decisions and provide context for assessing changes over time.

Investigate your own data

Guardian Live provides access to account- and fleet-specific driver and vehicle data, allowing you to explore recorded events in the context of your operations.

Looking beyond total event counts to consider normalised event rates, recurring times and days, and event types can help identify when and where fatigue or distraction occurs more frequently. These patterns are a starting point for further investigation, rather than conclusions in themselves.

Consider whether recorded trends align with operating schedules, journey lengths, rest practices, routes or other operational factors. Where available, analyst-reviewed events can provide additional context for individual detections.

Consider what the findings may mean for your operations

Where trends warrant further exploration, consider how existing safety measures and operational practices relate to the findings.

Areas for consideration include:

- **Operating practices:** Whether rosters, journey schedules, rest opportunities or other working conditions may be associated with recurring event patterns.
- **Safety initiatives:** How findings might inform fatigue management, distraction prevention, driver education or other safety measures.
- **Guardian Live features:** How event notifications, details and footage can support a better understanding of individual circumstances. Fleets can also elect to receive a phone call for a confirmed fatigue event, in accordance with their agreed Fatigue Intervention Plan.
- **Collaboration:** How fleet managers, safety teams and other stakeholders can consider findings alongside operational knowledge and existing risk management processes.

Combining Guardian insights with operational information can help fleets identify where further attention may be warranted.

Measure progress over time

Understanding safety performance involves more than tracking whether recorded events increase or decrease. Changes in event rates need to be considered alongside kilometres travelled, system coverage, fleet composition and changes to detection capabilities.

An increase or decrease in recorded events is a starting point for investigation, not a conclusion. Similarly, a reduction in event rates does not, on its own, establish that an intervention has been effective.

Guardian Live's daily and weekly digest reports and monthly insights report provide opportunities to review account- and fleet-level trends. Tracking relevant measures consistently alongside operational changes can help fleets understand how trends evolve and whether further investigation or adjustments are warranted.

Engage and support your drivers

Guardian insights may provide useful context for driver education and engagement. Sharing relevant trends can inform conversations about fatigue, distraction and safe driving practices.

Where appropriate, discussing individual events with drivers can also help establish the circumstances involved. Available event footage can support constructive conversations, reflection and driver education.

Approaching these conversations as opportunities for understanding and improvement, rather than focusing solely on individual events, can help build a shared commitment to safety.

Reinforce safe practices

Guardian's ability to detect events depends on the in-cab sensor being correctly positioned and able to see the driver. Misalignment or tampering can affect system performance.

Fleets can reinforce the importance of correct sensor positioning through driver communications and existing processes. Seeing Machines also provides resources to explain how Guardian works and address common questions or concerns.

For further information, contact your Account Manager or Distributor.

By exploring recorded trends, considering their operational context and reviewing changes over time, fleets can use Guardian insights as one source of information within their broader approach to managing driver risk and improving fleet safety.

Glossary

- **Circadian rhythms:** Natural internal processes that regulate the sleep–wake cycle and influence patterns of alertness across the day.
- **Confirmed fatigue event:** A microsleep or drowsiness event confirmed by a Guardian Centre analyst after reviewing the footage.
- **Drowsy episode:** A continuous period in which Guardian detects early drowsiness (an elevated drowsy state), counted from when the state begins.
- **Fatigue Intervention Plan:** An agreed, customer-specific process that guides the Guardian Centre in contacting nominated customer representatives following a confirmed fatigue event.
- **Generation 3:** The current generation of Guardian hardware. Early drowsiness and attention sharing figures in this report come from Generation 3 systems only.
- **Guardian Centre:** A team of trained analysts working 24/7 to review and classify high-risk events captured by Guardian and initiate Fatigue Intervention Plans where applicable.
- **Guardian Live:** A secure online platform that provides companies with access to driver and vehicle data and fleet insights.
- **Karolinska Sleepiness Scale (KSS):** A nine-point scale used in sleepiness research to assess levels of sleepiness, from fully alert to extremely sleepy and struggling to remain awake.
- **Operating hours:** The total time vehicles were operating during the reporting period, whether moving or stationary.
- **Trip:** A single vehicle ignition cycle. Trip statistics in this report cover trips of 10 minutes or more with early drowsiness detection running.

Help shape future reports

As we strive to improve and refine future Guardian Insights Reports, we invite you to share your feedback and suggestions. Your input will help ensure our reports remain informative and relevant to your interests and needs. Please take a moment to reach out and help us shape what matters most to you.


seeingmachines