

“AVs: Statement of Safety Principles”

techUK response

September 2026

Introduction and Summary

techUK is the UK's technology trade association, representing over 1,100 member companies. Our Self-Driving Vehicles Working Group brings together automated driving system developers, vehicle manufacturers, suppliers, software and connectivity providers, insurers and infrastructure operators.

We welcome the draft Statement of Safety Principles and support the three central choices the Department has made: keeping the SoSP tightly scoped, setting the safety standard at equivalent to a careful and competent human driver, and building the Principles from observable behaviours. Running through our response is a general support for the Principles as drafted, but with asks for:

- Further guidance – recognising that guidance will be vital in providing clarity and colour to the SoSP
- The Principles to be applied flexibly – making judgements on classes of behaviour in a way that is proportionate and contextual
- AVs – which will deliver significantly better safety outcomes overall – to be assessed fairly in the context of a careful and competent human driver
- International alignment to be protected

The purpose of the SoSP

Question 1: Do you agree or disagree with the purpose of the statement of safety principles being to provide guidance to the Secretary of State in considering the self-driving test?

AND

Question 2: Why do you think this? If relevant, what alternatives would you prefer?

techUK agrees that the Statement of Safety Principle (SoSP) should serve its statutory purpose of guiding the Secretary of State's assessment of the self-driving test at authorisation and of framing the annual assessment of fleet performance. A tightly-scoped SoSP gives a clear, stable benchmark against which to build safety cases.

We see no need to expand the uses of the SoSP where they are proposed, and certainly not without clear guidance and an assurance that it would not duplicate other regulatory processes. On the current proposals, the call for evidence highlighted a number of

alternative functions and we agree with the Department that these are better addressed elsewhere. They would be more properly located in authorisation conditions, the in-use regulatory scheme, NUIC operator licensing, and international type approval requirements. Duplicating them in the SoSP would create ambiguity and duplication.

We recommend that:

- If the SoSP is to have wider uses, this should be made very clear. This is especially true given the potential for the SoSP to be used in a post-authorisation enforcement manner. The final document should state explicitly how it will be used beyond authorisation, so that its role is predictable rather than discretionary.
- The relationship between the two statutory purposes be confirmed. The section 1 purpose is vehicle-specific, concerning the authorisation of a particular vehicle type or feature, whereas the section 38 purpose is sectoral and continuous, concerning the ongoing performance of the self-driving fleet. We would ask that the final statement confirm that a finding under the section 38 monitoring duty does not in and of itself put an existing authorisation in question. Explicit confirmation of this would give developers and investors useful certainty.
- The SoSP be located clearly in the authorisation process. By the time a developer seeks authorisation, it will already have prepared a structured safety case for type approval. It is not currently clear whether assessment against the SoSP draws on that safety case, supplements it, or requires separate evidence, nor at what point in the process the assessment happens. A greater understanding of the relationship between the SoSP and the wider authorisation process would materially reduce uncertainty.

On all these recommendations, further guidance would be of great benefit to the sector, and is likely to be the best mechanism to address these – and any other – concerns and clarificatory questions raised, rather than reflecting all of them in detail in the SoSP directly.

Setting the safety standard in the SoSP

Question 3: Do you agree or disagree with the safety standard for automated vehicles being set at equivalent to careful and competent human drivers?

AND

Question 4: Why do you think this? If relevant, what alternatives would you prefer?

We broadly agree. If a single standard is to be selected, we believe that this is the right one. There are points to be made in its favour:

- It is genuinely demanding, since “careful and competent” is higher than the *average* driver, who may be distracted intoxicated, etc.

- It aligns internationally with the ADS regulations adopting the same concept. This in turn lowers entry costs, enables cross-border deployment, and keeps Great Britain an attractive first market. Divergence now would be costly.
- No measurable alternative exists. There is no consensus on what "higher than careful and competent" would mean exactly.

However, we would qualify our support in the following ways:

- The comparator is not a workable benchmark. "Careful and competent" is borrowed from road traffic law, where it is a retrospective, case-by-case test applied to human conduct after the event. Being designed to test human failures, it is very difficult to apply properly as an engineering benchmark. It should be applied very flexibly, not as though a settled body of measurable criteria sits behind it. Future regulations will play a role in making up for some of this shortfall, for which there will certainly be a need.
- The issue of measurement remains unresolved by this consultation. The consultation says what the standard is, but not how or when compliance with it is demonstrated, what evidence is needed to do so, or where that evidence comes from (especially true given that no deployments at scale in Great Britain and the acknowledgement that trials will not generate statistically significant volumes in the short to medium term). It is not reasonable to compare one developer's performance directly against another's where operations differ materially.
 - We would support a gradual approach to establishing both baselines and measurement, and we would ask for further guidance as deployments scale.
- A careful and competent human driver does not drive perfectly. Where a vehicle performs in such a manner as a careful and competent human driver conceivably would, that should satisfy the standard. In wider practice, we expect AVs to perform to a significantly *higher* level of safety, of course.

We would also add a point on communication. Public acceptance of "careful and competent" increases markedly when it is made clear that it exceeds the average human driver. The final SoSP and supporting communications should state this prominently.

Using behaviours within the SoSP

Question 5: Do you agree or disagree with using behaviours as the basis for developing the principles within the statement of safety principles?

AND

Question 6: Why do you think this? If relevant, what alternatives would you prefer?

We agree in principle. Behaviours are broadly observable, testable and explainable, and can also be evidenced through simulation, track and real-world testing. They correspond to

categories long-used to assess human drivers and align with the behavioural competency work of other parts of government in relation to road use.

However, we would attach these implementation safeguards:

- Behavioural assessment should be technology-neutral and proportionate to the ODD, focussing on what the vehicle does rather than how it is engineered.
- Behaviours should map onto the structured safety case required by the UNECE ADS regulations, so developers evidence each behaviour once rather than duplicating material across regimes.

Identifying careful and competent behaviours

Question 7: Do you agree or disagree with the identified behaviours of vehicle control, hazard anticipation, perception and reaction, obeying traffic rules, interaction with other road users, and adaptability to road conditions?

AND

Question 8: Why do you think this?

We agree with this. The five behaviour themes broadly describe the dynamic driving task. They also correspond to categories long-used to assess human drivers, which supports public understanding and regulatory continuity. However, no exhaustive definition of “careful and competent” driving is possible, because the number of possible driving scenarios is infinite. Therefore, these behavioural themes are best understood as lenses through which to analyse vehicle behaviour, rather than particular benchmarks that may or may not be achieved.

Question 9: In your view, what, if any, other behaviours do you think are missing and why?

techUK does not recommend adding to the number of behaviours, and we believe that the behaviours listed sufficiently cover all relevant aspects of driving. Any gaps that might be claimed are likely filled by strengthening the relevant Principle.

Overall improvement in road safety

Question 10: Do you agree or disagree that the proposed principles on improving road safety for all and interacting safely with other road users provide sufficient protection for all road users?

AND

Question 11: Why do you think this? If relevant, what alternatives would you prefer?

We agree that the proposed principles create a reasonable protection for all road users. Additionally, we support the equalities-led framing.

Two implementation points will determine whether the protection is tangible.

- First, at early fleet scale, group-level statistical conclusions will be fragile. The consultation recognises this in relation to its wider rejection of numerical targets. Individual incidents should not necessarily be read as group-level signals.
- Second, any assessment scenarios should be co-designed with, industry, industry representatives, and road safety organisations. This would both improve the scenarios and support public confidence.

However, the framing understates the benefit. Both questions ask only whether any group is disadvantaged but, for several groups, self-driving vehicles offer material improvement. We and others have raised these benefits in previous responses and would highlight them again here.

The draft safety principles

Question 12: Do you agree or disagree with the inclusion of Principle 1: self-driving vehicles should use the road in accordance with traffic laws and the Highway Code?

AND

Question 22: Do you have any comments, including textual amendments, on Principle 1?

We agree with its inclusion, with the following comments:

- The Highway Code mixes legal requirements with advisory guidance. The SoSP or its supporting assessment guidance should assess adherence to Principle 1 in light of this. Additionally, advisory provisions should be assessed as context-dependent, as they are for human drivers, since rigid adherence to advisory rules can itself create additional risk.

- "Temporary restrictions" should be read widely, to encompass collisions, obstructions, the directions of police officers and other authorised officials, temporary speed limits, emergency closures, and vulnerable road users in unexpected positions. Further guidance for the sector on this point would be of great benefit.
- Compliance in practice depends on supplying vehicles with accurate, current information about restrictions. We encourage the Department to accelerate the digitisation of traffic regulation orders, to work with local authorities on signage and roadworks data quality, and to ensure officials directing traffic are trained to give consistent signals.

Question 13: Do you agree or disagree with the inclusion of Principle 2: self-driving vehicles should maintain control over the vehicle at all times?

AND

Question 23: Do you have any comments, including textual amendments, on Principle 2?

We support its inclusion, with the following suggested clarifications:

- Smoothness should not penalise decisive action. For example, emergency braking or assertive avoidance can be abrupt, entirely correct, and within the "careful and competent" standard.
- "At all times" has the potential to disregard the transition period for user-in-charge features and would benefit from being stated in the final text to avoid ambiguity.

Question 14: Do you agree or disagree with the inclusion of Principle 3: self-driving vehicles should predict, detect and respond to hazards proactively?

AND

Question 24: Do you have any comments, including textual amendments, on Principle 3?

We agree with Principle 3's inclusion in principle, but we are concerned that the word "proactively" may produce excessive conservatism, leading to unintended consequences for other traffic. In engineering practice, a requirement to be proactive is often observed by acting earlier and with greater margin, meaning earlier and sharper braking and a resulting higher variation in speed and trajectory. The result is an AV that disrupts other traffic, hesitates where a careful and competent driver would not, and whose caution itself creates risk by unduly disrupting other traffic. It is also unclear how proactivity would be measured, and likely to be read differently by different assessors.

We would, therefore, recommend the deletion of the word “proactively”, and believe that the remaining formulation of “predict, detect and respond to hazards” adequately describes what an AV should do. We also suggest that “contextually appropriate” be replaced by “proportionate”, which carries an established meaning in regulatory assessment.

Finally, it is important that “prediction” does not envision a level of prediction beyond that available to a careful and competent driver. Guidance on the sort of environmental cues or scenarios that an AV should be expected to recognise would be welcome.

Question 15: Do you agree or disagree with the inclusion of Principle 4: self-driving vehicles should adapt driving to road and weather conditions?

AND

Question 25: Do you have any comments, including textual amendments, on Principle 4?

We agree in principle with its inclusion, but note that adaptation has two distinct modes, which the text of the principle should account for:

- Adjusting driving strategy within the conditions that the vehicle is designed for; and
- Recognising that conditions have exceeded the ODD and responding safely, which may mean a minimum risk manoeuvre rather than continued adaptation.

This second point is particularly relevant for the comparison with careful and competent human drivers, who very often do not recognise hard-to-spot changes in condition, such as black ice or heat-related surface degradation. This Principle should recognise this shortcoming in human driving and not hold AVs to a higher standard.

Question 16: Do you agree or disagree with the inclusion of Principle 5: self-driving vehicles should behave in a predictable manner?

AND

Question 26: Do you have any comments, including textual amendments, on Principle 5?

Predictability is important to ensure safe mixed-traffic operation and to secure public confidence. Other road users should be able to read a vehicle's intentions from its signals, positioning and speed, at least to the same extent as this is possible by observing a careful and competent human driver, and it is important that this is reflected in the Principle. We support the Principle's inclusion, with the following comments:

- Predictability is best assessed against the signalling and positioning conventions of the Highway Code – which Principle 1 already requires – rather than against an open-ended expectation.
- A requirement for predictability must also allow for decisive defensive action or adaptation to sudden hazards (as outlined in other Principles), as these can be abrupt, quick to develop, and require adaptive behaviour. These may not be *per se* predictable, but are necessary for safety and in-keeping with the standards of a careful and competent driver.
- For modern automated driving systems as for human driving, completely identical outputs in identical scenarios are neither expected nor possible, and so should not be considered a meaningful test. Therefore, assessment should judge predictability as being within a consistent class of behaviour, rather than as being specific and exact actions.
- On "communicating its intention through appropriate signals", we suggest that guidance confirm that intent is communicated through conventional signalling, road positioning and speed profile, and that no specific external human-machine interface technology be mandated.
- Other road users include other self-driving vehicles, and the Principle should not differentiate between the two. This Principle addresses outward communication, while Principle 8 addresses inbound signals, and the two should cross-refer.

Question 17: Do you agree or disagree with the inclusion of Principle 6: self-driving vehicles should improve road safety in Great Britain for all?

AND

Question 27: Do you have any comments, including textual amendments, on Principle 6?

We support the general framing of this Principle, along with the ambition, which is a key reason to support the introduction of AVs. We have, however, comments on the text:

- "Avoiding any disproportionate increase in risk to any group of road users" should be assessed at fleet- and system-level over reasonable deployment times, rather than inferred from individual events.
 - At early deployment scale, single incidents can be incorrectly read as group-level signals, and the consultation itself recognises this. Any assessment methodology should be published.
 - Furthermore, not every incident is the fault of all vehicles involved. The SoSP and other official practice must be clear as to what is – and is not – to be recorded as an AV safety failure.
- We welcome the recognition that, where safety risks occur, the vehicle should seek to minimise the potential for harm. This correctly avoids demanding the impossible standard of zero risk.

- Assessment must account for how and where vehicles are deployed. Different deployments in different ODDs by necessity carry different risks and cannot be directly compared. Assessment of road safety improvements must be applied contextually and understood holistically, across the whole ecosystem and across a reasonable time period.

Question 18: Do you agree or disagree with the inclusion of Principle 7: self-driving vehicles should interact safely with all road users, particularly those most at risk?

AND

Question 28: Do you have any comments, including textual amendments, on Principle 7?

We support this Principle's inclusion and note that interaction with vulnerable road users is precisely where automated systems can exceed human performance, with continuous all-round perception and no distraction or fatigue. However, we wish to make the following points:

- The word "safely" should not appear in the Principle. This is one of two Principles in which the word "safe" appears in the headline text, whereas the others are expressed in behavioural terms. In a statement whose statutory purpose is to define what safe driving means, defining a principle by reference to the word "safe" risks circular reasoning, giving assessors and developers less to work with. We recommend it be recast behaviourally.
- "Being considerate" is particularly vague and could do with operational meaning being set out in assessment guidance. As drafted, it is open to wide interpretation, and consideration cannot be evidenced without a shared understanding of what it requires.

Question 19: Do you agree or disagree with the inclusion of Principle 8: self-driving vehicles should interact safely with emergency services?

AND

Question 29: Do you have any comments, including textual amendments, on Principle 8?

We support its inclusion, with the following points:

- This is the other Principle with the word "safe" in the headline text. We reiterate our concerns about this, on the same teams as for Principle 7 (above).
- "Other authorised officials" should be defined in guidance so developers know the set of people and organisations that they must recognise and respond to. This will be

performed in reference to Highway Code Rule 219 and the expected behaviour of careful and competent human drivers.

- This obligation is necessarily two-sided: vehicles can only respond reliably to reasonably consistent human behaviour. We, therefore, encourage the Department to work with other stakeholders on consistent stop and direction protocols, including the standardised hand signals on which traffic officers are already trained.
 - It is also important that, as with other Principles, bad actors may attempt to deceive a vehicle into action by impersonating an official. Given that careful and competent human drivers can also fall victim to this, AVs should be given the same leeway.
- The Principle should also confirm that responding to instructions can include circumstances where lawful directions conflict with default traffic rules, and that this would not mean a violation of Principle 1. There could also be scope for a defined fallback scenario in the face of conflicting or unverifiable instructions.

Question 20: Do you agree or disagree with the inclusion of Principle 9: self driving vehicles should only operate autonomously within the limits under which they are specifically designed to function?

AND

Question 30: Do you have any comments, including textual amendments, on Principle 9?

The ODD is the foundational safety construct for automated driving, and this Principle correctly reflects it, in line with the approach of the UNECE ADS regulation. We agree with this Principle's inclusion, with the following comments:

- The explanatory text should distinguish between two different situations that both require a safe response: a fault or degradation arising within the ODD on the one hand, and the vehicle approaching or leaving the conditions for which it is designed on the other. These are different failure modes calling for different strategies and evidence.
- "Appropriate strategies" should remain technology-neutral and proportionate to ODD complexity. If every extension or software update requires assessment from first principles, iterative improvement becomes prohibitive and safety improvements reach the road more slowly. Additionally, this should make clear that 'appropriate strategies' refers to a vehicle reverting to minimum risk conditions.

Question 21: Do you agree or disagree with the inclusion of Principle 10: self-driving vehicles should be designed to take account of the specificities of the territory where they will operate?

AND

Question 31: Do you have any comments, including textual amendments, on Principle 10?

We support the objective of this Principle, but we would like some clarity as to the separation between it and Principle 9. We would understand Principle 10 as relating principally to the evidence base for aggregate safety, and therefore as supporting the monitoring duty under section 38 by enabling meaningful comparison between automated vehicles and human drivers. On that reading it performs a different function from Principle 9, which concerns the performance of an individual vehicle within its designed limits. We would welcome the final statement making that distinction, and the relationship between the two Principles, explicit.

We would also ask from clarity on what is meant by "territory". The term is capable of being read as the national territory of Great Britain or as the specific geographic area in which a vehicle is proposed to operate. We consider the latter reading the more workable, since it aligns more closely with the ODD against which a vehicle is in any event validated, particularly given that driving in different areas of Great Britain presents materially different requirements. If, on the other hand, this is to be read as national territory, the Principle risks privileging GB-collected data over data from comparable operating environments elsewhere which is, in fact, often more representative of the intended area of operation. We would encourage government to define "territory" here as the defined geographic area of operation.

It follows that international evidence must be admissible where transferability, similarity and relevance are demonstrated. A requirement for exclusively GB-collected datasets would depart from the UNECE approach, raise entry costs without a safety gain, disadvantage UK-based developers seeking reciprocal treatment abroad, and delay the road safety benefits that deployment brings.

Principles that we haven't included

Question 32: In your view, should any currently excluded areas be included in the statement of safety principles?

AND

Question 33: Which areas should be included? Why do you think they should be included?

We support the drafted scope of the SoSP and do not consider creating additional principles to be a priority. More wide-ranging principles inflation risks diluting the statement's statutory function and its public legibility, and the refinements we suggest elsewhere in this response sit within the existing themes. Additionally, it is important that public concerns – such as cyber resilience – must not be inferred, or capable of being presented, that because these sit outside the SoSP they are considered unimportant or left unprotected. We would, therefore, ask the Department – as part of its public communication – to publish a public-facing single issue map, setting out the instrument by which each major concern and issue is covered.

Costs and benefits

Question 34: In your view, what costs should be taken into consideration when assessing the impact of the SoSP?

techUK and its members support a demanding safety standard, and the one proposed is demanding. The costs worth the Department's attention are those incurred without producing safety benefits. These include:

- Duplication – if evidence demanded under the SoSP is beyond or parallel to the structured safety case required by the UNECE ADS regulation, developers will bear double compliance costs for no safety gain.
- Divergence – any departure from international alignment raises market entry costs, slows authorisation and diminishes the UK's attractiveness as a first market. This applies to apparently minor differences in wording as much as to substantive divergence, since each requires a separate evidence stream.

Question 35: In your view, what benefits should be taken into consideration when assessing the impact of the SoSP?

The principal benefit is earlier and safer deployment of a technology with substantial casualty reduction potential. A clear, stable, internationally-aligned SoSP delivers investment certainty, lowers entry costs, supports the UK's position as a leading market, and underpins the growth and employment benefits identified for connected and automated mobility. A publicly-legible statement also builds the public trust on which adoption – and therefore every other benefit – depends, particularly where communication makes clear that “careful and competent” exceeds the average driver. We would also ask the Department to reflect the network level benefits of automated operation, even though they are harder to quantify.

Further benefits include improved mobility and independence for those who cannot drive, clarity for the insurance market and, through international alignment, the exportability of UK-validated systems and expertise.

Closing comments

Question 36: do you have any other comments?

The principles risk assuming a static system. In practice, automated driving systems develop continuously, capability improves, software is updated, and ODDs are extended. The consultation acknowledges the SoSP can be amended but ask that the accommodation of change be explicit. The statement should confirm the principles apply to a system that will develop over time and should say how and how often it will be reviewed, what the amendment process is, how much notice will be given, and how transitional arrangements work for vehicles already authorised. Assessment guidance for all Principles will also be of great importance and help to the sector.

ENDS