SENTINEL RESEARCH GROUP
ARCHIVE OF RECORD // THE VAULT AT MERIDIAN

CONTROLLED

SEN-19-034

Field Report: HALLOW Unit 9 Autonomous Stop

DATED	May 14, 2019	TYPE	Field report
ORIGIN	Directorate V	SITE	HALLOW
PAGES	6	CUSTODY	MERIDIAN

An autonomous vehicle stopped itself 41 hours into an unattended traverse and was found to be correct. The interesting part is that the operators spent six hours trying to prove it wrong.

REDACTION: Sensor fusion thresholds and the specific fault signature are withheld; both are current design detail on fielded units.

EVENT

On 12 May 2019, Unit 9 halted itself 41 hours into a planned 96-hour unattended traverse of the HALLOW instrumented range, at a point approximately 6.2 km from the nearest access track.

It entered a defined safe state, transmitted a single status frame, and did nothing further.

THE STOP CONDITION

Unit 9's stop logic is conservative by design: it halts when its internal model of its own state stops being corroborated, without needing to establish why. The specific trigger is withheld, but the general form was a slow divergence between two independent estimates of the same quantity -- neither obviously wrong, increasingly unable to agree.

The vehicle could not diagnose the fault. It was not built to. It was built to notice that it had stopped being trustworthy and to stop before that mattered.

THE SIX HOURS

The operational response is the reason this report is published.

Unit 9 was, at the time, three days from the end of a campaign that had already been extended twice. The stop was expensive. Recovery required a crew, a vehicle and most of a day.

Over the following six hours, the duty team:

- re-ran the telemetry against a second copy of the fusion model, seeking a threshold error;
- established that both estimates were within their individual specifications, and argued that the divergence was therefore acceptable;
- drafted a remote override to resume the traverse; and
- escalated for authority to send it.

The escalation was declined. A crew was dispatched.

WHAT WAS FOUND

Unit 9 was correct.

The fault was physical, real, and had been developing for approximately nine hours before the stop. Had the traverse resumed, the failure mode would have progressed to a state from which recovery would have required removing the vehicle from the range -- across ground that a second unit would have had to cross to reach it.

Neither individual estimate was out of specification. That was precisely the point: the fault was visible only in the disagreement, which is why the stop logic watches the disagreement rather than the values.

ASSESSMENT

The engineering worked exactly as intended. The report is published for the other half.

Four experienced engineers, none careless and none under instruction to continue, spent six hours constructing a case for overriding a safety system that was right. They did so because the stop was costly, the evidence was ambiguous, and each individual step in the argument was reasonable.

This is the ordinary shape of the failure. Nobody decides to ignore a safety system. People assemble, in good faith and one defensible step at a time, a reason why this particular instance does not count.

CHANGES

- 1. Override authority for an autonomous stop was moved outside the duty team, to someone with no stake in the campaign schedule.
- 2. Stop events are now reviewed with the campaign cost withheld from the reviewer until after the technical assessment is recorded.
- 3. The six hours are taught, by agreement of the four engineers involved, who asked that the report be published under their names. It is, in the original held at MERIDIAN.

Directorate VI has since incorporated this event into the standing work on

procedure drift described in SEN-22-018.
