

IREC 2025 Post-Competition DFMS Analysis

Carl Anderson

October 16th, 2026

Purpose

The purpose of this document is to pick up where the 2024 post-competition analysis left off. It repeats that work on the data from the 2025 IREC Digital Flight Management System (DFMS), and adds a few things the 2024 data couldn't tell us, along with some initial conclusions & possible prescriptions. In 2024 the analysis ran on the Trello action history. The 2025 DFMS is a database, and this time the data is a full dump of it: 4,058 timestamped status changes across 145 flight records, plus the flight records, safety issues, motors, transmitters, tracker positions and the weather station feed. That is enough to answer most of the 2024 questions again, and as with last year, it is only as good as what was entered into the DFMS.

Executive Summary

Broadly, 2025 looks a lot like 2024 on paper, but the constraint has moved. The DFMS recorded 139 launches this year against 102 in the 2024 DFMS (118 by the LCO desk count), which already beats the roughly 125 teams the 2024 analysis projected for a four day launch window. Registration & FSR were again not the problem. Every team was checked in by 4PM Monday and 80% had completed FSR by the end of Tuesday (77% last year), with FRR not starting in earnest until Wednesday, when 113 teams were sitting in the Ready for Final Safety status waiting for it (fig. 1). FRR was the bottleneck in 2024, and this year it was a bottleneck in a different way. The reviewers had room: teams were paged in within a minute of being called, the average review took 8.2 minutes (10.1 last year), only 1.2 reviews were open at a time, and the FRR waitlist never held more than 2 teams (fig. 6). What held FRR back was how fast teams were called. During the launch hours it completed about 3.8 teams an hour, and the 20 team salvos I think we should be aiming for need about 7.

That mattered because there was a lot of time with nothing leaving the rail. On 11–13 June there were 14 stretches of more than 30 minutes between launches, about 17 hours in all (fig. 7). Launches came in salvos (ten bursts of five or more launches held half of all launches, the biggest being 14 launches in 30 minutes) and a typical salvo went out about two hours after the last one. I looked at what limited the size of those salvos, and it was mostly the supply of teams. A full 20 team salvo was staged, meaning through FRR and waiting in Quarantine or Salvo, at only 5 of the 15 dispatches of six or more teams (fig. 9), and the pool was 20 or more for just 9% of launch hours. At the other 10 dispatches the range sent about 85% of what was staged. FRR opened 2.2 hours before the first large salvo on 11 June, when 113 teams had been ready since Tuesday night, and only on 12 June did it open the 3 hours the 2024 report asked for, which was also the day the pool filled (table 3). At the 5 dispatches that did have 20 or more staged the range sent 11 to 16 and left 61 teams behind. I found no sign that tracker frequency conflicts, or a need for hybrid or two-stage pads, decided who was left behind. The pad rows fill front to back and the back rows are backfilled when a salvo is large (fig. 10), so I can't show that pads capped those salvos, and all 61 teams left behind were still in Quarantine, so they had never been called into the salvo. The decision was made at the call, and the data doesn't say why it stopped at 11 to 16. Of the 289 times a team was moved from Salvo back to Quarantine, 163 happened within a minute of the call and every one of those was entered by the same staff id, so I read them as the DFMS being used to ping teams and not as real recalls. The other 84 touched 44 teams, and I couldn't tie them to wind, frequency conflicts or avionics issues, so I can't say why they happened.

That led to what I think is the most useful result for next year, which is salvo size. Looking at the 15 dispatches of six or more teams, the time until the next salvo barely depended on how many teams were in it (about 106 minutes plus 2 minutes per team), and roughly a quarter of dispatched teams bounced back before launching whatever the size. I built a simulation on those numbers, resampling the 2025 salvos 800 times, and in it the best salvo is the 20 pad limit in 73% of runs and 17 teams or more in 99%. Salvos of 14, 17 and 20 would have cleared all 139 teams in about 24, 21 and 19 range-hours instead of the 32 we used. That assumes the teams are there to fill the salvo, and in 2025 they often weren't. Dispatching everything that was staged, up to 20, would have averaged 13 teams a salvo instead of 10, which the simulation puts at about 26 range-hours. A full 20 team salvo every time also needs the uprange to have 20 staged before each dispatch. The result only flips if bounce rises by more than about 2 points for every extra team, and the upper bound the data allows is 1.6. Nothing bigger than 16 was ever dispatched though, so anything above that is extrapolation and I would want it tried in steps instead of taken on faith.

On the review side, about half of teams (69 of 143) were sent back from FSR at least once, and each send-back took a median of 5.3 hours to come back around. The number of safety issues a team got tracks send-backs closely, which is how issues cost teams time, but it doesn't predict how the flight went: teams whose boost failed had 2.00 issues on average against 1.87 for nominal boosts. FSR reviewers also differ a lot in how often they send teams back, from 10% to 71% among the 12 busiest, which is far more than chance would give ($p = 0.002$ after accounting for the day, review type & class). Whether the stricter reviewers are catching real defects I can't tell from this data, but it is the sort of thing a shared rubric would help with.

Three quarters of flights had a nominal boost, and 43% had a nominal boost and recovery. Recovery was the weaker half: of the 95 nominal boosts with a known recovery, 13 failed and 22 more deployed main at apogee. The 30k COTS class did worst (52% hard failures against 27% elsewhere), but it flies almost entirely on O class motors and boost failures climb with motor size, from 14% on L to 32% on O, so I think motor size is the likelier explanation than the class. On the radio side, 141 teams registered 425 transmitters, 54% of them in the 33 cm band, and about 30 teams have declared ranges that overlap at 915 MHz. The other bands are not crowded across their width: 70 cm and 5.8 GHz each have one hot spot (12 and 13 teams) and 2 m, FRS and 2.4 GHz are nearly empty. 19 of the 53 teams in the live video challenge had no video transmitter registered at all.

As in 2024 the data has its problems. One test record had to be thrown out, the status log never says why a team was recalled, and the tracker log only covers 24 of the 139 flights. I believe useful conclusions can still be drawn, but they are best read as broad trends and not as microscoping any one team. I have the following prescriptions for next year:

- **Dispatch everything that is staged, up to 20, in every salvo.** At the 5 dispatches with 20 or more teams staged the range sent 11 to 16 and left 61 behind, and all of them were still in Quarantine, so the limit was in how many teams were called. Pad rows don't explain it, since the back rows were backfilled when a salvo was large. Simulated salvos of 14, 17 and 20 would have saved about 8, 11 and 13 range-hours, and sending everything staged (about 13 a salvo) about 6. I would step up through 14, 17 and 20 and log dispatch size & bounce, so the next report can measure the effect instead of simulating it.
- **Stage ahead of the range.** Keep the 2024 prescription of opening FRR at least 3 hours before the first salvo, which happened on 12 June and not on 11 June (2.2 hours, with 7 teams staged at the first large salvo) or 14 June, and then call teams against a target and not on demand: aim to have 20 teams through FRR and waiting before the previous salvo's last launch. 13 June had only 1.1 hours but started from 18 teams already staged overnight. The 2024 prescription of doubling the reviewers is not what was missing this time, since the reviewers had room (1.2 reviews open on average, 4 at most, 113 teams waiting on Tuesday night). FRR completed 3.8 teams an hour in the launch hours against the roughly 7 an hour that 20 team salvos need.
- **Instrument the DFMS.** Stop using call-and-undo to signal teams, log a reason for each real recall, keep a history of pad assignments and not just the final pad, and log the positions from every tracker (24 of 139 flights were covered this year).
- **Make the safety review process more efficient and more thorough, starting at the video safety review.** Efficiency: missing checklists, vent holes, CP/CG marks and frequency labels made up 66 issues (24%) and only 3 were critical, so they are things a reviewer can check on video before the event and keep out of the FSR send-back loop. Thoroughness: recovery hardware is the largest issue category (28%) and recovery is where nominal boosts most often went wrong, so asking for the ground or ejection test record (seven issues already asked for one) is a cheap check to see if it moves recovery outcomes. Consistency: FSR send-back rates ran from 10% to 71% among the 12 busiest reviewers, and a shared rubric and the occasional paired review would hold teams to the same bar. The video safety review comes before the event and isn't in the DFMS data, so I can't say how much of this it already does.
- **Don't build salvos around frequency crowding, and register the video downlinks.** The declared plan is crowded in the 33 cm band: about 30 teams overlap at 915 MHz, and 39 transmitters from 26 teams are declared at exactly 915.000 MHz. I found no sign that frequency conflicts changed recalls or who was dispatched, so I would not build salvos around them. What is still cheap is requiring live video downlinks to be registered, since 19 of the 53 teams in the live video challenge had none.

Analysis

Stages

Figure 1 is a burndown chart in the same format as last year, where each line represents the number of teams left to complete a major step: check-in, FSR, FRR, pad queue and flight. Check-in replaces registration, and is the first time a team shows up as Ready for Safety. Read the same way as before, one stage is bottlenecking the next when their lines sit on top of each other. The FSR line comes down over Monday and Tuesday and flattens at 31 teams on Tuesday night, so FSR was essentially done before launching began. The FRR, pad queue and flight lines don't move until Wednesday, and from there they come down together. That looks like the 2024 signature for FRR holding everything else back, and in a sense it is, since nothing downstream moved until FRR did. The difference is why. The gap between the FSR and FRR lines on Tuesday night is 113 teams who had finished FSR and were waiting to be called. They weren't waiting on reviewers: teams were paged in under a minute, the FRR waitlist never held more than 2 teams, and reviews ran about 1.2 at a time. FRR ran at the pace it was called, and that pace decided how many teams were staged in front of the range (figs. 7 and 9).

Figure 2 is the equivalent of the 2024 fig. 2, the average number of minutes between consecutive teams completing each step, with the same trimming (gaps above $\mu + 3\sigma$ are dropped to get rid of the overnight waits). The grey bars are last year's values where the step has a clear match, and the salvo stages are new, because the 2024 board had no equivalent. Check-in is identical at 2.4 minutes. The conference day steps are slower this year, around 19 to 20 minutes between teams against 8 to 16 last year, which I believe is




Figure 1: Major Launch Operation Stages Burndown

mostly my method: each team’s time for a step is its last pass through it, so teams that were sent back have their times pushed later and the gaps spread out. On launch days FRR completions came every 16.0 minutes against 18.5 last year, and launches every 14.2 minutes against 13.2. That puts FRR output about 13% slower than launching, where in 2024 it was 40% slower. Table 1 puts the two years side by side.




Figure 2: Average Time Between Teams Completing Stages (grey: 2024 report)

Table 1: Stage Measures, 2024 and 2025

Measure	2024	2025
Mean minutes between teams completing registration (2024) or check-in (2025)	2.4	2.4
Teams through FSR by the end of Tuesday	77%	80%
Mean minutes between FRR completions	18.5	16.0
Median minutes between FRR completions	6.27	7.2
Mean minutes in FRR review	10.1	8.2
Median minutes in FRR review	5.26	6.2
Mean minutes between launches	13.2	14.2
Flights recorded in the DFMS	102	139
Flights counted by the LCO desk	118	not in the dump

Review: FSR and FRR

Figure 3 gives the same view as last year’s fig. 3 of how long teams spent in the queues and in the reviews themselves. The top row is FSR: the total time a team waited in the queue (median 109 minutes, with 15 teams over 1000 minutes because their review ran into the next day), the review itself (mean 31 minutes), and the time a team spent fixing things after a send-back (median 5.3 hours, though the mean is 16.8 because some teams took days). The bottom row starts with the gap between finishing FSR and starting FRR, and the tall bar around 45 hours is the Tuesday night to Wednesday wait before FRR began. FRR itself was quick. Teams were paged in about six seconds after being called and the review took a mean of 8.2 minutes (median 6.2), and the median completion gap of 7.2 minutes means roughly a minute of waste per team, about the same as the 5.26 against 6.27 minutes I saw in 2024. Reviews over two hours, as before, or under thirty seconds are dropped as data entry errors. Table 2 has the numbers.




Figure 3: FSR & FRR Statistics

Table 2: Time Spent in Each Stage

Stage	Mean	Median	90th percentile
FSR queue wait per team (minutes)	201	109	892
FSR review (minutes)	30.8	32.3	61.3
Remediation after a send-back (hours)	16.8	5.3	47.5
FSR complete to FRR start (hours)	40.6	45.5	66.7
FRR queue and page (minutes)	0.1	0.1	0.2
FRR review (minutes)	8.2	6.2	16.6
FRR complete to launch (hours)	8.4	3.7	22.4
Final pad queue to launch (minutes)	184	84	212

FSR is where the review time goes, and send-backs are the reason. 69 of the 143 teams that entered FSR (48%) were sent back at least once, 84 times in total (57 teams once, 9 twice and 3 three times). Figure 4 is the flow diagram for the 2025 gates, the equivalent of the 2024 Sankey for the remediation loop. Teams run left to right from check-in through how FSR went, how FRR went and whether they launched, and each ribbon keeps the colour of its team’s FSR path so the sent-back teams can be followed to the end. 74 teams (51%) passed FSR with no send-back and 68 (47%) completed it after at least one. Two never completed it, and one of those had been sent back, which is why the 69 above is one more than the 68 in the diagram. Nearly everyone cleared FRR without a return (124, 86%). 17 teams were returned to an earlier stage after FSR was finished: 8 from Ready for Final Safety back to Ready for Safety, 3 from FRR itself (4 times), and 6 from Quarantine back to Ready for Final Safety. Seven of the 17 had also been sent back at FSR, and 16 of them launched. In all 139 of the 144 teams launched (97%). The other 5 were 2 that stopped in FSR, 2 that never cleared FRR and 1 that cleared FRR and didn’t fly. 85% of the flown teams with a critical safety issue were sent back at least once, against 70% of teams with only non-critical issues and 4% of teams with none, so I believe the issues are how a team ends up in that loop.




Figure 4: Team Flow Through Review Gates

Launch day

Figure 5 is similar to the last two and shows how long teams waited after FRR to get into the pad queue, and then how long from the pad queue to launching. The top plots have a long tail, mostly teams that cleared FRR late in the day and flew the next, and the bottom two are zoomed in under 6 hours and under 4 hours. The median team waited 2.0 hours after FRR to reach the pad queue and then 84 minutes to launch (mean 184), and 57 teams moved backwards at least once within the pad sequence, for example from racking back to transport. Overall the median team waited 3.7 hours from clearing FRR to launching, and 74% launched the same day they cleared.




Figure 5: Pad Queue Statistics

Figure 6 is the queue congestion chart. It shows how many teams were in each state over the whole event, and the thing to take from it is that the wait on launch days was after FRR and not in it. Quarantine and salvo together held up to 31 teams, the pad process up to 25, and the FRR waitlist & review almost nothing (2 at the most). The big grey block is Ready for Final Safety, the 113 teams that finished FSR and waited to be called. I am not sure why, but about a dozen teams also sit in the In FSR status for days at a time. I suspect those are cards left in place while teams worked through remediation, and it is the reason I drop reviews over two hours.

On launch days the question is what the range was waiting for. Figure 7 lays out each launch day. The top strip of each pair is the launches, with the ones that belong to a salvo (five or more with no gap over 10 minutes) in blue and gaps of more than 30 minutes shaded. Below it is the number of teams staged for




Figure 6: Queue Congestion

the next salvo, meaning through FRR and sitting in Quarantine, Salvo or Paged Salvo, and the number already on the pads, against a line at 20 teams for scale. The 14 gaps on 11–13 June add up to about 17 hours, 60 to 68% of each day’s launch window (table 4), which is an upper bound on hold time because it includes the staging between salvos, so it can’t be compared with the 10–30% weather hold assumed in the 2024 projection. In the middle of those gaps a median of 12 teams were already on the pads, so a good part of each gap is a salvo being set up and not nothing happening. What the gaps show is how small the pool behind that salvo was. The median staged pool at the middle of a gap was 6 teams, it was under 10 in 11 of the 14 and 20 or more in only 1. Over all the launch hours (07:00 to 19:00, 11–13 June) it was 5 or more teams 67% of the time, 10 or more 34%, 14 or more 26% and 20 or more just 9%, with a median of 7 and a mean of 9.2. The pool only reached 20 on the morning of 12 June, when 31 teams were marked through FRR in a single hour and it peaked at 31, and for a short while at the start of 13 June. The pads held a mean of 9.8 teams.

So the pool, not the reviewers, was short, and it was short because FRR was called slowly. Over the active hours of the launch days FRR completed a mean of 3.8 teams an hour (median 3, 90th percentile 5), and sustaining a 20 team salvo every cycle takes about 7. Launches came in salvos. Ten bursts of five or more launches, with no gap over 10 minutes, hold 68 of the 139 launches (49%), the largest being 14 in 30 minutes on 12 June. Launches from the same pad were a median of 173 minutes apart and never closer than 21 minutes. Fixed solid pads in rows A to D took 123 flights (row A 63, B 34, C 22, D 4), the two stage pads 9 and the hybrid pads 7. Class doesn’t explain the skew, since every class flew most from row A and least from row D. I read it as rows being filled front to back, with the back rows kept for other flights and backfilled when needed, and the salvo size section checks that.

The 2024 report asked for FRR to open at least 3 hours before the first salvo. The waiver opening times aren’t in the data, so table 3 uses the first dispatch of six or more teams each day as the first salvo. The 3 hours was met on 12 June, when FRR began at 04:48 and the pool at that salvo was 31 teams, the largest of the event. It wasn’t met on 11 June. 113 teams had been ready since Tuesday night, but FRR began at 07:32, 2.2 hours ahead, and the first large salvo had only 7 teams staged. 13 June had 1.1 hours but started from 18 teams already staged overnight, and on 14 June only 8 teams went through FRR at all.

Table 3: FRR Lead Time Before the First Large Salvo

Launch day	FRR reviews began	First salvo of six or more	Lead (hours)	Teams staged at that salvo
11 June	07:32	09:44	2.2	7
12 June	04:48	08:35	3.8	31
13 June	04:48	05:55	1.1	25
14 June	06:26	07:33	1.1	6



Figure 7: Launches, Staged Teams and Teams on the Pads by Launch Day

Table 4: Launch Days

Launch day	Launches	First to last launch	Window (hours)	Idle spells >30 min (hours)	Real recalls
11 June	33	10:30–19:04	8.6	5.8	9
12 June	47	08:15–17:26	9.2	5.5	22
13 June	45	06:56–16:07	9.2	5.7	49
14 June	14	06:46–11:54	5.1	4.3	4

Recalls need some care. There were 289 returns from salvo to quarantine, but 163 of them (56%) came within a minute of the call and every one of those was entered by the same staff id, which reads as a call-and-undo habit and not a stand-down (fig. 8). Counting only teams that sat in Salvo for at least 5 minutes leaves 84 real recalls touching 44 of the 139 flown teams (32%), in 32 bursts of which 4 sent back five or more teams at once. Wind doesn't explain them: the 10 minutes before a burst averaged 6.3 mph against 5.9 mph across launch hours ($p = 0.38$), and launches took place in a 5 minute mean wind of up to 13.8 mph, with gusts to 20.7. Real recalls are also concentrated, with the 14 teams having the most holding 60% of them and 6 being the most for any one team. I expected frequency crowding or avionics problems to explain which teams, and the data doesn't support it. Teams whose tracker shared a centre frequency with another team were recalled as often as the rest (32% against 33%, $p = 0.72$), and teams with a wiring, radio or GPS issue logged in review were recalled 37% of the time against 29% ($p = 0.32$). The 17 teams with the least tracker overlap were recalled less (12% against 34%, $p = 0.09$), the only hint in that direction. Teams that registered a video transmitter were recalled less often as well (20% against 36%, $p = 0.05$), for a reason I can't see in the data.




Figure 8: Time in Salvo Before a Return to Quarantine

Salvo size

The part of this I find most useful for planning is salvo size, because it is the one thing the range can change on the day. A salvo here is the group of teams moved from Pad Queue to In-Transport together. There were 41 of those, 15 with six or more teams, and those 15 carried 80% of the launches so they are what I fitted. Inside a salvo the pad stages move in lockstep: every member enters racking and Ready to Fly within about 2 minutes of the others, with Ready to Fly about 43 minutes after dispatch, and then teams launch one at a time about every 3.4 minutes. The next salvo went out a median of 121 minutes after the last. That interval fits 106 minutes of fixed cycle plus 2.1 minutes per extra team (95% interval -3.4 to $+7.6$, 11 intervals). About 24% of dispatched teams bounced back before launching, and that share did not rise with size (-0.7 points per extra team, interval -2.9 to $+1.6$). Setup (dispatch to the first launch) takes about 57 minutes and does grow a little with size among these salvos, 2.1 minutes per team ($p = 0.05$), which is why I didn't want to just assume the overhead was fixed.

That fit says what a dispatch costs once it is made, and it takes for granted that the teams are there to fill it. Figure 9 checks that for the 15 dispatches of six or more teams. The blue bar is the salvo that went out, and the grey bar on top is the teams that were staged at that moment and not sent, so the full height is what was available. Only 5 of the 15 had 20 or more available and 7 had 14 or more, and the median available was 12 against a median dispatched of 10. When fewer than 20 were staged (10 dispatches) the range sent about 85% of them on average, so at most dispatches the supply of staged teams was the limit. The 5 dispatches with 20 or more staged are the ones that need explaining: the range sent 11, 14, 14, 13 and 16 of 31, 30, 21, 25 and 22, and left 61 teams behind.




Figure 9: Teams Dispatched Against Teams Staged at Each Large Salvo

Two explanations I could test for the left-behind teams are that they needed a different kind of pad and that tracker frequency conflicts kept teams apart. Neither holds up. 59 of the 61 were assigned solid pads, the same kind as the teams that went, so it wasn't hybrid or two-stage pads. And the dispatched sets had about as many tracker conflicts as random draws from the same staged pool: 21 pairs sharing a centre frequency against 23.0 expected ($p = 0.39$), and 111 pairs with overlapping declared ranges against 118.5 ($p = 0.31$), across the 12 dispatches that left anyone behind. If the desk had been keeping conflicting trackers out of the same salvo there would be fewer. Adding the teams already on the pads from the earlier salvo to the comparison gives the same answer (55 conflicting pairs against 58.9 expected, $p = 0.37$, and 204 against 211.3, $p = 0.36$).

Pad assignment is the other thing I could look at (fig. 10). The DFMS only keeps each team's final pad, so I use that as the assignment. The 20 solid pads are four rows of five, and over the whole event row A

took 63 of the 123 solid flights, row B 34, row C 22 and row D 4. That pattern is what I would expect if the rows are filled front to back, with the back rows kept for other flights and backfilled when a salvo needs them, and the data fits that. Dispatches of 11 or more teams (7 of them) averaged 4.0 teams from rows C and D, against 0.5 for the 8 smaller ones, so the back rows were used when a salvo was large. At the 5 dispatches with 20 or more staged, rows A to C were nearly full (about 5 teams from each of A and B and 4 or 5 from C) and row D took 0 to 2. What the data can't show is whether a salvo was ever held to its size by the pads. 43 of the 61 left-behind teams (70%) have the same final pad as a team that went, but with backfilling that doesn't mean they couldn't have been moved, and teams were sometimes moved to a different pad after being assigned, which is why a row can show more than five dispatched teams. The other thing I could check is whether the left-behind teams were simply not ready. All 61 were still in Quarantine at the dispatch, so none had been called into Salvo, which puts the decision at the call and not at the pads. A third of them (33%) had cleared FRR less than 30 minutes earlier, against 7% of the teams that were sent, so fresh FRR completions explain some of it. That is mostly the first dispatch on 12 June, which came about 35 minutes after the 08:00 bulk of completions. At the other four dispatches the left-behind teams had been through FRR for a median of 2 to 15 hours. 56 of the 61 waited more than an hour for their own dispatch (median 146 minutes). So not enough teams through FRR explains one dispatch of the five, and for the rest I can't say why the range called 11 to 16 when 21 to 31 were staged. Whether a Quarantine team was physically ready to fly isn't recorded, and neither are recovery crews, the waiver window, weather or caution about bounce, so I can't separate those.




Figure 10: Staged Teams by Assigned Pad Row at the Five Dispatches With 20 or More Staged (rows are each team’s final pad)

A closed form curve with a fixed overhead can only ever say bigger is better, so I built a simulation that puts the costs the event log can measure back in. It launches all 139 teams from a ready pool in salvos of at most K teams. Each cycle takes the fitted interval with a 25 minute random spread, each team bounces back to the pool with the observed probability, and the last salvo is whatever is left. The parameters are redrawn 800 times by resampling the 2025 salvos, so the bands include how little data there is. Simulated hours run about 30% above the 32 range-hours we actually used, so fig. 11 is scaled to 32 hours at the 2025 mean salvo of 9.8 teams.




Figure 11: Simulated Range-Hours Needed by Salvo Size

At the 20 pad limit the best size was 20 in 73% of runs, 17 or more in 99% and 14 or more in all of them.

Salvos of 14, 17 and 20 would have cleared the 139 teams in about 24, 21 and 19 range-hours (80% interval 16 to 22 at 20) against 32, saving 8, 11 and 13 hours, which is about 1.4 launch days at 9.5 hours a day. The returns diminish: 14 already captures 60% of the gain from 20. Against the launch window, salvos of 12 or more would have cleared all 139 teams inside the 28.5 hours of three launch days, and two launch days (19 hours) is only reachable at 20 and even then at about even odds, so I wouldn't promise it. It flips only if bounce grows with size. If each extra team added 3 points to the bounce rate the best size would be 15 or 16, still about 9% faster than 2025, and at 5 points it would be 11. The upper bound from the data is 1.6 points. With the bounce and cycle slopes both at their upper bounds, sizes 14 to 20 take within 3% of the time of size 10, so trying them would cost nothing.

There are four cautions. First, there are only 15 salvos and none above 16, so the bands are wide and 20 rests on the fitted lines. Second, the model keeps range clearing and recovery inside the fixed cycle, so if recovery time per rocket grows with salvo size the gain is smaller (the cycle upper bound covers that). Third, it ignores the daily waiver windows. Fourth, and the ones 2025 speaks to most directly, it assumes the pool is never empty and the pads are never short. The staged pool was at 20 or more for 9% of launch hours, and the biggest salvos stopped at 11 to 16 teams for reasons the log doesn't give. Dispatching whatever was staged, capped at 20, would have averaged 13 teams a salvo instead of 10, which the simulation puts at about 26 range-hours at a fixed size of 13, and that needs nothing more from FRR. Full 20 team salvos on every cycle are the bigger prize, and they need FRR to deliver about 7 teams an hour against the 3.8 it averaged. The reviewers have the room for that, as the next sections show. I would trial it in steps through 14, 17 and 20 while logging dispatch size and the staged count at each dispatch, which would replace these bounds with measurements.

Flight outcomes

Nominal boost and nominal recovery together describe 43% of flights, and every other path is a good boost with an off-nominal or unobserved recovery (32%), a boost failure (20%) or an unrecorded outcome (5%). Figure 12 follows every flight from its boost outcome on the left to its recovery outcome on the right, with the band width showing the number of flights. Of the 104 nominal boosts, 60 had a nominal recovery, 22 deployed main at apogee, 13 failed (9 ballistic, 4 with no main, shredded or tangled) and 9 were not recorded. The 28 boost failures were 13 shreds, 10 CATOs and 5 stability failures, and 16 of them came down ballistic, so for those flights the recovery outcome is mostly the boost outcome continuing. 7 flights have no outcome recorded at all.




Figure 12: Boost Outcome to Recovery Outcome (flights)

Figure 13a shows the rate of boost failure and hard failure (boost or recovery) by class, with 95% intervals. The 30k COTS class had the highest hard failure rate, 11 of 21 flights (52%) against 30 of 111 (27%) elsewhere ($p = 0.04$), driven by boost failures (43% against 16% in 10k COTS), though the intervals only just separate. I compared three classes, so 0.04 doesn't clear a corrected threshold of 0.017 and I would read it as a lead. The class also logged the fewest issues per team, which is either a cleaner set of rockets or lighter scrutiny, and the data can't separate the two. Motor size looks like the better lead (fig. 13b). 19 of the 30k COTS flights were on O class motors, and boost failures rise with the largest impulse class on the rocket: 14% on L (1 of 7), 17% on M (13 of 77), 24% on N (4 of 17) and 32% on O (10 of 31). O class flights outside the 30k COTS class failed 25% of the time (3 of 12) against 37% (7 of 19) inside it ($p = 0.70$), so the class adds little once motor size is known. O against L to N is $p = 0.13$ and the trend across classes $p = 0.08$, so this is a lead and not a result. CATOs did fly larger motors, with a median maximum total impulse of 17,055 N-s against 9,671 for nominal flights ($p = 0.03$), and COTS and SRAD motors failed at the same rate (21% and 20%). Rail exit velocity (median 30.5 m/s for nominal boosts against 32.4 for CATOs, 30.0 for shreds and

31.2 for stability failures) and wind (a 10 minute gust median of 10.2 mph for boost failures against 11.8 for nominal boosts, $p = 0.35$) did not separate boost outcomes.




Figure 13: Failure Rates by Class (a) and by Motor Impulse Class (b)

Tracker positions

The DFMS logged 3,232 tracker positions from 26 trackers, which covers 24 of the 139 flights (17%). Logging began at 08:53 on 12 June, and after that only 20 of the 86 Featherweight trackers on flown rockets (23%) produced any data. After removing invalid fixes (2%), 10 flights have a usable flight profile and 8 of the 18 trackers with five or more fixes never left the pad. For the 8 flights in the 10k COTS class with a profile, GPS apogee was between 23% below and 20% above the 10,000 ft target (mean -1%, standard deviation 14%), measured against each tracker’s own ground level (fig. 14). The other two were a 10k SRAD solid flight, 1% below, and a 30k multi-stage flight, 50% below. Crowded frequencies don’t explain the gaps: captured and uncaptured trackers had the same median overlap, 29 teams ($p = 0.96$), so coverage looks like a ground station or logging limit. It also leaves too few flights to check the LCO’s recovery labels. After logging began, 1 of the 9 flights labelled recovery unknown had a tracker record. Logging every tracker would make the DFMS a source of measured apogee, descent and drift for every flight.




Figure 14: GPS Apogee Against the Class Target (one row per flight)

Safety issues

Reviewers logged 270 issues on 87 flown teams. The 18 teams with five or more issues account for 47% of them, and the 33 teams with a critical issue hold 54%. Figure 15 shows what was flagged. Recovery hardware and deployment (76 issues) and wiring & electronics (67) are the two biggest categories, 53% of issues together. Motor & propulsion issues are the most severe, with 9 of 14 critical, while 1 of 35 checklist issues and none of the 15 venting issues were critical. Checklist, venting and marking issues (CP/CG, team number, frequency labels) total 66, or 24%, spread over 48 teams, and only 3 of them were critical. Fixes aren’t always documented either: 66% of issues carry a resolution note, against 56% of critical ones (37 of 66) and 69% of non-critical ones, a gap that is not statistically significant ($p = 0.07$).

Issue counts don’t separate outcomes (fig. 16a). Teams with a nominal boost averaged 1.87 issues and teams with a boost failure 2.00 ($p = 0.79$). For recovery the means are 1.85 for nominal and 2.26 for ballistic, no main, tangled or shredded descents ($p = 0.46$). Having a critical issue didn’t change boost success either (77% against 79%, $p = 0.81$), and teams cited for recovery hardware failed recovery 33% of the time (13 of 39) against 27% (22 of 81) for the rest ($p = 0.52$). With only 28 boost failures, the confidence interval




Figure 15: Safety Issues by Category

of -0.9 to $+1.1$ issues rules out large effects but not small ones. Issues do cost time, though. Teams that cleared FRR later had more issues logged (fig. 16b and c): 1.0 per team on 11 June, 1.8 on 12 June, 2.8 on 13 June and 4.9 on 14 June, with the share of teams holding a critical issue rising from 15% to 43% over the same days. Issue count and the order of FRR completion correlate at $\rho = 0.22$ ($p = 0.008$), and by the end of 12 June 61% of teams with a critical issue had cleared FRR against 80% of the rest. The cost shows up as being called later and not as a longer process, since the elapsed time from check-in to FRR was similar, a median of 74 hours with a critical issue and 71 without.




Figure 16: Safety Issues by Flight Outcome (a), and by the Day a Team Cleared FRR (b, c)

Review capacity and reviewers

FRR was run by very few people. Seven people completed the 145 FRR reviews: one did 54%, three did 96% and five did 99%. The flight records agree, with six people appearing as the final safety signer and the top three signing 51%, 25% and 19% of flights. One person completed 65% of the FRR reviews on 12 June and 75% on 13 June, and reviews overlapped very little: while any FRR review was open the average was 1.2 running at once, and the maximum was 4. FSR drew on far more people. 19 completed FSR reviews, the top five doing 53%, and as many as 24 reviews were open at once on the conference days (fig. 17a). I read the concentration as who happened to be free to do FRR and not as a staffing gap. FRR was not what limited the pace, since the full-pad salvos above need about 7 FRR reviews an hour, which one reviewer at 8.2 minutes a review can just supply, and the 3.8 an hour we got was set by when teams were called and not by who was reviewing.

FSR reviewers also differed in how often they sent teams back (fig. 17b). Twelve reviewers completed at least 10 FSR reviews each (203 of the 233 reviews that ended in a pass or send-back), and their send-back rates ranged from 10% to 71% against 36% overall. Each dot is one reviewer and the curved lines are the limits you would expect from chance alone given how many reviews they did, so dots outside them are further from the overall rate than luck explains. Five of the twelve sit outside the 95% limits and one outside the 99.8% limits. After shuffling reviewer labels within day, first or repeat review, and class group, a gap this large appeared in 0.2% of shuffles ($p = 0.002$), and about 17% of the variation in send-back outcomes lies between reviewers. If every reviewer above the median rate of 32% had matched it, 22 of their 72 send-backs (31%) would not have happened, unless the stricter reviewers were catching real defects, which these data can't show. Two thirds of re-reviews after a send-back (67%) went to a different reviewer, and those passed

86% of the time against 74% when the same reviewer re-checked, which is not statistically significant ($p = 0.23$). Reviewers are not identified anywhere in this report.




Figure 17: Who Did the Reviews (a), and FSR Send-Back Rate by Reviewer (b, one dot per reviewer, no labels)

Radio

141 teams registered 425 transmitters, a mean of 3.0 per team and a maximum of 10. Trackers are 297 of them (70%), followed by 44 handsets, 43 video links and 41 ground station links. Featherweight GPS units are the biggest named group (127, 30%), and LoRa is the most common modulation (206, 48%). By band, 33 cm carries 54% of transmitters (231), 70 cm 22% (95), 5 GHz and above 10% (44), 2.4 GHz 6% (24), FRS 4% (18) and 2 m 3% (12). Figure 18 splits the transmitters by manufacturer group and use, and the congestion of each band follows it.




Figure 18: Radio Environment: Transmitters by Manufacturer & Use

To see how crowded each band is, I counted for every frequency how many teams have a declared range covering it, using the declared centre frequency plus or minus half the bandwidth (fig. 19 and table 5). Every panel uses the same vertical scale, so the bands can be compared directly. These are declared ranges, not measured interference.

Table 5: Declared Congestion by Band						
Band	Transmitters	Teams	Peak teams	Peak at (MHz)	Span with 10+ teams	Peak without wide
33 cm (902–928 MHz)	231	121	31	915.0	89%	25
70 cm (420–440 MHz)	87	64	12	434.5	0.5%	10
2 m (140–152 MHz)	13	9	2	146.1	0%	2
FRS (462–468 MHz)	18	16	3	462.7	0%	3
2.4 GHz (2400–2500 MHz)	23	15	8	2400	0%	3
5 GHz (5150–5925 MHz)	40	31	13	5798	3.6%	0

The 33 cm band is the crowded one. 89% of it has 10 or more teams covering it, 7.5% has 20 or more, and the peak is 31 teams at 915 MHz (29 of them trackers). 39 transmitters from 26 teams are declared

at exactly 915.000 MHz (33 trackers and 6 ground station links), which is more than twice the next busiest centre frequency, 903.8 MHz with 12 teams (fig. 20). The band’s 231 transmitters sit on only 70 distinct centre frequencies, and 188 of them (81%) are on a centre that at least one other team also declared. 164 of the 203 trackers share their exact centre with another team’s tracker. Twelve declarations from 9 teams are 5 MHz wide or more and raise the floor, but without them the peak is still 25 teams and half of the band still has 10 or more. The commonest declared bandwidths are 2.5 MHz (93 transmitters) and 0.5 MHz (42), so most of the overlap is neighbouring channels and not only identical ones.




Figure 19: Teams With a Declared Range Covering Each Frequency, by Band (dotted line: 10 teams)

The other bands look very different. In 70 cm the trackers use narrow channels (median 0.1 MHz), so the crowding is a cluster and not a plateau: the peak is 12 teams at 434.5 MHz, 10 teams declared exactly 434.55 MHz, and only 0.5% of the 420–440 MHz span has 10 or more teams. The 2 to 3 team floor across the whole band comes from 5 declarations 5 MHz wide or more, one of them 50 MHz. Counting only narrow declarations, 57 of the 77 narrow 70 cm trackers (74%) have another team’s tracker range over them, and 43 of the 81 trackers share an exact centre. The 2 m and FRS bands are nearly empty, with a peak of 2 teams at 2 m and 3 at FRS, and the 2.4 GHz band peaks at 8 teams, all of which declared exactly 2400 MHz. Most of that comes from wide video and ground station declarations, since the peak is 3 teams without them. The 5 GHz band holds only video and ground station links, all 20 MHz or wider, and they cluster around 5.8 GHz: the peak is 13 teams at about 5798 MHz, 8 teams declared exactly 5800 MHz, and 3.6% of the 5150–5925 MHz span has 10 or more teams. Only 900 MHz is crowded across the whole band. 70 cm and 5.8 GHz have a single hot spot each, and 2 m, FRS and 2.4 GHz do not look crowded on declared ranges.

A LoRa link tolerates sharing better than the overlap suggests, and declared ranges can’t tell channel sharing from real interference. They do show that the declared plan is crowded in the two bands that carry nearly all trackers. In 33 cm every one of the 203 trackers has another team’s tracker declared over it, and 192 of 193 when the wide declarations are left out. In 70 cm it is 57 of 77 trackers. The transmitter and flight data can also be checked against each other. 135 teams appear in both, 4 flown teams have no transmitter on file and 6 teams with transmitters have no flown record. 50 flown teams (36%) registered a single tracker. Of the 53 flown teams flagged for the live video challenge, 34 have a video transmitter registered and 19 do not.




Figure 20: Teams Declaring the Same Exact Centre Frequency (six busiest per band)

Methodology

The 2024 analysis used the action history from Trello. The 2025 DFMS is a PostgreSQL database, and the data for this analysis is a full dump of it, plus the three SQL extracts pulled earlier (flight records, safety issues and transmitters), which the dump reproduces. The table that matters most is the status transition log: 4,058 rows over 145 flight records, each saying which record moved, from which status to which, who moved it and when. The timestamps are stored in UTC and the dump shows them with a -5 hour offset, so everything was converted to Mountain Daylight Time (UTC -6). The analysis was done in Python (pandas, SciPy) and the figures drawn in matplotlib as vector graphics. The dump’s tables of people and of team contact details were never read.

Table 6 is the 2025 equivalent of the 2024 events and states tables. Generally the statuses are meant to be visited in order, but teams can be sent backwards at almost every step, and the log shows it: 626 backwards moves, mostly Salvo to Quarantine, FSR to Ready for Safety, and rollbacks inside the pad sequence. To keep send-backs from making a team look finished early, a team’s time for a milestone is its last forward crossing: FSR complete is the last Safety Review to Ready for Final Safety, FRR complete the last Final Safety Review to Quarantine, pad queue the last Paged Salvo to Pad Queue, and launch the last Pre-Flight to Flight. Intervals between teams follow the 2024 method, with gaps above $\mu + 3\sigma$ dropped to remove the overnight waits. Reviews longer than two hours, as in 2024, or shorter than 30 seconds are dropped. One staff test record is excluded: it passed FSR, FRR and pad queue within 90 minutes of check-in on Monday and was reset on Thursday. The **quarantined** timestamp in the earlier flight record extract is the time a team first entered Quarantine, which is FRR completion and not a launch time, so launch times come from the transition log.

Table 6: DFMS Statuses

2025 DFMS status	Stage used in this analysis
Ready for Safety	Ready for FSR (check-in); later visits are FSR remediation
Safety Review Queue, Safety Review Paged	FSR waitlist
Safety Review	In FSR
Ready for Final Safety	Ready for FRR
Final Safety Queue, Final Safety Paged	FRR waitlist
Final Safety Review	In FRR
Quarantine	FRR complete, waiting for a salvo
Salvo, Paged Salvo	Called for a salvo
Pad Queue, In-Transport, Racking, Ready to Fly, Pre-Flight	Pad process
Flight, Landed	Launch and landing

The flow diagram sorts each of the 144 teams by how it finished FSR (no send-back, sent back and completed, or never completed), how it finished FRR (cleared with no return, cleared after a return to an earlier stage, not cleared, or never reached) and whether it launched, with a return being a move from Final Safety Review back to Ready for Final Safety, from Ready for Final Safety back to Ready for Safety, or from

Quarantine back to Ready for Final Safety.

For the salvo work, a dispatch is the set of teams moved from Pad Queue to In-Transport within 3 minutes of each other: 41 dispatches, 15 with six or more teams. A team bounced if its next move out of the pad sequence went back to pad queue, salvo or quarantine, and launched if it reached Flight within 6 hours. Setup is dispatch to first launch, launch spacing is the time from first launch to the 80th percentile launch divided by the launches in between, and the cycle interval is the time to the next dispatch of six or more teams on the same day. The slopes against salvo size are ordinary least squares fits with t based 95% intervals. The simulation launches all 139 teams in salvos of at most K teams, with each cycle taking the fitted interval (106 minutes plus 2.1 minutes per team) with normal noise of 25 minutes and a 30 minute floor, each team bouncing back to the pool with probability 24%, and the last salvo being whatever is left. The interval and bounce parameters are refitted on each of 800 resamples of the salvos, and results are scaled so the 2025 mean salvo equals the 32 observed range-hours. Sensitivity runs raise the bounce probability by 1 to 5 points per team above the mean and set the cycle slope to its upper bound.

The staged pool is the number of teams in Quarantine, Salvo or Paged Salvo, and the pads are the number in Pad Queue, In-Transport, Racking, Ready to Fly or Pre-Flight, both read from the status log at a moment in time. The teams available at a dispatch are the teams dispatched plus the teams that were staged 30 seconds before it and were not part of it. The dispatched teams have to be excluded from that count, because the desk moves a group through Pad Queue to In-Transport in one go and many of the dispatched teams are still in Salvo or Paged Salvo 30 seconds earlier. A left-behind team is one of those staged teams that was not dispatched. Launch hours are 07:00 to 19:00 on 11–13 June, sampled every 5 minutes. An idle spell is a gap of more than 30 minutes between consecutive launches on the same day, and the pool is read at the middle of the gap. FRR completions per hour are counted over the hours between 06:00 and 20:00 on launch days that had at least one completion, and the roughly 7 an hour needed for 20 team salvos is the launch rate the simulation implies at that size. Reviews open at once is the overlap of FRR review intervals. The FRR lead time is the first entry into Final Safety Review on a launch day against the first dispatch of six or more teams that day. Pad rows come from the DFMS pad table, where the solid pads are rows A to D of five pads each, and a team’s pad is its final pad because the DFMS keeps no history of pad changes. A left-behind team shares a pad with a dispatched team when its final pad is the same as one of theirs, and has an occupied pad when its final pad belongs to a team still in transport, racking, ready to fly or pre-flight from an earlier salvo. For tracker conflicts, two teams conflict when any of their declared tracker centre frequencies are within 0.01 MHz, or, in the wider test, when any declared tracker ranges overlap. Conflicting pairs inside each dispatched set are summed over the 12 dispatches that left anyone behind and compared with 400 random draws of the same sizes from the staged pool at each dispatch, and a second version also counts conflicts with the teams still on the pads from earlier salvos. A team’s time since FRR is measured from its last Final Safety Review to Quarantine move to the dispatch.

Recalls are returns from Salvo to Quarantine. A real recall is one after at least 5 minutes in Salvo, and returns within a minute are counted as reversals. Bursts are recalls within 10 minutes of each other. Wind before a burst is the mean of the 10 minutes before its start, from the weather station’s 266,034 packets (Peet Bros Data Logger Mode, wind speed in 0.1 km/h), compared with the same measure at every 10 minute slot between 08:00 and 18:00 on launch days. Team level recall tests use Spearman, Mann–Whitney and Fisher exact tests, and a permutation test across teams for video. Reviewer consistency covers the 233 FSR reviews that ended in a pass or send-back (22 people), keeping the 12 reviewers with at least 10. The chi-square statistic of send-back rates across reviewers is compared with 5,000 shuffles of reviewer labels within day, first or repeat review and class group, and the intraclass correlation is the one-way ANOVA estimate on the binary outcome. Tracker fixes need a GPS lock and plausible coordinates and altitude, apogee is the highest fix above the tracker’s own 5th percentile altitude, and the class target comes from the event class table.

Nominal boost and nominal recovery are the DFMS outcomes boost-success and recovery-ffp. A boost failure is a shred, a CATO or one of four stability outcomes, a recovery failure is a ballistic, no main, shredded or tangled descent, and a hard failure is either. Unknown and blank outcomes (4 boost, 16 recovery, 3 blank) are left out of the rates, and the intervals in figs. 13 and 16 are Wilson and Welch 95% intervals. Safety issue categories come from keyword rules on the free text note, first match wins, checked by reading every note. Boundary cases, mostly deployment electronics that could sit under wiring or recovery, were settled by the rules, so category counts are approximate. Critical is the DFMS flag on the issue. The 287 issue rows in the dump drop to 270 after removing back-of-queue notes, one further row and unflown rockets, and flown teams with no issue rows count as zero issues. Radio overlap uses the declared range (centre frequency plus or minus half the bandwidth), and two records with an empty range are left out. Band occupancy is counted on a 5 kHz grid (2.5 kHz for FRS, 0.5 MHz for 2.4 GHz and 1 MHz for 5 GHz), counting a team once wherever any of its declared ranges strictly covers the frequency, and a declaration narrower than a grid step still occupies its own centre. Wide declarations are those of 5 MHz or more, and two trackers share a centre when their declared centres are within 10 kHz. The tests are Welch t , Mann–Whitney, Fisher exact, Spearman and permutation tests, with a 95% interval on the Welch difference. Everything here is exploratory, and the three class comparisons use a corrected threshold of 0.017.

The data has its limits. Length, lift-off mass and effective rail length mix units and were not used, and rail exit velocity was used with entries above 75 read as ft/s and converted, with 4 implausible entries left out. 34% of issues have no resolution note, 3 flights are on a pad recorded as not defined, and Post-Flight and Complete were never used, so every flown team ends at Landed. Wind comes from one station with no data before 06:44 on 11 June, and 163 of the recalls were entered by a single staff id and treated as workflow. The simulation only has 15 salvos to learn from. The staged pool is only as accurate as the moves entered in the log, and a team moved late shows up late. The pre-event video safety review is not recorded in the DFMS, so nothing here can say what it already catches, and the waiver opening times aren’t in the data either. Pad assignments are recorded only as each team’s final pad. The 2024 caveat still applies, and the analysis can only be as good as the entries made in the DFMS, so it supports broad trends and not

conclusions about any one team.

This report contains no team numbers, team or school names, rocket names, staff identifiers or quoted notes. Reviewers appear only as unlabelled statistics, the tracker comment field (which can hold callsigns) was not read, and the free text issue notes, some of which name individuals, were categorised and not reproduced. The only identifiers used are class names.

AI Usage Statement

This document was compiled with the help of artificial intelligence (AI) tools, which were used for data processing, analysis, figure generation, and drafting of the text.