

Investigation Report

F-TCF Error — 2026 Newport Bermuda Race

Prepared for: Newport Bermuda Race Competitors and Followers

Prepared by: Bermuda Race Organizing Committee (BROC), BROC Technical Committee, US Sailing Offshore Office

Date: August 25, 2026

1. Executive Summary

On the morning of Friday, June 19, 2026, the Bermuda Race Organizing Committee (BROC) issued Forecast Time Correction Factors (F-TCFs) to the Newport Bermuda Race fleet that were subsequently found to be in error. The error was not raised until Tuesday, June 23, when a competitor flagged a questionable F-TCF discrepancy between two nearly identical boats. That report triggered an investigation that has found two distinct, independent problems:

1. **A human process failure on race morning.** Additional weather and current files were downloaded from within the routing software for later comparison purposes. Unknown to the operator, two conflicting GRIBS files remained active alongside the official data sources during the optimization run, producing corrupted, effectively randomized elapsed-time calculations. This was the dominant cause of the Friday morning error.
2. **A separate software defect.** The routing software produced different results when routing the fleet in a batch ("fleet mode") versus routing boats individually. This defect was discovered early during the investigation (by June 25) and has since been identified and corrected by the software's developer. While at the time this bug was discovered, it was the only known cause of F-TCF errors, it was later determined to be a very small contributor to the magnitude of erroneous F-TCF values.

Corrected, individually run F-TCFs were delivered to race officials on Wednesday, June 24, at which point the Race Committee corrected the results. Following several days of redress hearings, the International Jury directed publication of results on June 26, using the corrected

F-TCFs. Those results have since been confirmed accurate and reproducible through multiple independent re-runs.

This report documents the sequence of events relating to the erroneous F-TCFs, the root causes identified to date, an analysis isolating the relative impact of each cause, the steps taken to verify the corrected results, and the corrective actions BROOC is adopting to prevent recurrence. Consistent with our July 25th update to the fleet, BROOC and its Technical Committee are sharing what we learn so the broader offshore racing community can benefit. Section 8 also identifies broader policy questions this incident has raised — about F-TCF's fairness value relative to its complexity, and about the underlying ORR rating rule — which BROOC intends to address through separate reviews.

2. Background

2.1 F-TCF scoring

F-TCFs are time correction factors used for time-on-time (ToT) scoring. A boat's corrected time equals its elapsed time multiplied by its F-TCF. F-TCFs express the ratio of expected boat speed (or equivalently, elapsed time) between boats for the specific course and forecast conditions of a given race, rather than a fixed, condition-independent rating. For example, a boat rated 0.900 is predicted to sail the course at 90% of the speed of the “scratch” boat rated 1.000.

F-TCFs for the Newport Bermuda Race are determined as follows:

1. Boat-speed polars are calculated by the ORR velocity prediction program (VPP) for each boat and formatted for use in the routing software, Expedition.
2. The racecourse — start, finish, intermediate marks, and exclusion zones (as defined in the SI) — is entered in Expedition.
3. Wind and current forecast files (GRIBs) for the expected period of the race are loaded into Expedition. For the Newport Bermuda Race, these were pre-selected to be an HRRR wind file for the start/coastal area, a GFS wind file for the offshore body of the race, and an RTOFS file modeling the Gulf Stream.
4. For each boat, Expedition calculates the optimal route from start (at its published scheduled starting time) to finish — the route yielding the least elapsed time — combining the course, the forecast files, and the boat's polars.
5. Fleets can be routed in batch (“fleet mode”), which is faster and uses a PC's parallel-processing capability to route many boats simultaneously, or one boat at a time (“individual mode”).

6. Resulting routes are reviewed visually to look for implausible cases, and a sample of F-TCFs are compared against corresponding boat ORR certificate ratings as a secondary check.
7. Calculated elapsed times are exported to a spreadsheet, a scratch boat is selected for each division, and F-TCFs are calculated as the ratio of the scratch boat's elapsed time to each boat's elapsed time. These are published to the competitors and provided to the race's scoring platform.

2.2 Purpose of this investigation

The BROC, its Technical Committee, and its service providers — including US Sailing, which provided F-TCF generation as a service to the BROC, and operates the Expedition software to do so — are committed to determining the root cause(s) of the erroneous F-TCFs and to implementing corrective actions. That work began before the race had finished and has continued since. This report reflects our findings to date and the corrective actions we are putting in place. As stated in our July 25th update to the fleet, BROC, its committees, its selected rating organization, and its service providers are ultimately responsible for the results of the race and for the anomaly that affected it. This investigation is limited to the technical nature of the F-TCF anomaly. It does not explore any actions taken or not taken by the Race Committee or Technical Committee under the RRS, or the decisions or actions of the International Jury because of the discovery of erroneous F-TCFs.

3. Timeline of Events

Date	Event
Thu, June 18	A full trial run of the F-TCF process was executed by US Sailing's Offshore Office as if the race were starting that day, using standard procedures. Results were reviewed and deemed accurate. The instance of Expedition used by US Sailing for all services performed for the BROC in 2026 is V12.5.5
Fri, June 19 (race morning)	HRRR, GFS, and RTOFS (the pre-selected official forecast) files were downloaded per the operator's typical procedure (via email query to SailDocs). In addition, ECM and Mercator GRIB files were downloaded from within Expedition itself, to support a later separate comparison of forecast models. Only HRRR, GFS, and RTOFS were to be used in the generation of F-TCFs. The ECM and Mercator GRIBS for the same region and timeframe were downloaded to support post-race internal research on forecast model accuracy.

Date	Event
Fri, June 19 (race morning)	The fleet — approximately 140+ boats — was routed in Expedition's fleet mode, in two batches. Resulting routes were reviewed and judged plausible; a sample of resulting F-TCFs was compared against certificate ratings and judged plausible. A weather feature that morning (a ridge of high pressure drifting into the bottom of the course in the GFS model) was known to be diverting slower boats and not faster ones, providing a plausible, independent explanation for some deviation.
Fri, June 19	F-TCFs generated that morning were delivered by US Sailing's Offshore Office to the BROC. The BROC published the F-TCFs and the GRIB files to the notice board and loaded the F-TCFs into the scoring system before the start of the race.
Tue, June 23	A competitor requested US Sailing's Offshore Office to review the F-TCFs for two boats that were expected to have similar performance characteristics but had been assigned substantially different F-TCFs in the published scratch sheet. The competitor was the owner/operator of the two boats in question, Oakcliff Farr 40's. Re-running those two boats individually produced F-TCFs markedly different from those published on Fri, June 19. This targeted optimization run used only the official GRIB files (HRRR, GFS and RTOFS).
Tue, June 23	The full fleet was re-run in fleet mode, again in two batches as was done on Friday. F-TCF results again differed from Friday's — and also differed slightly from the two-boat individual re-run just completed. This re-run used only the official GRIB files (HRRR, GFS and RTOFS).
Tue, June 23	The discrepancy between fleet-mode and individual-mode results was escalated to Nick White, owner and developer of Expedition. Race officials, including the BROC vice chair, were notified that Friday's F-TCFs were erroneous.
Tue night–Wed, June 24	All 140+ boats were re-run individually (not in fleet mode) using only the HRRR, GFS and RTOFS GRIBs. Corrected F-TCFs were delivered to the BROC and race committee Wednesday morning.
—	Nick White identified and corrected a coding defect present in Expedition's fleet-mode routing (not present in individual mode) and released a corrected version of the software (V12.8.18). After this bug (and fix) were confirmed, US Sailing did not update to the

Date	Event
	latest version, and instead used individual optimization runs to generate the F-TCFs.
—	Because Tuesday's fleet-mode re-runs — which still contained the fleet-mode defect — were themselves very different from Friday's results, a second, larger factor was clearly responsible for the magnitude of Friday's error, but at the time was not understood.
Following days	The Race Committee updated the results based on the corrected F-TCFs received from US Sailing early Wednesday morning, June 24. Redress hearings were held before the International Jury; race results were revised more than once during this process.
Fri, June 24	The International Jury directed publication of final results using the corrected F-TCFs published on the afternoon of Wednesday, June 24.
Since June 26	The published results have been verified as correct and reproducible via multiple independent re-runs, performed by different people, using different instances of Expedition and explicitly defined settings. Subsequent inspection (during July and early August after the race) of Friday's and Tuesday's Expedition output/log files enabled identification of the “larger factor” that was apparent but inexplicable on June 24. The ECM and Mercator files downloaded Friday morning had remained active during the Fri, June 19 optimization run, so the routing engine worked from competing, conflicting wind and current inputs, producing erroneous elapsed times. The optimization runs performed on Tuesday and Wednesday used only the official wind and current GRIBs (HRRR, GFS and RTOFS). This has been confirmed by reproducing the Friday error. Re-running the fleet with GFS and RTOFS active alongside ECM and Mercator has reproduced results consistent with those issued Friday morning.

4. Root Cause Analysis

The investigation identified two distinct, independent causes contributing to the events of June 19–26. The first was the dominant cause of the magnitude of error in the F-TCFs issued on the morning of the start, June 19; the second was a separate software defect (bug) discovered early in the investigation that, while real, contributed very little to the magnitude of the F-TCF errors.

4.1 Root Cause 1 — Human Error – Competing GRIB file condition (primary cause)

On the morning of the start, Friday, June 19, in addition to the official HRRR, GFS, and RTOFS files used to generate official F-TCFs, the operator downloaded ECM and Mercator GRIB files from within Expedition to support a post-race comparison of forecast models. For reasons still not perfectly understood, the ECM and Mercator files were active in the Expedition session. The operator did not notice that the ECM and Mercator GRIBS were selected in the Weather Data selection view, but it is clear from reviewing the optimization output detail that indeed these models were included in the optimization run. When the fleet was then routed for official F-TCFs, the optimizer worked with two competing wind sources (GFS and ECM) and two competing current sources (RTOFS and Mercator) simultaneously, rather than the intended single, consistent set of inputs.

The resulting routes were visually plausible and passed the standard spot-check against certificate ratings closely enough to not raise immediate concern — a review made more difficult by the fact that a genuine, independent weather feature (the moving ridge of high pressure) was already expected to cause some boats' F-TCFs to diverge from their fixed wind speed certificate ratings during the race. The competing-data condition introduced a degree of randomness into the routing solutions that produced erroneous elapsed times and, in turn, erroneous F-TCFs across the fleet. When Expedition has GRIB files loaded for the same area, with the same data-fields, initialization-time, and resolution, it isn't predictable what data it uses at any point during the optimization routine. According to the Expedition Developer, "Expedition selects data at any time and location mainly based on, and in this order: Model run time and Resolution; which works well if you load (eg) HRRR and GFS. However, if you load EC and GFS from the same run time, then which gets used is effectively random."

By the time the Tuesday optimization runs were performed, as US Sailing's Offshore Office was attempting to investigate the reported anomalies, the instance of Expedition had been restarted one or more times, and the incorrect and competing GRIB files were no longer active in the optimization configuration. The operator has no recollection of specifically selecting or de-selecting the conflicting GRIB files in the "Weather Data" configuration. However, it is certain from inspecting the output logs that the Friday optimization run did, and that the Tuesday runs did not, include the conflicting GRIB files.

This condition was not caused by a known defect in the software but resulted from a gap in operating procedure (a human error). US Sailing's Offshore Office's standard practice for generating official race-morning F-TCFs is to fully close and restart Expedition immediately beforehand to refresh the index of available weather files and ensuring only the target data sources are selected for use in the optimization session. That reinitialization step was apparently not performed before the official Friday morning run.

4.2 Root Cause 2 — Software Defect - Fleet-mode routing “Bug” (secondary cause)

During the investigation on Tuesday, June 23, re-running the fleet in “fleet mode” produced results that differed from re-running boats individually under otherwise identical settings. This indicated a performance specific to Expedition's fleet (batch) routing that did not exist in its individual-boat routing. This discrepancy was previously undetected likely because, while fleet mode had been used to generate F-TCFs in prior races, the impact of the “bug” on F-TCFs is small enough to be non-obvious (see Section 5).

This discrepancy was immediately reported to Nick White, Expedition's owner and developer, who confirmed the bug (“I discovered the fleet routing wasn’t calculating the target splines for the polars it loaded” NW). White identified and corrected the underlying code and released a corrected version of the software (V12.8.18). Notably, this defect alone does not explain the magnitude of the Friday morning error: fleet-mode re-runs performed on Tuesday (before the bug fix, and therefore still containing this defect) produced results markedly different from Friday's original F-TCFs, indicating a separate, larger factor — the competing GRIB file condition described in Section 4.1 — was proven responsible.

4.3 Contributing factor — Insufficient verification of F-TCF values before release

Existing quality checks at the time — visual review of routed tracks and a sample comparison of F-TCFs against certificate ratings — were not sufficiently robust to detect the anomaly before F-TCFs were issued to the fleet. Because a real weather effect (the ridge of high pressure) was already expected to cause some divergence from certificate ratings, the checks in place did not clearly distinguish a genuine, forecast-driven deviation from one caused by corrupted input data. A more rigorous verification process — for example, immediate review of optimization log details to confirm the GRIBs used, and independent parallel runs compared for agreement before release, as has since been used to confirm the corrected results — would have provided a stronger opportunity to catch the error before the race started.

5. Isolating the Impact of Each Root Cause

Because Root Cause 2 (the fleet-mode “Bug”) was a defect in Expedition itself rather than a one-time procedural lapse, competitors have understandably asked whether it could also have affected F-TCFs in prior races that used fleet-mode routing. To address that question directly, BROC asked US Sailing's Offshore Office to isolate the impact of each root cause by comparing two sets of Expedition outputs for the 2026 fleet, all measured against the same baseline — the corrected individual-boat runs delivered Wednesday, June 24, which are treated as the most reliable reference:

1. Friday, June 19 fleet runs — affected by both root causes (the competing-GRIB-file condition and the fleet-mode defect), compared against the individual-run baseline.

2. Tuesday, June 23 fleet re-runs — affected only by the fleet-mode defect (the competing GRIB files were not active in that session), compared against the same individual-run baseline. This isolates the defect's impact on its own.

The charts below plot, for each boat, the ratio of its fleet-mode predicted elapsed time to its individual-mode predicted elapsed time (a ratio of 1.00 means the two runs agreed exactly). Figure 1 reflects the combined effect of both root causes; Figure 2 reflects the fleet-mode “Bug” in isolation.




Figure 1 — Friday, June 19 fleet runs vs. individual runs (combined effect of both root causes).




Figure 2 — Tuesday, June 23 fleet re-runs vs. individual runs (fleet-mode software defect in isolation).

The fleet-mode defect in isolation produced deviations roughly one-fifth the size of the combined Friday error, on both a mean and a maximum basis, and affected far fewer boats.

This matters for two reasons. First, it confirms that the extraneous, competing GRIB file data active during the Friday morning run (Root Cause 1) — not the fleet-mode “bug” — was responsible for the great majority of the scoring impact competitors experienced in the original results based on the F-TCF published on June 19. Second, because the fleet-mode defect's isolated effect is small and was present only in fleet-mode routing (not individual-mode routing), BROOC believes it is unlikely to have materially changed the outcome of prior Newport Bermuda Races or other events scored using F-TCF via Expedition's fleet mode. The competing-GRIB-file condition that caused the bulk of this year's error was specific to how the June 19 routing session was implemented and is not evidence of a similar error in prior races.

F-TCFs are not absolute numbers; each boat's F-TCF is calculated relative to its own division's scratch boat ($\text{F-TCF} = \text{scratch boat's elapsed time} \div \text{that boat's elapsed time}$). If an error inflates or deflates a boat's predicted elapsed time by roughly the same factor as another boat — that error largely cancels out in terms of corrected time order of finish. In such a scenario the F-TCFs for different boats are affected by similar amounts and in the same direction and thus will have minimal impact on corrected-time order of finish. In the case of the fleet-mode ‘bug’,

the error mostly ‘washes out’ because most boat's raw predicted elapsed times were incorrect by a similar magnitude and in the same direction. See Figure 3. This effect is likely a contributing factor to why the fleet-mode routing ‘bug’ was not discovered prior to this race. Note: there was testimony at a Redress Hearing that a competitor was aware of an ‘anomaly’ (between fleet mode and individual mode routings) in Expedition sometime 30 – 60 days prior to the race and stated that questions were raised about it in an online forum of Expedition users. The Expedition developer, however, reported to have first been informed of the possible ‘bug’ by the US Sailing’s Offshore Office alert on Tuesday, June 23.

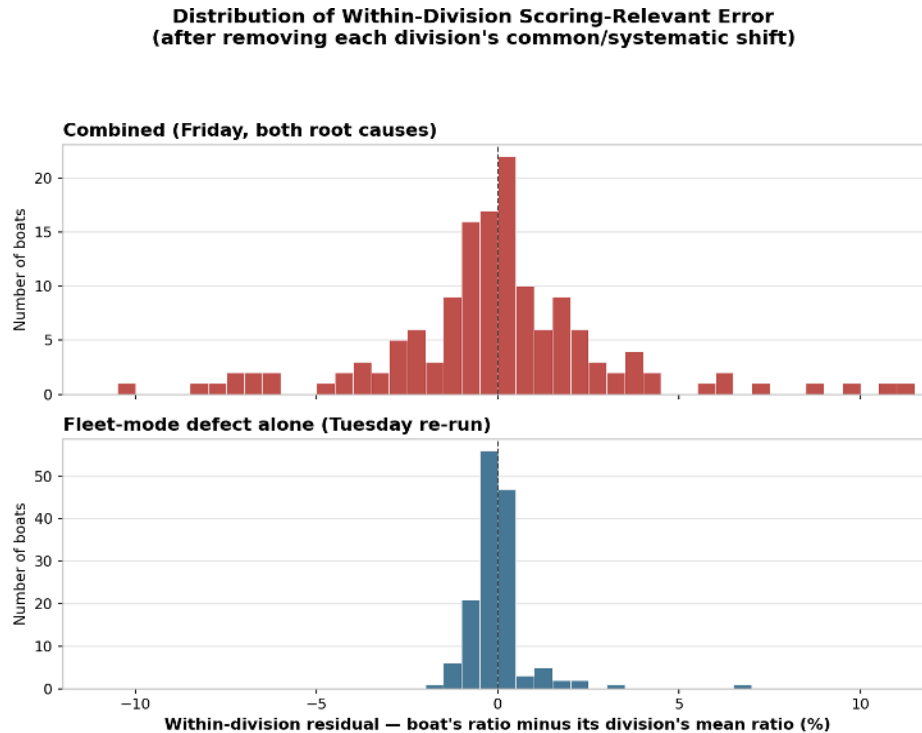


Figure 3 — Distribution of each boat's within-division residual (its ratio minus its own division's average ratio), combined error vs. fleet-mode defect alone. A value of 0% means the boat's error matched its division's average exactly and would not, by itself, change that boat's placing within its division.

Because most (not all) of the fleet-mode ‘Bug’ effect on F-TCFs was similar amongst the fleet, its net effect on finishing order was small. It also reinforces BROC's assessment that the competing-GRIB-file condition on the morning of June 19 (Root Cause 1) — not the fleet-mode ‘bug’ — was responsible for the great majority of the F-TCF impact competitors experienced.

6. Verification of Corrected Results

The F-TCFs used to score the race, as published per the International Jury's direction on June 26, 2026, have been confirmed correct and reproducible. This was established by confirming the correct models were used in the output logs, and by comparing the output of multiple independent runs, performed by different people, using different instances of Expedition, with explicitly defined settings. These runs consistently reproduced the same F-TCF values,

confirming that the corrected F-TCFs — generated using individual-boat mode with only the intended HRRR, GFS and RTOFS inputs active and published on Wednesday, June 24 — are sound.

7. Recommended Corrective Actions

Based on the root causes identified above, BROC's Technical Committee recommends the following corrective actions. Some are procedural and can be implemented immediately; others involve coordination with US Sailing and Expedition's developer and will be worked through as part of our continuing review.

1. **Mandatory reinitialization protocol.** Require Expedition to be fully closed and restarted immediately before any official race-morning optimization run, with no other downloads, comparison runs, or in-app operations performed in that session beforehand.
2. **Written standard operating procedure.** Codify the exact, step-by-step sequence for generating official race-morning F-TCFs — including which files to load, in what order, and how to confirm no extraneous data sources are active — as a checklist that any operator must follow and log. This includes documenting and publishing in advance all Expedition optimization configuration settings, and confirming the GRIBs used during the optimization runs through immediate inspection of the output log details.
3. **Independent parallel verification before release.** Require official F-TCFs to be generated on at least two independent Expedition instances, by different operators, using identical, explicitly defined settings, with results reconciled before F-TCFs are released to race officials or the fleet — the same method used to confirm the corrected June 26 results.
4. **Fleet-mode / individual-mode reconciliation check.** Until fleet-mode routing has been independently revalidated following the developer's fix, cross-check fleet-mode outputs against individual-mode outputs for a defined sample (or all) boats and require agreement within a specified tolerance before any F-TCFs are published.
5. **Developer engagement on software safeguards.** Work with Expedition's developer to add a visible warning or lockout when multiple competing wind or current GRIB files are simultaneously active for the same space and time, and to independently confirm, through regression testing, that the fleet-mode defect fix performs correctly before fleet mode is used in future races.
6. **Stronger, less subjective quality checks.** Replace or supplement the visual “plausibility” review with objective checks — for example, flagging any boat whose F-TCF deviates from a fixed wind-speed certificate rating beyond a defined threshold for manual review, with a documented process for distinguishing genuine forecast-driven deviation from data-integrity errors.

7. **Operator training and redundancy.** Ensure that any operator generating official race-morning F-TCFs is specifically trained on how Expedition indexes weather files and displays active data sources for optimization runs and consider requiring a second trained reviewer to check the live race-morning run in real time.
8. **Adequate time buffer.** Build sufficient time between generation of race F-TCFs and the published start for the verification/review process so these checks can be completed without time pressure from an imminent start.
9. **Retained audit trail.** Preserve full output and log files from every official F-TCF run for a defined retention period, to support faster diagnosis if a question arises in the future.

8. Related Policy Questions Raised by This Incident

This investigation was limited to the causes of the June 19 F-TCF error and the corrective actions needed to prevent that specific failure from recurring. The incident has understandably renewed broader questions from competitors about F-TCF scoring and the selected rating system (ORR) that was used. BROC considers these important — but distinct — questions that warrant their own review, separate from the root-cause findings above. They include the following.

8.1 Is F-TCF a valid approach to improving fairness in offshore racing?

Proponents of forecast-based scoring argue that it more fairly accounts for the weather conditions each boat sails through, substantially correcting for significant weather effects that a fixed, condition-independent rating cannot — particularly relevant on a course like the Newport Bermuda Race, which can see significant (and highly forecastable) weather and current variability along its length. Critics point to the added complexity, the dependence on forecast accuracy and software correctness, and this incident as evidence of the operational risk that complexity introduces. BROC intends to evaluate F-TCF's performance against both sides of this question as part of its ongoing review, weighing measured scoring outcomes across multiple races against the operational risks this incident illustrated.

8.2 Does the complexity of F-TCF — compounded by this error — outweigh its value?

This incident has sharpened concerns that F-TCF's operational complexity — reliance on specialized software, third-party forecast models, and manual procedures — creates risk that may outweigh its scoring benefits, particularly where competitors find it challenging to verify how their F-TCF was produced (though the BROC has endeavored to make available all details to do so). BROC will weigh this trade-off as part of its continuing review, including whether the additional safeguards and transparency measures in Section 7 can preserve F-TCF's benefits while reducing this risk to an acceptable level.

8.3 Is the ORR rating rule the right foundation for NBR, given perceptions it favors slower boats?

This is a distinct question about the underlying rating rule that generates each boat's baseline VPP and polars, separate from the F-TCF scoring method built on top of it. Some competitors have raised concerns that Newport Bermuda Race results statistically favor slower (and/or older) boats under ORR, and view ORR as a rule shaped by longtime stakeholders — including the CCA — that does not fairly treat newer, faster designs. BROC recognizes this as a significant perceived concern to a portion of the fleet and takes it seriously. Evaluating it fairly, including statistical analysis of results across rating bands and comparison against alternative rating rules, is a substantive undertaking that BROC and its Technical Committee intend to pursue separately from this report.

BROC will address these three questions through its ongoing rules and scoring review and will report to competitors separately on that work. They are noted here because competitors' confidence depends not only on correcting this specific F-TCF error, but on confidence in the underlying methodology as well.

9. Conclusion

The F-TCF error that affected the 2026 Newport Bermuda Race resulted from two independent causes: a procedural gap (human error) that allowed extraneous, competing forecast data to remain active during the official race-morning routing run, and a separate software defect in fleet-mode routing that was discovered and corrected (new version of Expedition released in late June 2026) during the investigation. As shown in Section 5, the human error — not the software defect — was responsible for the great majority of the F-TCF error impact. The proposed corrective actions (along with the already 'fixed' Expedition Bug) are intended and expected to prevent these causes (and foreseeable similar causes) from impacting future races scored with the F-TCF methodology.

Neither cause reflects a technical flaw in the underlying F-TCF methodology itself. In an upcoming report, we will explore comparisons of the 2026 results against those a fixed performance curve scoring (PCS) approach would have produced. Early analysis continues to show that F-TCF scoring is materially fairer to the fleet, particularly in a race prone to the weather conditions and current behavior seen this year.

BROC and its Technical Committee remain fully accountable for the results of the race and regret the disruption and uncertainty the F-TCF error caused competitors. We are committed to implementing the corrective actions described above, to continuing our review — including the broader policy questions in Section 8, and to sharing what we learn with the sailing community and other race organizers.

For competitors or others wishing to replicate the behavior of Expedition and the generation of F-TCF values that occurred during this race, and as described in this report, the following documents are available on the Bermuda Race Sailor Portal.

- [Official GRIBS used for optimal routing / F-TCF generation: HRRR, GFS, RTOFS](#)
- GRIBS inadvertently included in the Friday June 19 routing / F-TCF generation:
[ECMWF](#)
[MERCATOR](#)
- Scratch boat polar files (Expedition format):
Double Handed: [Scowling Dragon US43216](#)
Finisterre: [Gypsy Soul US43262](#)
Gibbs Hill: [Black Jack 100 US43242-1](#)
St. David's: [Temptation Oakcliff US41779](#)
- [Expedition configuration settings used in optimization runs for F-TCF generation](#)

The BROC wishes to thank the US Sailing Offshore Office, Jim Teeters, Stan Honey, and Nick White for their critical contribution to the timely correction of the erroneous data and for their thoughtful engagement in the incident review and corrective actions.

*Bermuda Race Organizing Committee (BROC), BROC Technical Committee, US
Sailing Offshore Office*