A Geometry-Based National Screen for Candidate Aerial-Firefighting Reload Water Sources

Todd O'Hara

TracPlus — Aerial Firefighting Operations & Analytics

Working paper, July 2026 · Configuration geometry-v0.1 · Conterminous United States

Abstract

Aerial firefighting coverage depends less on whether an aircraft can reach an incident than on how much suppressant it can deliver within the response window. For helicopters and fixed-wing scoopers, that quantity is set largely by where usable reload water sits: reload distance governs cycle time, the number of loads flown, and the area over which an asset can sustain delivery. Despite this, no openly available national inventory screens waterbodies for the surface geometry each aircraft type needs. We build one from the U.S. Geological Survey 3D Hydrography Program (3DHP) waterbody layer. The screen is deliberately simple and geometry-only. For each polygon we compute surface area, an approximate maximum length (the long side of the minimum rotated rectangle), an effective width (area divided by length), and a shape-regularity ratio, then apply per-aircraft thresholds held in a versioned configuration file. Across the conterminous United States we screened 3,713,560 waterbodies of at least 1,000 m². Of these, 3,641,020 (98.0%) meet a permissive helicopter dip or snorkel geometry, 56,636 (1.5%) a Fire Boss scooping geometry, and 30,140 (0.8%) a CL-415-class geometry. The output is a reproducible, joinable layer that carries location, dimensions, and per-aircraft status for every source, intended as the geometry input to Standards-of-Cover models of sustained delivery. It screens mapped surface geometry only. Depth, current water level, obstacles, access, and operational approval are out of scope, and the thresholds are screening assumptions rather than validated operational limits.

Keywords: aerial firefighting; Standards of Cover; hydrography; reload sources; spatial screening; 3DHP.

1 The coverage problem

Standards-of-Cover analysis defines an effective response as the resources that can deliver a required effect to an incident within a target time. For aerial firefighting that framing has a specific consequence. A helicopter or scooper does not drop once and finish. It flies a cycle: deliver, return to water, reload, deliver again. Its contribution across a response window is the sum of those loads, so the location of usable water becomes a primary driver of capability.

The effect is large. A helicopter working next to a usable dip pond can complete several cycles in half an hour. Move the same aircraft away from water and it may put down only its first load before the window closes. Reach is necessary, but it is not the whole story: two aircraft with the same base, speed, and initial payload can cover very different ground because the water beneath them differs.

Modelling this nationally means knowing, for any point on the map, where an aircraft of a given class could plausibly reload. That requires a national inventory of waterbodies screened against the surface geometry each aircraft needs: open dip water for helicopters, a long clear run for scoopers. We build and release that screened inventory, and set out how it feeds a reload-cycle and sustained-delivery model. The screen is configuration-driven and reproducible; the per-source layer records location, dimensions, shape quality, and per-aircraft status; and the limitations are stated plainly rather than buried.

1.1 Background

The value of such a layer is easiest to see against the aviation-effectiveness and fleet literature. A RAND assessment of Forest Service large-aircraft requirements framed fleet choice as minimising total social cost, and its cost-minimising portfolios leaned heavily on water-scooping aircraft, conditional on suitable water being close at hand [3]. The Forest Service Aerial Firefighting Use and Effectiveness (AFUE) study, built from thousands of observed drops, found that helicopters and scoopers produced higher rates of effective drops than large airtankers, and that effectiveness falls as fires grow [4]. Both results point the same way. The payoff from rotary and scooping aircraft is real but conditional on nearby water, which is exactly the quantity a reload-source layer makes explicit. Operational response-time doctrine is also well established: federal airtanker contracts set a fifteen-minute reaction standard, and CAL FIRE designs its air-attack network so that a tanker can be over a fire within roughly fifteen to twenty minutes [5]. Coverage has a long modelling tradition in location science, from set-covering to maximal-covering formulations [6], which the sustained-delivery model of Section 5 works toward. On the data side, 3DHP is the successor to the National Hydrography Dataset and the first systematic national remapping of surface water since the 1:24,000-scale topographic era [1].

2 Data

The source is the USGS 3D Hydrography Program (3DHP) waterbody layer. Each polygon carries a stable 3DHP identifier, an optional GNIS name and code, a feature type (river, canal, lake, or ocean/Great Lake), a feature date, and a mapped surface area. We process the conterminous United States extract and keep the source identifiers and names throughout, so every screened record traces back to its origin.

One caveat belongs up front, because the method rests entirely on polygon geometry. 3DHP became the operational program in FY2024, after the NHD was retired on 1 October 2023 [2]. Its database was first populated by migrating existing NHD features into the new data model, and elevation-derived hydrography from lidar and IfSAR is being added basin by basin and released on a quarterly cycle. In most of the country, then, the waterbody geometry we screen is still inherited NHD mapping, much of it from the 1:24,000-scale era, rather than the newer lidar-derived product. Our results are a snapshot of the FY2025 staged product (extract dated 12 January 2026) and will shift as 3DHP is rebuilt.

3DHP records where mapped water exists. It does not say whether a body is deep enough, currently wet, clear of obstacles, or permitted for use. We exclude ocean and Great Lake features, which are not the target of an inland reload screen. Of 5,747,966 CONUS waterbody polygons, 3,713,560 were at least 1,000 m² with valid geometry and were carried forward. We keep every one of them with its classification, so that later work can join observed activity to any mapped body regardless of how it screened.

3 Method

The screen is intentionally simple, explainable, and configurable. It models surface geometry only, and makes no attempt to infer depth, seasonality, obstacles, terrain, or approval. Every threshold lives in a versioned configuration file rather than in the processing code, so the screen can be recalibrated against known operational sources without changing the method.

3.1 Derived geometry

We project each polygon to an equal-area coordinate system and compute four quantities. Surface area (m^2) comes from the projected polygon. Approximate maximum length (m) is the longer side of the polygon's minimum rotated rectangle; we read it as an upper bound on available dimension, not as a proven straight scoop run. Effective width (m) is area divided by length, which we prefer to the rectangle's short side because it is less prone to overstating the usable width of an irregular body. The shape ratio (polygon area over bounding-rectangle area) is kept as a flag for highly irregular shapes rather than as a filter. We also store a representative interior point for each source, which lets downstream work run nearest-source joins without decoding full polygons.

3.2 Classification and thresholds

We screen three aircraft families against surface geometry: helicopters, the Fire Boss, and the CL-415 class.

The thresholds (Table 1) sit deliberately above each aircraft's nominal water-contact distance. A CL-415 scoops a full 6,137-litre load in a water-contact run of roughly 410 m, but a usable site needs much more than the contact length. Descent, the run itself, and a climb-out clear of shoreline obstacles bring the practical requirement to about 1,300 to 1,400 m for a body ringed by 15 m obstacles, over water at least about 2 m deep [7, 8]. The Fire Boss (AT-802F) scoops a smaller load of roughly 3,000 litres in 12 to 15 seconds and needs a shorter run [9]. Our YES thresholds (1,500 m for the CL-415 class, 1,000 m for Fire Boss) sit above these operational minima on purpose, because a mapped polygon dimension does not prove the run is straight, unobstructed, or into wind; the MARGINAL band relaxes length toward the operational minimum. The width thresholds (100 to 150 m) are a lateral manoeuvring and safety margin rather than a published figure. We treat them, like every threshold here, as screening assumptions in a versioned configuration, not as aviation standards. Depth, the other hard scooper requirement, is outside the reach of a geometry screen (Section 6).

Table 1. Screening thresholds, configuration geometry-v0.1. Width is the effective (area ÷ length) width. Thresholds are screening assumptions set above nominal water-contact distances, not certified operational limits.

Aircraft class	Status	Min. length (m)	Min. width (m)
Helicopter (Type 1/2/3, dip / snorkel)	YES	—	15
Fire Boss (AT-802F)	YES	1,000	100
Fire Boss (AT-802F)	MARGINAL	700	75
Large scooper (CL-415 class)	YES	1,500	150
Large scooper (CL-415 class)	MARGINAL	1,000	100

The helicopter screen adds a minimum area of 1,000 m² and applies no length requirement, which keeps small ponds and tanks that may still support bucket or snorkel work. Because a rotated-rectangle length overstates the straight run available on long, sinuous features, river-type polygons are excluded from the scooper classes by default. All classes share a single YES / MARGINAL / NO vocabulary so that results filter and compose the same way.

4 Results

Across the conterminous United States the screen separates a very large inventory of small dip-capable ponds from a much smaller set of scooper-capable waterbodies. The funnel from raw inventory to scooper candidates spans more than two orders of magnitude (Figure 1), and the counts are given in Table 2.

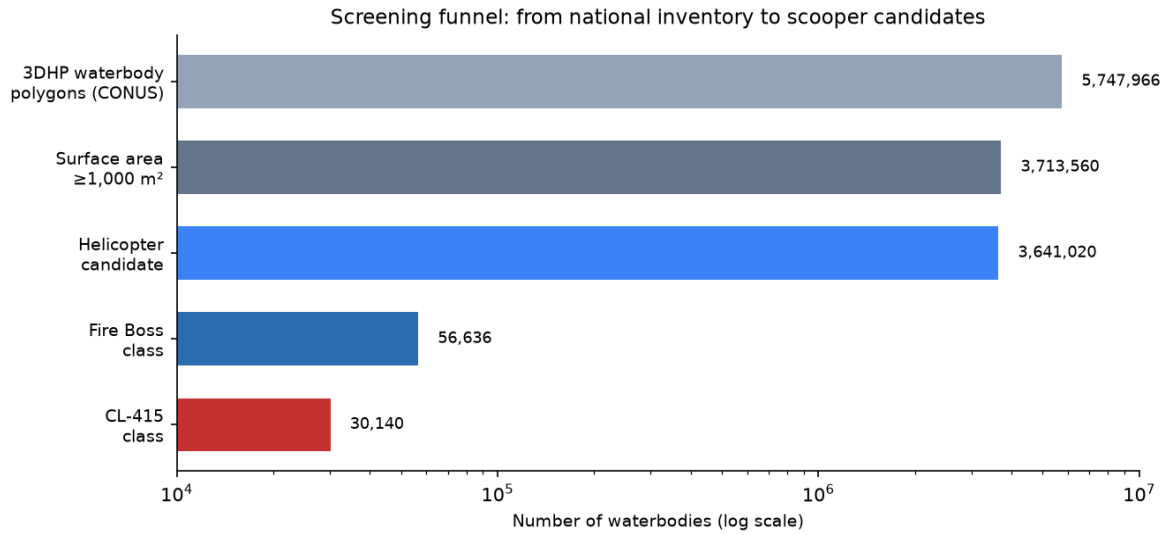

Figure 1. Screening funnel from the national 3DHP inventory to fixed-wing scooper candidates (logarithmic scale). Helicopter-capable bodies dominate the inventory. Scooper-capable bodies are rare, and are geometrically the largest lakes and reservoirs.

Table 2. Screening results, CONUS. Percentages are of the 3,713,560 screened waterbodies. CL-415-usable is a subset of Fire Boss-usable. The two 30,140 entries are not a duplication: CL-415-usable and Fire Boss-YES coincide because the CL-415 MARGINAL threshold (≥ 1,000 m, ≥ 100 m) is identical to the Fire Boss-YES threshold, and no waterbody in this dataset is ≥ 1,500 m long yet only 100 to 150 m wide, so the sets are the same.

Category	Count	Share
Waterbodies screened (≥ 1,000 m²)	3,713,560	100.0%
Helicopter candidate (dip-capable)	3,641,020	98.0%
Fire Boss usable (YES + MARGINAL)	56,636	1.5%
of which Fire Boss YES	30,140	0.8%
CL-415 usable (YES + MARGINAL)	30,140	0.8%
of which CL-415 YES	14,914	0.4%

The screened inventory holds 198,217 km² of mapped inland water surface, excluding ocean and Great Lake features, and is overwhelmingly lakes (3,648,041), with far fewer rivers (56,668) and canals (8,851) clearing the area threshold. The scooper-capable subset is spatially uneven (Figure 2). It is dense across the glaciated northern tier and the reservoir-rich West, and thin through arid and karst regions where few bodies offer a long, wide run.

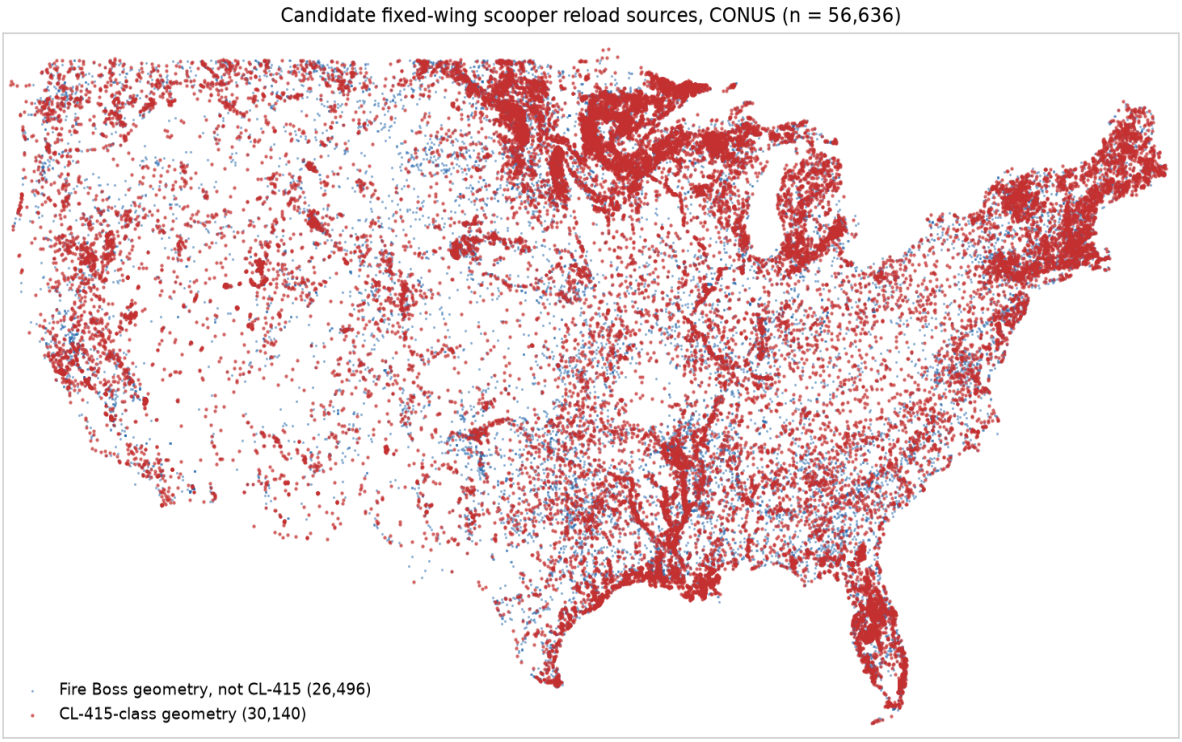

Figure 2. Candidate fixed-wing scooper reload sources across CONUS (n = 56,636). Blue marks bodies that meet Fire Boss geometry but not CL-415 (26,496); red marks CL-415-class bodies (30,140). The two groups are disjoint and sum to the 56,636 scooper-usable total in Table 2. Each point is a source's representative interior point.

5 From screen to response capability

The screened layer is an input, not a result. It supplies the geometry term in a model of sustained delivery, which this paper sets out but does not evaluate. For an aircraft of class c operating near an incident, reload cycle time is governed mainly by the distance to the nearest waterbody whose screened status admits class c, written d_near(c) and supplied by this layer:

$$t_cycle(c) \approx 2 \cdot d_near(c) / v_c + t_reload(c) + t_drop$$

The number of loads deliverable within a response window T, and the cumulative volume, follow from it:

$$loads = \lfloor (T - t_transit) / t_cycle(c) \rfloor \qquad volume = loads \cdot payload_c$$

Swept across the landscape, these give an effective-coverage surface for each asset and aircraft class: the area within which it can deliver at least a target volume inside the window, which is the quantity a Standards-of-Cover assessment actually needs. Because every source carries a representative point, the

nearest-source distance is a fast point query per class, so the model can be computed nationally. Evaluating it against real bases and demand is the subject of separate work.

The same layer supports calibration rather than assumption. Observed reload events, where an aircraft repeatedly descends to and departs from one location, can be matched to screened sources: to confirm which geometries are used in practice, to surface used sources the screen missed, and to retire screened sources that see no activity. Because the layer stores the geometry metrics that drive classification, thresholds can be re-tuned against such evidence without reprocessing the source data.

6 Limitations

This is a screen of mapped surface geometry. It is conservative by design and incomplete by design. The following are out of scope in this version and must not be read into a positive result: water depth and draft; current and seasonal water level; drought and reservoir drawdown; shoreline vegetation; wires, bridges, and other obstacles; boats and surface activity; terrain and approach paths; salinity; and environmental, incident, or landowner restrictions. A mapped length is an upper bound on available dimension, not a proven straight, clear scoop run, which is why scooper thresholds sit well above nominal contact distances and why rivers are excluded from the scooper classes by default. The thresholds themselves are screening assumptions and remain uncalibrated pending comparison against a sample of known helicopter, Fire Boss, and CL-415 sources. Finally, as Section 2 notes, the underlying 3DHP geometry is largely inherited NHD mapping and is being rebuilt, so both the counts and individual polygons will change over time.

Interpretation. A positive result denotes a potential reload source based on mapped surface geometry only. Current water availability, depth, obstacles, permissions, and operational suitability have not been assessed. The layer is a screening aid, not operational approval.

7 Implications

A common, reproducible, national baseline of geometry-screened reload sources changes what aerial firefighting coverage analysis can ask. Coverage has often been reasoned about in terms of reach, meaning whether an asset can arrive in time. A reload-source layer lets the question become sustained delivery: how much can be put on the incident, and over how large an area, before the window closes. That supports concrete decisions about where to base or pre-position rotary and scooper assets, where reload scarcity caps effective response regardless of flight time, and how aircraft mix trades against the water available in an operating area.

It also connects directly to the effectiveness evidence. AFUE found scoopers and helicopters more often effective than large airtankers, and RAND's cost-minimising fleets lean on scoopers where water is close [3, 4]. Both findings are conditional on that nearness of water. A national reload-source layer quantifies the condition, which makes it a practical input to fleet-mix and basing arguments rather than a purely descriptive dataset.

Because the method is geometry-only and configuration-driven, it extends cleanly. Modelled depth from global lake datasets and satellite-derived surface-water permanence can be joined to address the screen's two largest blind spots, whether a body is deep enough and whether it is reliably wet [10, 11]. The same screen applies to any national hydrography, which makes a globally consistent reload-source layer a tractable next step for a fire season that is increasingly international.

8 Conclusion

Effective aerial firefighting coverage depends on the water within reach of the aircraft, not only on the distance to the fire. A deliberately conservative, geometry-only screen of the national 3DHP inventory yields a reproducible, joinable reload-source layer. Across CONUS, 3.64 million waterbodies meet a helicopter dip geometry and 56,636 meet a fixed-wing scooper geometry, out of a screened total of 3.71 million bodies of at least 1,000 m². The screen makes no claim of operational suitability. Its value is a national, recalibratable baseline onto which depth, permanence, observed reload activity, and ultimately a full model of sustained delivery can be added.

References

[1] U.S. Geological Survey. 3D Hydrography Program (3DHP) — Waterbody dataset, FY2025 staged product, conterminous United States. Public domain.

[2] U.S. Geological Survey. 3D Hydrography Program: transition from the National Hydrography Dataset (NHD retired 1 October 2023; 3DHP first operational year FY2024).

[3] RAND Corporation (2012). Air Attack Against Wildfires: Understanding U.S. Forest Service Requirements for Large Aircraft. Monograph MG-1234.

[4] U.S. Forest Service (2020). Aerial Firefighting Use and Effectiveness (AFUE) Report.

[5] National Wildfire Coordinating Group. Standards for Airtanker Operations (PMS 514); and CAL FIRE air-attack program

[6] Church, R. L., & ReVelle, C. (1974). The maximal covering location problem. Papers in Regional Science, 32, 101–118.

[7] Bombardier / Canadair CL-415 product specifications.

[8] National Interagency scooper operations guidance (approved scooper operating plans; minimum water-body dimensions and obstacle-clearance criteria).

[9] Air Tractor / Fire Boss, LLC. AT-802F Fire Boss specifications.

[10] Messager, M. L., Lehner, B., Grill, G., Nedeva, I., & Schmitt, O. (2016). Estimating the volume and age of water stored in global lakes using a geo-statistical approach. Nature Communications, 7, 13603 (HydroLAKES).

[11] Pekel, J.-F., Cottam, A., Gorelick, N., & Belward, A. S. (2016). High-resolution mapping of global surface water and its long-term changes. Nature, 540, 418–422.

Data and reproducibility.

Derived from public-domain USGS 3DHP data with a configuration-driven Python screen (thresholds versioned as geometry-v0.1). Outputs are written as GeoParquet with a representative interior point per source for spatial joins. Prepared as a working paper by TracPlus.

Disclaimer: not for operational use.

This document and the associated dataset are a geometry-based screening exercise, provided for research and informational purposes only. They are not an operational product and must not be used for flight planning, dispatch, navigation, water-source selection, or any in-flight, tactical, or on-incident decision. A positive result does not mean a waterbody is safe, legal, accessible, deep enough, currently wet, free of obstacles, or approved for firefighting use; none of those factors were assessed. The underlying mapping is third-party data that may be incomplete, inaccurate, or out of date, and everything here is provided "as is", without warranty of any kind, express or implied, including any warranty of accuracy, completeness, or fitness for a particular purpose. The operational suitability of any water source must be independently verified with the relevant aviation and land-management authorities before any use.