SOL ENGINEERING SERVICES

The CF-JET Project in Plain Language

A summary of the project, the key concepts and techniques with examples from the research, and answers to frequently asked questions

Prepared for program managers, engineers, university co-investigators, industry partners and other readers who are new to the technical material

Classification	Unclassified. Prepared from open-source information.
Date	27 November 2026

The Project in Brief

CF-JET is a proposed cooling plate for power electronics that produce a great deal of heat in a small space. Hybrid-electric Army vehicles convert 100 to 400 kW of electrical power, and at conversion losses of 2.15 to 3 percent a single stage gives off 2 to 12 kW of heat (computed). The plate that removes this heat must work with hot coolant, a small pressure budget and little room. CF-JET answers with jets of liquid that strike a carbon fiber plate, return ports placed between the jets, a body made by 3D printing and a thin electroplated metal face on the coolant side.

This summary explains the project for readers without an engineering background. It describes the project and its design point, explains the main concepts and techniques with examples drawn from the research behind it, and answers the questions that readers ask most often. Numbers come from that research. A value marked as computed was calculated from published figures, and a value marked as an assumption is a design choice of the proposal.

- **The target is a measurable number.** The cooler must remove 1.5 kW from a 60 mm by 60 mm area while the plate stays about 30 °C warmer than the coolant (Table 1).
- **The target is ambitious and falls inside the published range.** It converts to 72 mm^2·K/W, which is 1.4 to 2.8 times better than conventional finned and swaged-tube plates and within the 9 to 158 mm^2·K/W that published jet and microchannel results have reached.
- **Three parts of the design have no published measurement at full size.** These are a printed jet plate with interleaved returns near 36 cm^2, a carbon fiber insert that conducts 50 W/m·K through its thickness, and the combination of the two.
- **AI is proposed to reduce the number of simulations and prototypes.** The methods work on synthetic problems and published cases, and they help only when each prediction reports how uncertain it is.
- **Testing, funding and rules decide whether the design reaches hardware.** A written test definition can change a measured result by more than the instruments do, and the funding route decides who can be paid and what security rules apply.

1. What the Project Is and Why It Matters

1.1 The project designs a cooler for electronics that make more heat than conventional coolers remove

Power electronics convert electricity from one form to another. An inverter in a hybrid-electric combat vehicle changes battery power into the alternating current that powers a motor, and every conversion loses a few percent of the power as heat. Silicon carbide (SiC) devices handle higher power and higher temperature than silicon devices, and they pack that heat into a smaller area. A cold plate is mounted against the electronics and removes the heat in flowing liquid. When the cold plate falls short, the electronics must be derated, enlarged or shut down.

CF-JET is the name of the effort proposed by SOL Engineering Services, LLC to build a better cold plate for these loads. The research behind it examined the heat the electronics produce, the materials, the printing process, the AI design tools, the testing and the funding paths. Figure 1 shows how the pieces relate.

Figure 1. The CF-JET project at a glance

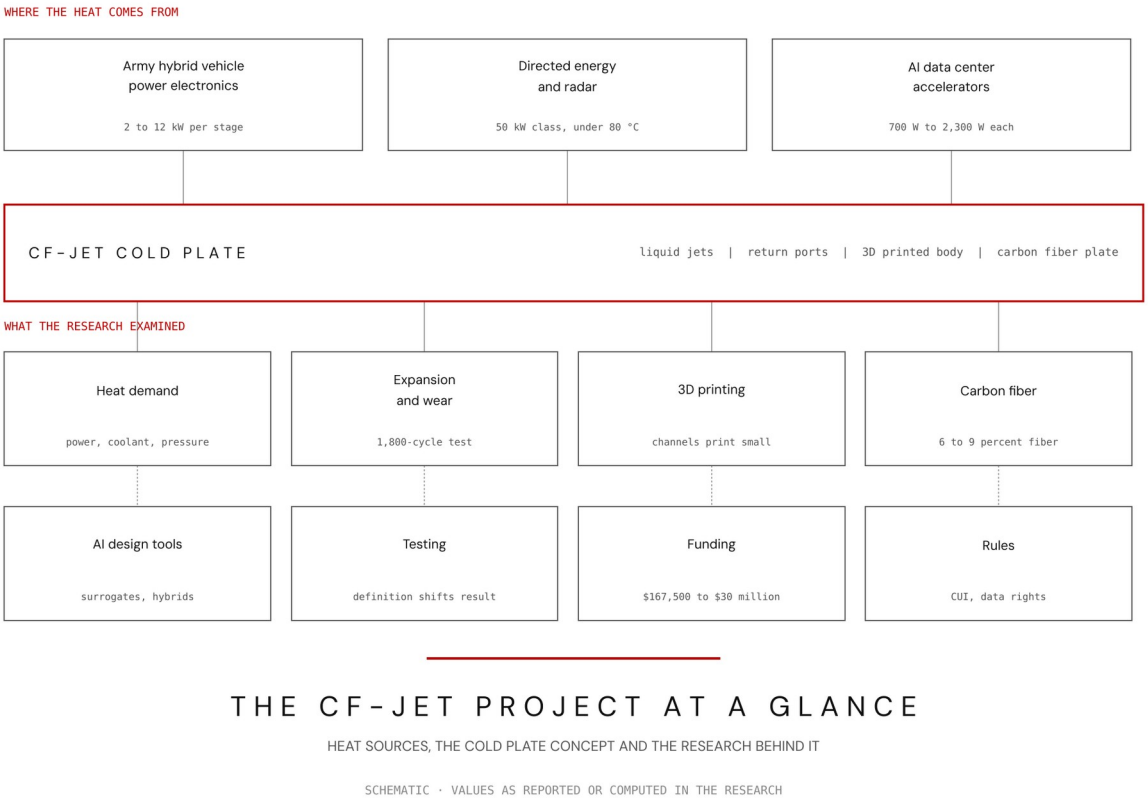

Schematic. Boxes show where the heat comes from, the cold plate concept and the areas of research that support it. Values are those reported or computed in the research.

1.2 The design point states the job in numbers that anyone can check

Every cooler is judged against a design point, a set of conditions that the cooler must meet. The CF-JET design point appears in Table 1 with a plain-language reading of each value. The 41.7 W/cm² average heat flux and the 30 K temperature rise at full load are computed from the other values, and the rest are design assumptions of the proposal.

Table 1. The CF-JET design point in everyday terms

Quantity	Value	Plain meaning	Evidence
Heat to remove	1.5 kW	About the output of a household space heater, delivered to an area smaller than a playing card	Design assumption of the proposal
Cooled area	60 mm by 60 mm (36 cm²)	A square about 6 cm on each side	Design assumption of the proposal
Average heat flux	41.7 W/cm²	Heat leaving each square centimeter of the plate	Computed from 1,500 W over 36 cm²
Coolant	50/50 water-glycol at 65 °C inlet	The same kind of mixture used in car cooling systems, entering hot	Design assumption of the proposal
Flow rate	4 to 8 L/min	Two to four large kitchen pitchers of liquid every minute	Design assumption; pitcher comparison computed
Pressure drop	30 kPa or less	The push the pump must supply to move coolant through the plate	Design assumption of the proposal

Quantity	Value	Plain meaning	Evidence
Thermal resistance	0.02 K/W	Each watt of heat raises the plate temperature 0.02 K over the coolant temperature	Design assumption of the proposal
Temperature rise at full load	30 K	The plate surface stays about 30 °C warmer than the coolant at 1.5 kW	Computed from 1,500 W times 0.02 K/W
Resistance per unit area	72 mm²·K/W	The same target restated so that coolers of different sizes can be compared	Computed from 0.02 K/W times 3,600 mm²

Computed entries use 1,500 W, 36 cm² and 0.02 K/W as stated. The pitcher comparison assumes a pitcher of about 2 L. Interpretation: the design point is a proposal assumption that the research tests against published data.

1.3 Hot coolant and a small pressure budget make the job harder than it looks

A cooler works because the coolant is colder than the part. When the coolant is already hot, less temperature room remains. Public Army documents give 65 °C as the current coolant inlet temperature for electronics and 80 to 105 °C as the future range, and the vehicle loop offers only 2 to 5 psi, or 14 to 34 kPa, of pressure drop to the electronics. In one published example, a cold plate that removes 1,340 W per module with 25 °C coolant keeps only about 47 percent of that capacity with 105 °C coolant. These conditions explain why the CF-JET design point starts at 65 °C and limits pressure drop to 30 kPa.

2. How the Cooler Works

2.1 Liquid jets strike the hot plate and spent liquid leaves through ports between the jets

In a conventional cold plate, liquid flows along channels beneath the part, much as water flows along a pipe. In a jet-impingement plate, many small nozzles aim liquid straight at the hot surface. Fresh, cool liquid reaches the surface at high speed, and the heat leaves faster than it does in a slow channel flow. A review cited in the research found local heat transfer in jets up to three times that of confined channel flow at the same maximum velocity.

A problem appears when many jets share one surface. Liquid that has already struck the surface must travel to an exit, and on its way it flows across the jets farther along. This cross-stream movement weakens the heat transfer of those jets. An air-based correlation predicts that a row of jets with ten rows upstream keeps 65 percent of its heat transfer in an in-line array and 55 percent in a staggered array (computed). CF-JET places return ports between the jets so that spent liquid leaves nearby, and simulations in the research report lower average temperatures and better uniformity with such ports. Figure 2 shows the layers and the liquid paths.

Figure 2. Cross-section of the CF-JET concept

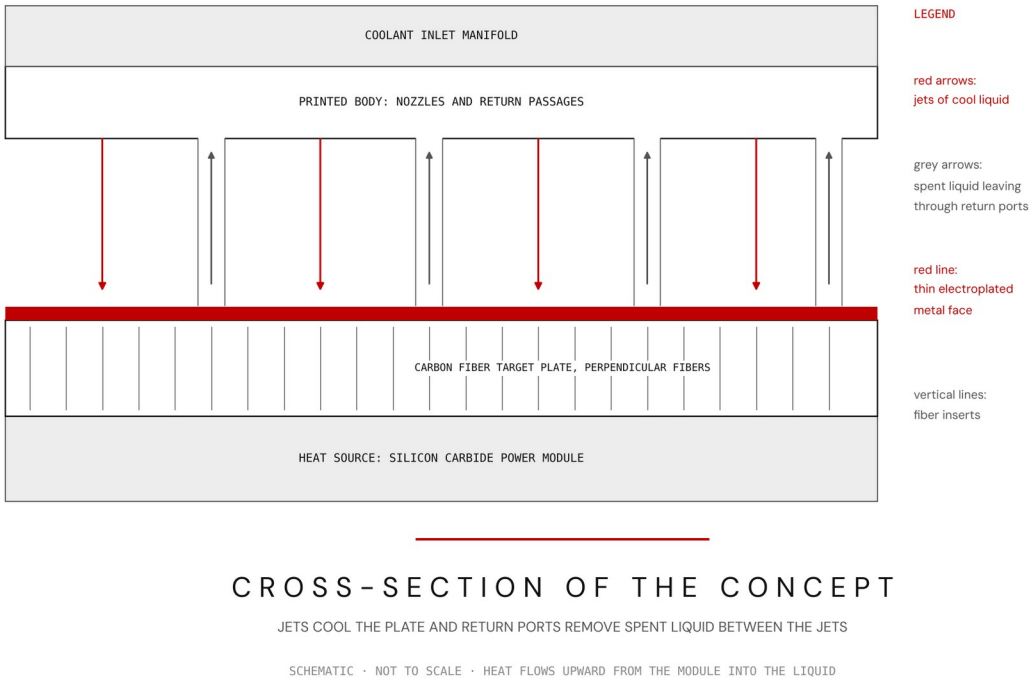

Schematic, not to scale. Arrows show liquid entering through jets and leaving through return ports. Fibers in the target plate are drawn perpendicular to the plate.

2.2 Each of the three main parts does one job

- **The printed body** holds the jets and return passages in one piece. Additive manufacturing can form manifold and nozzle geometry that machining cannot reach easily.
- **The carbon fiber target plate** receives the jets and moves heat from the electronics to the liquid. Pitch fiber conducts 140 to 1,000 W/m·K along its length by vendor value, so fibers standing perpendicular to the plate give heat a direct path through it.
- **The thin electroplated face** seals the coolant side of the plate. A copper layer of 10 to 25 µm adds only 0.03 to 0.09 percent to the thermal resistance budget (computed).

3. Key Concepts Explained with Examples

3.1 Heat flux is heat per area, and small areas inside the module reach far higher values than the average

Heat flux is the amount of heat power passing through each unit of area, written in W/cm². The CF-JET plate averages 41.7 W/cm². The average hides the hot spots. If twelve silicon carbide dies of 0.25 cm² each share the 1.5 kW load, each die passes 500 W/cm², and with six dies each passes 1,000 W/cm² (computed). The highest measured flux in the sources reviewed is about 170 W/cm² at the National Renewable Energy Laboratory (NREL) in 2024, and the highest reported die-level flux exceeds 1,000 W/cm².

3.2 Thermal resistance and the 72 mm²·K/W budget turn the cooling job into a spending plan

Thermal resistance states how many degrees the temperature rises for each watt of heat, in K/W. At 0.02 K/W and 1,500 W the plate stays 30 K warmer than the coolant (computed). Dividing by the area gives 72 mm²·K/W, a form that lets coolers of different sizes be compared. Every part of the heat path spends some of that budget: the liquid film at the surface, the plate itself, the face and the joints. A part that spends more than its share leaves too little for the others, and a published example shows how quickly this happens (Figure 3).

3.3 A thin plastic skin could spend more than the entire budget

Carbon fiber in epoxy conducts heat poorly through its thickness. Standard laminates reach only 0.76 to 2.8 W/m·K, so the research asks what the design needs. A fiber conductivity estimate shows that fiber standing perpendicular to the plate must fill roughly 6 to 9 percent of the volume to reach 50 W/m·K. At 36 W/m·K a 1 mm insert spends 28 mm²·K/W of the budget and a 2 mm insert spends 55 mm²·K/W (computed). A 50 µm skin of unfilled plastic on the jet side would spend 172 mm²·K/W, which exceeds the whole 72 mm²·K/W by itself (computed). A copper face of 10 to 25 µm spends almost nothing. This example is why the design uses an electroplated metal face and why the carbon fiber insert is the first item to test.

Figure 3. Share of the 72 mm²·K/W budget spent by candidate plate features

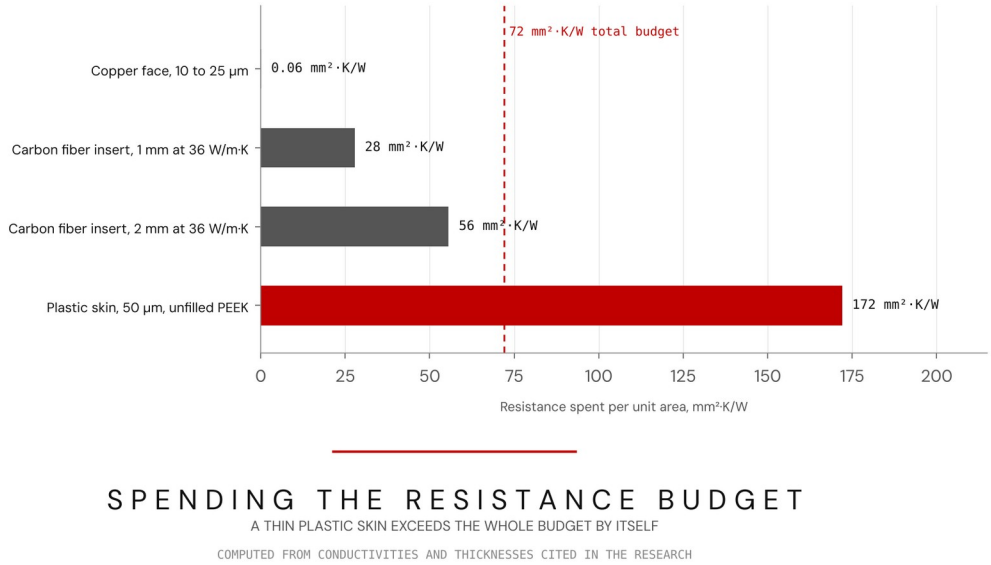

Computed from the conductivities and thicknesses cited in the research. The plastic skin value uses the conductivity of unfilled PEEK, and the insert values assume 36 W/m·K. The copper face is drawn at its upper estimate of 0.0625 mm²·K/W.

3.4 Jet impingement and crossflow decide how many jets help

A shower head aimed at a window sprays a few spots of water, and the water that flows down the glass weakens the spray farther down. A jet array behaves in the same way on a hot plate. The research recommends spacing the return ports so that the ratio of crossflow to jet flow stays at about 0.3 or lower, and it recommends measuring the weakening on a sub-array that uses the real water-glycol coolant. For the design point, about 100 to 150 jets on 3,600 mm² imply a jet spacing of 4.9 to 6.0 mm (computed).

3.5 Thermal expansion mismatch cracks joints that heat and cool repeatedly

Every material grows when heated, and different materials grow by different amounts. Copper expands at 16.5 to 17 ppm/K, aluminum at 22.2 to 25 ppm/K and silicon carbide dies at 3.4 to 4 ppm/K. Copper and aluminum therefore expand four to eight times as much as the die. When these layers are bonded together and cycled hot and cold, the joints between them stretch and shear, and solder fatigue and bond-wire lift-off follow. In NREL modeling the strain energy per cycle is 0.18 MPa for aluminum silicon carbide, 1.94 MPa for copper and 3.93 MPa for aluminum. A tested aluminum silicon carbide heat sink survived 1,800 cycles from −40 to 150 °C without failure, and its authors report a life improvement of at least 18 times. This is the reason the CF-JET plan measures the expansion of the finished target plate, including the plated face.

3.6 Printed channels come out smaller than drawn, and small errors multiply

A 3D printer builds metal parts layer by layer, and the finished channels differ from the drawing. In one study of Inconel 718, vertical channels drawn at 1 mm measured 988 µm and horizontal ones measured 805 µm, and channels of 500 µm were partly blocked at build angles of 22.5° or less. A nozzle that is 19.5 percent undersize raises the ideal pressure drop by a factor of 2.4 (computed), because pressure drop rises steeply as the opening narrows. The research therefore recommends building jet orifices with their axes along the build direction and measuring the finished diameters by computed tomography, which is an X-ray scan that builds a three-dimensional image. Jets under about 0.56 mm in diameter exceed the 30 kPa limit at 8 L/min with the fewest jets and a discharge coefficient of 0.7 (computed).

3.7 A surrogate model stands in for a slow simulation

A computer simulation of liquid flow and heat transfer can take hours for each design. A surrogate model is a fast mathematical stand-in trained on the results of a modest number of simulations, much as a recipe card summarizes many trials in the kitchen. Published surrogates for jet and cold plate hardware match their simulations within 0.2 to 2 percent after 20 to 100 simulations. In a benchmark computed for the project on a synthetic problem, a Gaussian process surrogate placed 98.6 percent of test points within ±5 percent after 200 training points and returned a prediction with its uncertainty in 0.12 ms. The benchmark tests the software and the number of training points needed, and it makes no statement about CF-JET hardware.

3.8 The design space is far larger than the number of prototypes anyone can build

A published NREL program on jet-impingement heat exchangers received $1.6 million over fiscal years 2011 to 2013 and had characterized one prototype by March 2013, when it approved the second and third. Nine design variables at three levels each give 19,683 combinations, which is 6,561 times the three prototypes of that program (computed). Hardware alone cannot explore such a space, and simulation and surrogates must do most of the exploring (Figure 4).

Figure 4. Design combinations compared with prototypes in a published program

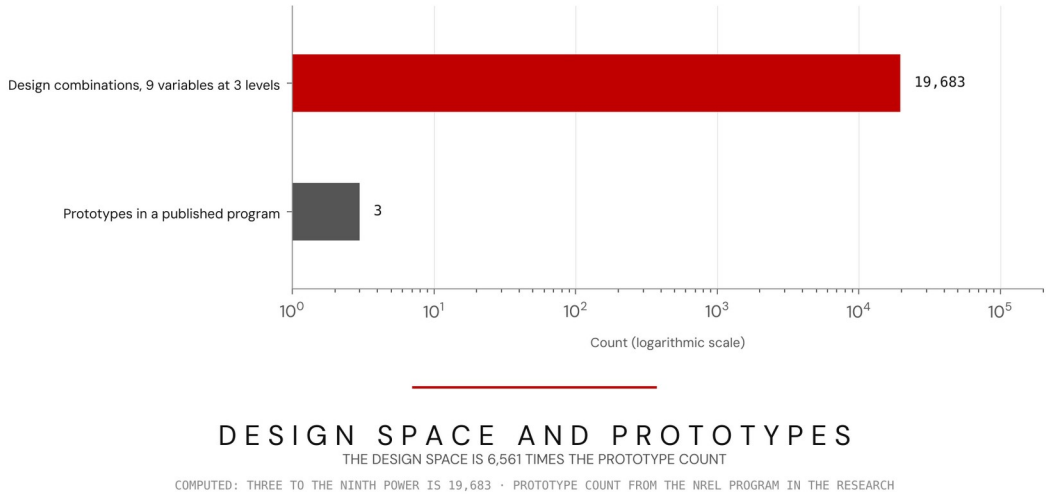

Computed: nine variables at three levels give 3 to the ninth power, or 19,683, combinations. The prototype count is from the NREL program described in the research. The scale is logarithmic.

3.9 Multi-fidelity learning mixes cheap sketches with a few expensive photographs

A coarse simulation is quick and approximate, and a fine simulation is slow and accurate. Multi-fidelity learning uses many coarse cases and a few fine ones and lets the model learn how the two differ. Two preprints report 57 percent of wall time saved in a burner design campaign, and a cost model with a 126-to-1 cell-count ratio places a campaign of 120 coarse and 20 to 80 fine simulations at 10.5 to 40.5 percent of the cost of 200 fine simulations (computed).

3.10 Uncertainty intervals are useful only when they are calibrated

A weather forecast that states a 70 percent chance of rain should be right about 70 percent of the time. A prediction with a 95 percent interval should contain the true value 95 percent of the time. In the project benchmarks, Gaussian process intervals with a nominal 95 percent level covered 95.8 to 99.1 percent of test points, and ensembles of five neural networks covered only 52.8 to 75.1 percent. Published deep learning studies report that calibration worsens outside the training data. A tool that is confident and wrong misleads an engineer more than a tool that admits a wide interval.

3.11 A hybrid of physics rules and learning extrapolates better than either alone

Physics supplies rules that hold everywhere, such as energy balance, and data supplies the corrections where the rules are approximate. A hybrid model adds a learned correction to a codified physics rule. On a synthetic benchmark outside its training region, the hybrid predicted thermal resistance with 3.1 percent mean absolute error, against 7.0 percent for a purely learned Gaussian process, 15.1 percent for an ensemble of five networks and 15.7 percent for the rule alone. The benchmark rule error is smooth, which favors a hybrid, and the hybrid intervals covered only 70.9 percent of far-region observations. The method needs validation on hardware before anyone relies on it.

3.12 Error budgets combine independent errors, and the combined error sets the safety margin

When two independent errors act together, they combine as the sides of a right triangle combine, by the root of the sum of squares. A surrogate error of 5 percent and a simulation error of 15 percent give 15.8 percent in combination (computed). A design that the surrogate predicts to meet 72 mm²·K/W then needs a predicted margin of about 10.2 percent when conduction and spreading use 18 mm²·K/W (computed). The simulation error itself is large: a review of turbulence models for jet impingement reports local errors of 15 to 30 percent for the better variants, and a 2023 air benchmark finds the k-ω SST model within about 2 to 6 percent at the stagnation point.

3.13 The definition of the test often changes the answer more than the instruments do

Instruments alone can verify the resistance of a 0.02 K/W cold plate at 1.5 kW to about ±3.0 percent (computed). Published rig studies show that the choice of case temperature and the joining pressure shift the same quantity by 3.3 to 42.9 percent. At the uncertainty of one university rig, a measured 0.0195 K/W has a 31.9 percent chance of exceeding 0.02 K/W, and at the instrument-only level the chance is 4.8 percent (computed). The research therefore calls for a written test definition, a measurement of the joint between the plate and the heat source, and an acceptance limit with guard band. Comparing predictions with measurements to tell 95 percent coverage from 85 percent coverage needs 46 validation points (computed).

4. What the Research Found

4.1 The findings fall into nine topics, and each has its own evidence

Table 2 summarizes the findings by topic with the type of evidence behind each. Most numbers come from published, peer-reviewed or Government sources. Benchmarks computed for the project test software on synthetic problems and make no claim about CF-JET hardware.

Table 2. Findings at a glance

Topic	Finding	Meaning	Evidence
Heat demand	Hybrid Army vehicles convert 100 to 400 kW of power, and one conversion stage gives off 2 to 12 kW of heat. Coolant temperature and heat load both set the limit.	The cooler must work with hot coolant and a small pressure budget.	Public Army briefings and published studies; heat values computed
Expansion and wear	Copper expands four to eight times as much as silicon carbide. An aluminum silicon carbide plate cycled 1,800 times between −40 and 150 °C without failure.	Matching how fast materials grow protects the electronics from cracking.	Peer-reviewed modeling and one tested heat sink
Conventional coolers	Finned and swaged-tube cold plates reach about 100 to 200 mm²·K/W. Published jet and microchannel results fall between 9 and 158 mm²·K/W.	The CF-JET target of 72 mm²·K/W is aggressive against conventional plates and inside the published range.	Patent background and peer-reviewed studies; ratios computed
3D printing	The best printed cooler measured reached 21 to 35 mm²·K/W on 1 cm². CF-JET needs 36 times that area. Printed channels come out smaller than designed.	Printing is proven for small coolers and unproven at the CF-JET size.	Peer-reviewed measurement; scale-up computed
Carbon fiber plate	Fiber must stand perpendicular to the plate at about 6 to 9 percent of the volume to conduct 50 W/m·K. No published measurement of such an insert was found.	The plate is the least proven part of the design.	Peer-reviewed and vendor data; requirement computed

Topic	Finding	Meaning	Evidence
AI design tools	Published fast stand-in models match simulations within 0.2 to 2 percent after 20 to 100 simulations. A hybrid physics-and-learning model reached 3.1 percent error outside its training region on a synthetic test.	AI can cut the number of simulations and prototypes if it reports its own uncertainty.	Published studies and benchmarks computed for the project
Testing	Instruments alone can verify the resistance to about ±3.0 percent, and the test definition alone can shift the result by 3.3 to 42.9 percent.	A written test definition matters more than the thermocouples.	Standards and published rig studies; uncertainty computed
Market	Estimates of the 2024 market for data center liquid cooling range from $0.87 billion to $4.97 billion. Accelerator power rose from 700 W to 2,300 W in four years.	Commercial demand sets price and performance expectations.	Analyst estimates, trade reports and manufacturer data
Program and rules	Funding ceilings range from $167,500 to $30 million. Several university programs fund the institution only. CUI work needs 110 security requirements.	The funding route decides who can be paid and what rules apply.	Government solicitations; law firm and trade commentary

Entries restate findings from the research in short form. Computed values are marked as such in the text. Interpretation: the meanings are plain-language readings of the findings.

4.2 Demand for cooling is growing quickly in vehicles, directed energy and data centers

Army vehicle programs publish power levels and little thermal detail, and the Ground Vehicle Systems Center lists a power density target of 12 kW/L against 3 kW/L today. The commercial side moves faster. Reported accelerator power rose 3.3-fold in four years, from 700 W for the H100 in 2022 to 2,300 W for Vera Rubin in 2026 (computed), and an analyst projects leading-edge parts beyond 4,000 W in 2029. At the CF-JET resistance of 0.02 K/W, those powers give coolant-to-case rises of 46 K and 80 K (computed, Figure 5).

Figure 5. Accelerator power compared with the CF-JET design power

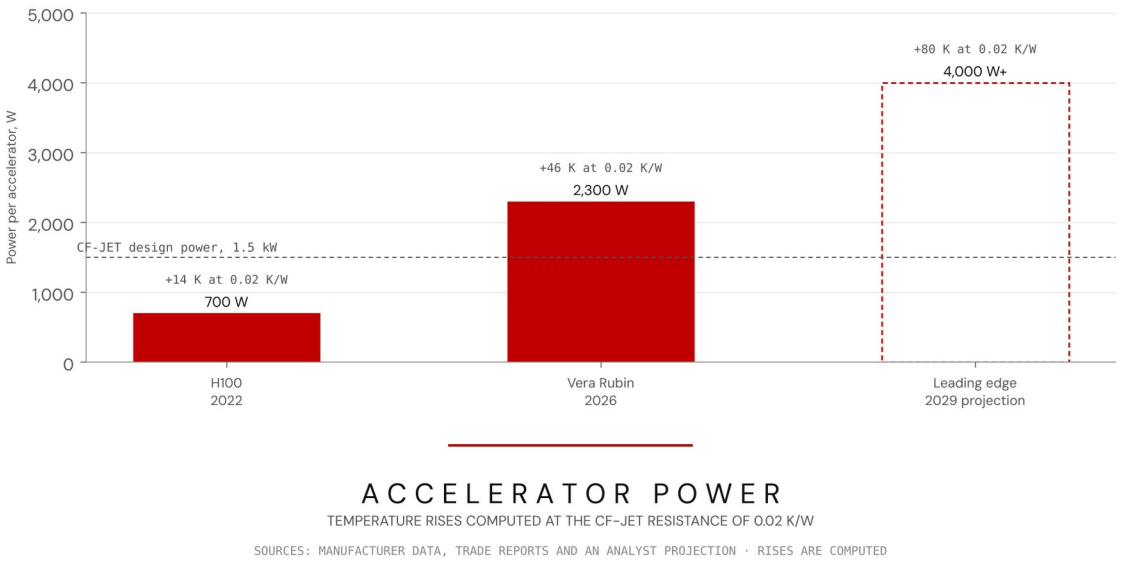

Values: H100 in 2022, 700 W; Vera Rubin in 2026, 2,300 W; leading-edge parts in 2029, over 4,000 W as an analyst projection.

Temperature rises are computed at 0.02 K/W. The dashed line marks the CF-JET design power of 1.5 kW.

4.3 Vendors claim more than the CF-JET design point, and none of the claims has independent test evidence

CoolIT reports a 4,000 W single-phase cold plate with resistance under 0.009 °C/W, and Accelsius reports 4,500 W on a thermal test vehicle. The CF-JET average flux of 41.7 W/cm² compares with a vendor claim over 300 W/cm² and a Department of Energy module target over 400 W/cm². These are vendor claims that the research could not verify. The CF-JET proposal targets military conditions, such as hot coolant and low pressure drop, in a printed form.

4.4 The market estimates differ widely because each counts something different

Analyst estimates of the 2024 market for data center liquid cooling range from $0.87 billion to $4.97 billion, a 5.7-fold spread (computed). One estimate counts investment, another counts manufacturer revenue and others count solutions and services. Two sources put cold plates alone near $0.3 billion in 2024. Large companies are buying cooling suppliers: Eaton agreed to buy Boyd Thermal at $9.5 billion, Ecolab to buy CoolIT Systems at $4.75 billion, and Schneider Electric bought a 75 percent interest in Motivair for $850 million.

4.5 Funding and rules differ by vehicle, and several university programs fund institutions only

Transition funding ceilings range from $167,500 to $30 million, and the Small Business Administration guideline for a Phase II award is $2.24 million. The 2026 pilot program for historically Black colleges and universities and other minority-serving institutions funds institutions only and bars subawards and equipment. A small business reaches those institutions through an STTR award, a subaward under a broad agency announcement or a cooperative agreement with an ARL laboratory. Fundamental research rules cover basic research anywhere and applied research on campus, and printed hardware falls outside them. Work with controlled unclassified information (CUI) must meet the 110 security requirements of NIST SP 800-171 whatever the status of the cybersecurity certification program.

5. What Remains Unknown

5.1 The central design claims lack full-size measurement, and the gaps are stated in each area

The research states its gaps with its findings. Table 3 lists the main questions, what the sources establish and what they leave open. Filling these gaps is the purpose of the first hardware tests the research recommends: a single jet in the working fluid, a carbon fiber insert measured at 1 to 3 mm thickness with the plated face attached, and an array with the return-port arrangement of the full plate.

Table 3. Open questions and the evidence behind them

Question	What is known	What remains unknown	Evidence
Does the full plate reach 0.02 K/W?	The target converts to 72 mm²·K/W, inside the 9 to 158 mm²·K/W published range.	No measurement of a printed jet array with interleaved returns near this size was found.	Interpretation of published data
Can fiber inserts conduct 50 W/m·K?	Vertically aligned pitch fiber in thin elastomer reaches 35.2 to 51.9 W/m·K.	No measurement of a millimeter-thick insert at 6 to 9 percent fiber was found.	Peer-reviewed measurement; gap stated

Question	What is known	What remains unknown	Evidence
Does expansion matching improve life in this design?	Modeling favors aluminum silicon carbide by 11 to 22 times in strain energy, and one test showed at least 18 times life.	No published life test varies only the cold plate expansion in a silicon carbide module.	Peer-reviewed modeling; gap stated
Do the AI methods work on a real cooler?	Methods work on synthetic problems and on smooth, low-dimensional published cases.	No study covers a jet array with 100 to 150 jets, interleaved returns and a composite plate.	Benchmarks computed for the project; gap stated
What do the Army requirements say about the cooler?	Public documents give power density and coolant temperature targets.	They contain no module-level heat flux, resistance, flow or pressure drop requirement.	Government documents; gap stated
Which rules apply to the university work?	Fundamental research rules protect publication of basic research on campus.	No source classifies a cold plate for export control or states proof-of-ownership requirements for software baselines.	Regulation and guidance; gap stated

Entries are drawn from the evidence-gap statements in the research. Interpretation: a gap means that no source retrieved addressed the question, and a source that was not retrieved may exist.

6. Frequently Asked Questions

6.1 Questions about the project

What is CF-JET? CF-JET is a proposed liquid cold plate that combines jets of coolant, return ports between the jets, a 3D printed body, a carbon fiber target plate and a thin electroplated metal face. SOL Engineering Services, LLC proposes it in response to Army Research Laboratory announcements.

Who would use it? Army ground vehicle programs that need compact power electronics are the primary application examined in the research. The same heat problem appears in directed energy systems, radar modules and data center accelerators, and the research documents a large commercial market for data center cooling.

What is a design point? A design point is the set of conditions a design must meet. For CF-JET it is 1.5 kW of heat on a 60 mm by 60 mm area, 50/50 water-glycol coolant at 65 °C, a flow of 4 to 8 L/min, a pressure drop of 30 kPa or less and a thermal resistance of 0.02 K/W.

Does the cooler exist today? The cooler is a concept. SOL's methods and software are described as proposed or demonstrated by SOL, and the research found no independent publication on them. No peer-reviewed measurement of a printed jet plate with interleaved returns near this size was retrieved.

Why does the design point use coolant at 65 °C? Public Army documents give 65 °C as the current inlet temperature for vehicle electronics and 80 to 105 °C as the future range. Hot coolant leaves less temperature room, so a cooler that works at 65 °C is a realistic test of the design.

6.2 Questions about the technology

Why use jets of liquid? Jets deliver fresh, cool liquid straight to the hot surface and reach the highest local heat transfer of the common single-phase methods. The average over a whole plate is far lower than the local peak, and crossflow weakens downstream jets, so the design adds return ports.

Why add return ports between the jets? Spent liquid must leave the surface. Return ports let it leave close to where it landed, which reduces the crossflow that weakens other jets. Simulations in the research

report lower average temperatures, better uniformity and modest pressure drops with such ports, and a printed manifold design lowered pressure drop by 37.9 to 70.2 percent against a reference microchannel sink.

Why 3D print the body? Printing forms manifolds, nozzles and return passages in one piece, and the research found support for one-piece printed metal bodies. The limits are accuracy and size: printed channels come out smaller than designed, and the best measured printed cooler covered 1 cm², which is one thirty-sixth of the CF-JET area.

Why use carbon fiber? Pitch carbon fiber conducts heat along its length at up to several hundred W/m·K and expands very little. These properties help the plate move heat, provided the fibers stand perpendicular to the plate. No published measurement of a millimeter-thick insert at the needed fiber fraction was found, so the plate needs early testing.

Why plate the coolant face with metal? A metal face of 10 to 25 µm seals the coolant side and adds almost nothing to the thermal resistance. A 50 µm plastic skin would spend more than the whole 72 mm²·K/W budget. No erosion rate for a polymer under a slow liquid jet was found, so the face is a design safeguard that testing must confirm.

How heavy would it be? The research gives no CF-JET mass. A 60 mm by 60 mm by 3 mm plate has a mass of 96 g in copper and 29 g in aluminum (computed). No public source found links cold plate mass to thermal resistance for vehicle or aerospace use, so the research recommends reporting conductance per unit mass alongside pressure drop.

Is a resistance of 0.02 K/W realistic? The target converts to 72 mm²·K/W. Conventional finned and swaged-tube plates reach about 100 to 200 mm²·K/W, and published jet and microchannel results fall between 9 and 158 mm²·K/W. The target is therefore aggressive against conventional plates and inside the range that published results have reached on smaller areas.

6.3 Questions about the AI methods

What does AI do in this project? AI methods learn from simulation and test data how design choices affect resistance and pressure drop, so that engineers can search the design space without executing a simulation for every option. The research proposes surrogate models, multi-fidelity learning and a hybrid of physics rules with learned corrections.

Why not simply train a large neural network? Hardware programs produce a handful of test articles, and a large network needs far more data. The Army research announcement ARL-BAA-0038 states the problem directly: the Army has limited data to train or adapt AI systems and expects domain-expert rules and uncertainty-aware methods to help.

What is a surrogate model? A surrogate is a fast stand-in for a slow simulation. It is trained on a modest number of simulation results and then predicts new cases in a fraction of a millisecond. Published surrogates match their simulations within 0.2 to 2 percent on smooth, low-dimensional problems.

How can an engineer trust a prediction from AI? Trust comes from checks that the research proposes: an uncertainty interval with verified coverage, a detector that warns when a design lies outside the data, a guardrail that rejects predictions with wide intervals and a decision record that an engineer signs. In the synthetic benchmark, the Mahalanobis distance detector separated far-region points with an area under the curve of 0.999, and an interval-width guardrail cut far-region errors over 10 percent from 22.9 to 6.0 percent of points.

Does AI replace testing? Testing remains required. Every surrogate prediction must be compared with measurements, and the research computes that telling 95 percent coverage from 85 percent coverage needs 46 validation points. The AI methods reduce the number of simulations and prototypes, and hardware results confirm or reject their predictions.

6.4 Questions about testing and quality

How will the cooler be tested? The research recommends first measuring one jet in the working fluid to find the simulation error, then testing a carbon fiber insert with the plated face attached, then an array with the return-port arrangement of the full plate. Each test needs a written definition of the case temperature, a measured joint resistance and a guard-banded acceptance limit.

Why do two laboratories report different numbers for the same cooler? The definition of case temperature and the pressure that joins the cooler to the heat source change the answer. Published rig studies show shifts of 3.3 to 42.9 percent from these choices alone, and one rig reported 34.2 percent higher resistance at 10 psi than at 15 psi.

Which university laboratories could help? Auburn documents an X-ray computed tomography system. Calhoun, Mississippi Gulf Coast and East Mississippi document additive manufacturing laboratories, and Alabama A&M, Tuskegee and Jackson State document printing and materials equipment. No retrieved source documents a flow loop, laser flash or thermal cycling equipment at an HBCU or community college in either state.

6.5 Questions about funding, partners and rules

Who would pay for the work? The research identifies ARL broad agency announcements, SBIR and STTR awards and the transition programs that follow research. Ceilings range from $167,500 to $30 million, and the Small Business Administration guideline for a Phase II award is $2.24 million. The ceiling depends on the year and the solicitation.

Can a small business work with a historically Black college or university? Yes, through a vehicle that pays both. The HBCU and minority-serving institution programs that were retrieved fund institutions only. An STTR award requires the research institution to perform at least 30 percent of the work, and a subaward or a cooperative agreement are other routes. SOL's HVAC modeling work with Alabama A&M University students is an existing relationship of this kind.

Must university research be published? Basic research at budget category 6.1 is fundamental research at any performer, and applied research at 6.2 is fundamental on a university campus. Advanced development at 6.3 and physical items such as a printed cold plate fall outside these rules. A university that accepts a sponsor's publication restriction loses fundamental research status for that work.

What are CUI and CMMC? Controlled unclassified information is sensitive Government information that is not classified. DFARS 252.204-7012 requires contractors that hold it to meet the 110 security requirements of NIST SP 800-171 and to report cyber incidents within 72 hours. The Cybersecurity Maturity Model Certification program added assessments of that compliance, and on 13 July 2026 the Government suspended its third-party assessment phase while leaving the self-assessment phase in place.

Who owns the software and models? The contract records ownership through assertions. An offeror lists the data and software it delivers with restricted Government rights and gives a basis for each, and the Government assumes no restriction on anything left off the list. Under an SBIR or STTR award, data stay protected for 20 years from award and then pass to Government purpose rights. The research found no

solicitation that requires documentary proof of ownership of an algorithm, and it found no settled rule for AI models trained on Government-funded data.

6.6 Questions about the market

Is there a commercial market for this cooler? The research documents a large and fast-moving market for liquid cooling of data center accelerators, with estimates from $0.87 billion to $4.97 billion in 2024 depending on scope. It also documents vendor claims of higher power and flux than the CF-JET design point. These claims lack independent test evidence, and the research could not verify them.

Why are large companies buying cooling suppliers? The sources report that power and fluids groups are acquiring thermal suppliers as accelerator power rises. Eaton agreed to buy Boyd Thermal, Ecolab to buy CoolIT Systems and Schneider Electric bought a 75 percent interest in Motivair, at stated prices that total $15.1 billion (computed).

7. Where to Find More Detail and What the Terms Mean

7.1 Each topic has a full research report with sources

Table 4 lists the topics and the titles of the research reports that treat them in depth. Each report states its sources, its derived values and its evidence gaps.

Table 4. Topics and the research reports that cover them

Topic	Research report title
Army electronics heat and coolant limits	Thermal Management Demands of Silicon Carbide Power Electronics in Army Ground Vehicles
Expansion mismatch and wear-out	Thermal Expansion Mismatch and Power Module Reliability
Limits of conventional coolers and jets	Limits of Conventional Liquid Cold Plates and Jet-Impingement Crossflow
Designing with few tests	Engineering Design Under Sparse Data
3D printed jet plates	Additively Manufactured Jet-Impingement Cold Plates with Interleaved Returns
Carbon fiber target plates	Carbon Fiber Composite Target Plates for High-Heat-Flux Cooling
Fast stand-in models and optimization	Surrogate Modeling, Multi-Fidelity Learning, and Optimization for Thermal-Fluid Design
Physics rules plus learning	Uncertainty-Aware Neuro-Symbolic AI for Engineering Decisions
Testing and quality checks	Test Infrastructure, Validation, and Quality Assurance for AM Thermal Hardware
Markets and funding paths	Markets, Transition Pathways, and Commercialization for Advanced Cold Plates
Partners, data rights and compliance	Program, Teaming, Data Rights, and Compliance Considerations

Titles as issued. Interpretation: the grouping of topics is a reading aid.

7.2 A short glossary explains the technical terms

Table 5. Glossary of terms used in this summary

Term	Plain meaning
Additive manufacturing	Building a part layer by layer from metal powder or plastic, commonly called 3D printing
Cold plate	A metal or composite plate with liquid passages that removes heat from electronics
CFD	Computational fluid dynamics, a computer simulation of fluid flow and heat transfer

Term	Plain meaning
CTE	Coefficient of thermal expansion, how much a material grows when heated
CUI	Controlled unclassified information, sensitive Government information that is not classified
Heat flux	Heat power passing through each unit of area, in W/cm²
Jet impingement	Cooling by aiming liquid jets straight at a hot surface
Pressure drop	The loss of liquid pressure as it passes through the cooler
Silicon carbide (SiC)	A semiconductor that handles high power and high temperature in vehicle inverters
Surrogate model	A fast mathematical stand-in for a slow simulation, trained on simulation results
Thermal resistance	Temperature rise per watt of heat, in K/W
Uncertainty interval	A range around a prediction that states how far off the prediction might be

Definitions are plain-language readings of the terms as the research uses them.