

POWER PRIMER

Behind the Meter Power for Semiconductor Fabs

How chips are made, why a fab draws the power of a small city, and what that means for water, reliability, and cost. An educational overview from Atlas Energy Solutions.



Every semiconductor — the processor in a phone, the memory in a data center, the controller in a car — is made in a fabrication plant, or "fab." Fabs are among the most complex facilities humans build, and among the most energy-hungry: a single leading-edge fab can draw as much continuous electricity as a small city, every hour of every day, for the twenty-plus-year life of the plant. As chip manufacturing expands in the United States, that power demand has stopped being a technical detail and started showing up in debates about industrial policy, grid planning, and community relations. This short primer explains the connection from the ground up: how a fab works, why it needs the power it does, and how the choice of power source shapes water use, reliability, and cost.

Written for policymakers, journalists, community stakeholders, and industry newcomers who want the real engineering picture rather than headline numbers.

01 / INSIDE THE FAB

How a semiconductor fab actually works

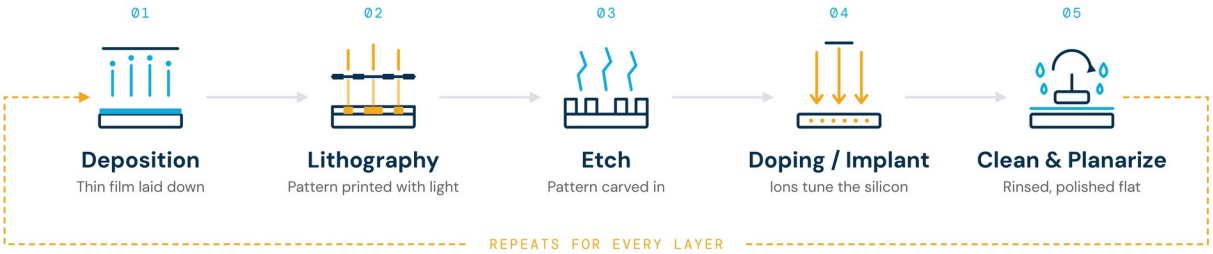
To understand fab power, you first have to understand what a fab does. Chipmaking isn't assembly. It's the repeated, atom-scale patterning of circuitry onto silicon wafers, inside an environment far cleaner and more tightly controlled than a hospital operating room.

THE WAFER JOURNEY

A fab takes a polished silicon disc (a wafer) and builds integrated circuits on its surface through hundreds of sequential steps. A leading-edge chip can require **more than 1,000 process steps** and take three to four months to complete. The core steps repeat in cycles:

- **Deposition.** Ultra-thin films of conducting or insulating material are laid down across the wafer.
- **Lithography.** Light is projected through a patterned mask to print circuit geometry onto a light-sensitive coating. Leading-edge fabs use extreme ultraviolet (EUV) machines — among the most complex tools ever built.
- **Etch.** The exposed pattern is carved into the film with plasma or chemicals.
- **Doping / implant.** Ions are driven into the silicon to tune its electrical behavior.
- **Clean and planarize.** The wafer is rinsed with ultrapure water and polished flat before the next layer begins.

These five stages repeat, layer upon layer, until the full three-dimensional circuit is complete. Every cycle demands precise temperature, ultra-clean air, ultrapure water, exotic gases, and — underpinning all of it — utterly stable electrical power.



The core wafer cycle. Each pass adds one layer; a leading-edge chip can require more than 1,000 process steps.

THE CLEANROOM IS THE FAB

The wafers travel through a cleanroom: a sealed space where airborne particle counts are held thousands of times lower than outdoor air. The most sensitive zones (EUV lithography) run at ISO Class 1–2, permitting only a handful of particles per cubic meter. Maintaining that environment requires enormous volumes of filtered, temperature- and humidity-controlled air, recirculated continuously. **The air-handling system that does this is the single largest consumer of electricity in the fab, ahead of any of the chipmaking tools.**

200–600	~30 MW	24/7/365
MW of continuous power draw for a leading-edge fab	Rule-of-thumb electric load per $1B of fab investment	Fab load runs flat, all year, for decades

WHAT SURPRISES MOST PEOPLE

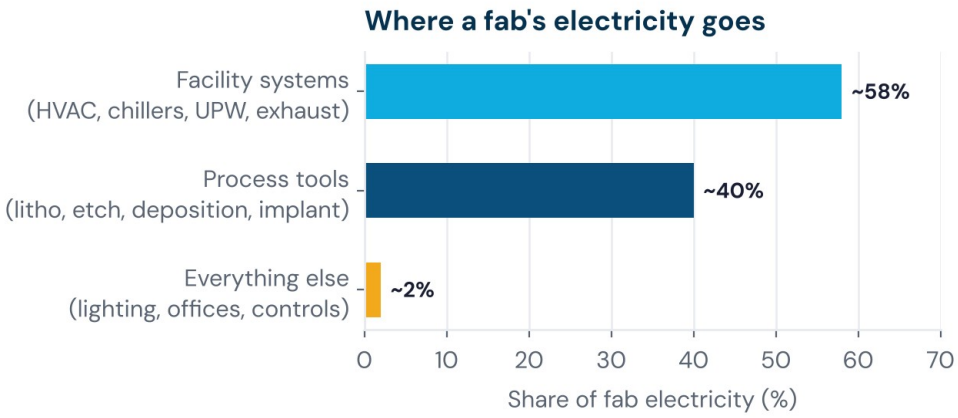
Most people assume the chipmaking machines use most of the power. They don't. The facility systems that maintain the environment — air handling, chillers, ultrapure water, gas and exhaust — consume more electricity than the process tools themselves. And most of that runs at full power whether or not a single wafer is being processed. A fab is, in energy terms, a giant environmental-control machine with chipmaking tools inside it.

02 / THE LOAD

Where the power goes

Fab electricity splits into two big blocks, and everything else barely registers. Measured studies of working fabs tend to land on similar proportions: facility systems take **roughly 55–60%**, process tools **about 40%**, and lighting, offices, and controls the last few percent.

LOAD BLOCK	SHARE OF TOTAL	WHAT'S IN IT
Facility systems	~55–60%	Cleanroom HVAC (the largest single item), chillers, ultrapure water, process cooling water, exhaust/abatement, compressed dry air, gases, vacuum
Process tools	~40%	Lithography (incl. EUV), etch, deposition, implant, metrology — the machines that pattern the wafer
Everything else	~1–3%	Lighting, offices, controls



Illustrative, based on published fab energy studies. Proportions vary by node, climate, and facility design.

Within the facility block, cleanroom HVAC dominates — and cleanroom air conditioning runs at roughly **ten times the energy density** of ordinary comfort air conditioning, because the air is filtered to extreme standards and fully exchanged many times per hour.

WHY THE LOAD NEVER RESTS

What makes fab power unusual is how constant it is. Three things explain that:

- **Support systems run continuously.** Vacuum pumps, purge-gas flows, water heaters, and abatement units typically run at full power even when the tool above them sits idle waiting for the next wafer lot. The environment must be held constant at all times.
- **The cleanroom cannot be turned down.** Letting particle counts, temperature, or humidity drift — even briefly — risks contaminating in-process wafers worth millions. HVAC runs flat, 24/7/365.
- **There is no seasonal relief.** Unlike an office or a factory, a fab has no nights, weekends, or off-season. Its load curve is nearly a flat line, which is unusual and important for how its power should be supplied.

THE LOAD-TO-INVESTMENT RULE OF THUMB

A new fab adds roughly 30 MW of electric load for every $1 billion of capital investment. A $20B leading-edge project therefore implies a continuous load on the order of several hundred megawatts. That's why power availability, more than incentives or workforce, is usually the first hard filter on where a fab can go.

POWER QUALITY, NOT JUST POWER QUANTITY

A fab needs clean power as much as it needs abundant power. A voltage sag lasting only 10–20 milliseconds — far too brief for a person to notice a light flicker — can trip process tools, scrap in-process wafers, and force lengthy tool recalibration. Industry estimates put the cost of a single serious power-quality event in the **hundreds of thousands to tens of millions of dollars**. The industry standard SEMI F47 defines how much voltage sag tools must ride through, precisely because grid disturbances are common and costly.

10–20 ms	$10M+	SEMI F47
A voltage sag long enough to trip tools and scrap wafers	Upper-range cost of a single serious power-quality event	The standard for how much sag a tool must ride through

03 / THE TRADE-OFF

Grid power vs. on-site power: water, reliability, and cost

A fab can draw its electricity from the utility grid, generate it on-site (behind the meter), or combine both. Neither choice is right in every case, but the trade-offs are concrete and measurable. **One clarification first:** grid and on-site aren't really opposites. Most fabs that generate on-site stay grid-connected anyway, for backup and flexibility. The question is usually which source is primary, not whether to pick one and abandon the other.

DIMENSION 1 — WATER

Water and power are tightly linked in a fab, and the link runs both ways. It helps to separate two very different water stories.

The fab's own water use. Water is the largest-volume material in chipmaking. An average fab can use up to about 10 million gallons of ultrapure water per day, and large multi-fab campuses can approach the daily water demand of a city of a million people. Producing 1,000 gallons of ultrapure water takes roughly 1,400–1,600 gallons of municipal water. This water use is inherent to making chips and is largely independent of how the fab is powered.

The water embedded in the power. This is the part the power source controls. Grid electricity is not water-neutral: conventional thermoelectric power plants (coal, nuclear, and many gas plants) withdraw and consume large volumes of water for cooling. So a fab drawing grid power is indirectly responsible for water consumed at distant power stations. The water intensity of that power depends entirely on the regional generation mix.

WATER FACTOR	GRID POWER	ON-SITE POWER
Fab process water	Same either way — inherent to chipmaking	Same either way — inherent to chipmaking
Water for generation	Embedded in grid mix; thermoelectric plants are water-intensive	Depends on technology; some on-site options use little or no cooling water
Local visibility	Hidden — consumed elsewhere, rarely counted against the fab	Local and accountable — easier to measure and manage on-site

So the power source doesn't change how much process water a fab needs. What it changes is the water footprint of the electricity, and whether that footprint sits on-site where it can be measured and managed or stays hidden somewhere in the regional grid.

DIMENSION 2 — RELIABILITY

Fabs are less forgiving about reliability than almost any other kind of facility. The question isn't whether the power is on. It's whether the power is continuous, clean, and stable down to the millisecond.

FACTOR	GRID POWER	ON-SITE POWER
Exposure to outages	Subject to weather, wildfires, grid faults, and regional capacity events	Insulated from most grid-side disruptions; islanded operation possible
Voltage-sag events	Grids see frequent sags from faults and load switching; each risks wafer scrap	Dedicated generation can be engineered for tighter power quality
Curtailment risk	May face demand-response or curtailment during grid stress	Self-supplied load is not subject to grid curtailment
Control of destiny	Reliability depends on the utility's system and neighbors	Reliability is controlled and maintained by the operator

Grid power isn't unreliable in absolute terms; modern grids are highly available. But a fab's tolerance for disturbance is so low, and the cost of each event so high, that even a grid with "good" reliability statistics can impose millions in annual losses through brief sags that most customers would never notice. **On-site generation doesn't eliminate risk; it trades the utility's maintenance and redundancy for the operator's own. What changes is who controls reliability.**

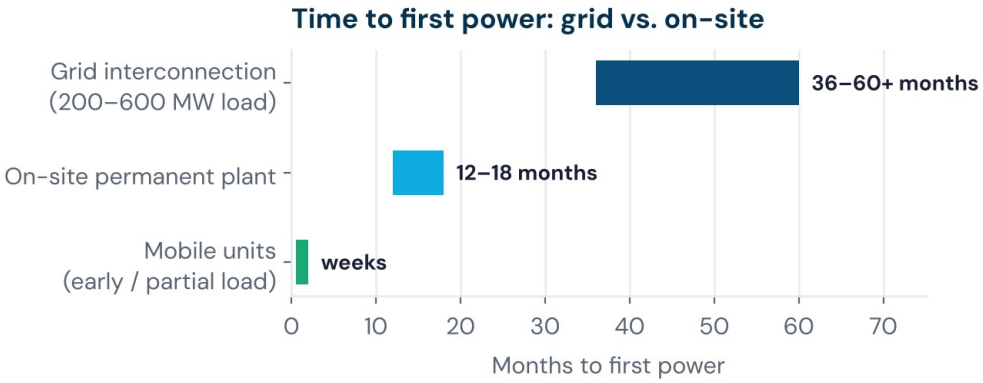
DIMENSION 3 — COST

Electricity is one of the largest ongoing operating costs of a fab, so how it is priced and how volatile it is feed directly into chip economics. Cost comparison has to weigh both the price of energy and the risk around that price.

COST FACTOR	GRID POWER	ON-SITE POWER
Up-front capital	Low — utility owns generation; fab pays for interconnection	Higher — generation must be built or contracted
Price stability	Exposed to wholesale market volatility and rate changes	Can be contracted to a known, stable long-term cost
Demand charges	Large demand charges on a flat, always-on load	Avoided or reduced for self-supplied load
Transmission cost	May fund grid upgrades; costs can flow to ratepayers	Reduces or defers transmission build-out
Time to power	Interconnection for large loads can take 3–5+ years	Can be deployed on the fab's construction schedule

The cost picture is mixed. Grid power usually wins on up-front capital; on-site power usually wins on price certainty, demand-charge avoidance, and speed to energization, which matters a great deal in the current U.S. buildout. Which matters more depends on the project's

timeline, the local grid's condition, and how much the developer values predictability over the plant's decades-long life.



Illustrative, based on public analyses of large-load interconnection timelines. Actual timelines vary by region and utility queue.

THE TIMING FACTOR THAT RESHAPES THE WHOLE COMPARISON

In today's U.S. environment, interconnecting a 200–600 MW load can take three to five years or more, because fabs, AI data centers, and other large loads are all competing for the same transmission capacity. When a fab building is finished before the grid can serve it, the cost comparison changes character entirely: the relevant question is no longer just price per megawatt-hour, but whether the multi-billion-dollar plant can produce anything at all on schedule. This is why on-site generation is increasingly discussed as a bridge to grid power, a long-term primary source, or both.

04 / THE ALTERNATIVE

Why an on-site, behind-the-meter approach fits a fab

"Behind-the-meter" (BTM) power means generating electricity at or next to the facility that uses it, rather than drawing it across the public grid. For a load that is huge, flat, and unforgiving, this approach can offer meaningful advantages:

- **Speed.** On-site generation can be built on the fab's own schedule, typically reaching first power in twelve to eighteen months, and mobile units can carry an early or partial load within weeks while the permanent plant is assembled.
- **Reliability and control.** An islanded plant has no utility to fall back on, so generation is built with N+1 or 2N spare capacity and paired with battery storage that rides through the millisecond disturbances that scrap in-process work.
- **Cost predictability.** Power shows up as a contracted operating cost per kilowatt-hour instead of a substation, an interconnection fee, and a depreciation schedule on the customer's balance sheet.
- **Scalability.** Modular generation grows with the load, so a second phase is a contract amendment rather than a new interconnection study.

- **Water.** Closed-loop, radiator-cooled generation consumes almost no cooling water, which keeps the water embedded in the electricity from competing with the ultrapure water the process itself requires.
- **Community benefit.** Because it doesn't draw on the local utility, on-site power is not competing with homes and small businesses for capacity or turning up in residential rates, and it does not require new long-distance transmission corridors across other people's property.

On-site generation isn't automatically the right answer everywhere. **Technology follows the site, not the other way around.** Fuel availability, emissions limits, whether there is process heat worth capturing, the shape of the load, and how much land is available all pull the configuration in different directions. High-efficiency natural gas reciprocating engines are one well-proven option: a Caterpillar CG260-16, to use a familiar example, produces roughly 4 MW per unit at about 44% electrical efficiency, holds that efficiency at part load because idle units simply switch off, reaches around 80% total efficiency where waste heat can serve a process load, and runs about 80,000 hours before its first overhaul. Turbines, storage, renewables, grid supply, and hybrids of all of them have their place in the mix. The engineering decides.

3–5+ yrs	12–18 mo	N+1 / 2N
Typical time to achieve grid interconnection for a 200–600 MW load	Typical time to first power for a permanent on-site plant	Spare generation designed in for islanded operation

05 / THE ATLAS APPROACH

A faster path to power — and a partner who keeps it on

Atlas Energy Solutions develops behind-the-meter power as an integrated service, working with customers across five stages — **Scope, Design, Contract, Build, and Operate**. The stage that matters most once a plant is running is the last one: Atlas is a long-term **operator**, and treats reliable, dedicated power as the product it delivers. Atlas does not sell generators. We engineer, finance, build, own, and operate the power plant, and the fab buys the electricity it produces at a contracted rate. The customer runs its plant; we run the power.

For a fab, power isn't a convenience — it's the environmental-control system, the cleanroom, and every wafer in process. So the operating model is built around keeping it on:

- **Dedicated, on-site generation.** Power is generated for the facility and sized to its load, so the fab isn't competing with other users or exposed to events elsewhere on the grid.

- **Redundancy designed in, not borrowed from the grid.** N+1 or 2N spare capacity, paired with battery storage that carries the plant through the millisecond sags that would otherwise scrap in-process work.
- **Proactive, condition-based maintenance.** Equipment is serviced on a planned schedule rather than run to failure — the approach designed to keep availability high over the decades-long life of a fab.
- **Dedicated staffing and 24/7 remote monitoring.** Around-the-clock monitoring, fuel supply and scheduling, and on-site response are part of the service rather than a staffing problem handed to the customer. Someone is watching it at three in the morning.

The commercial side is designed to match. Through an **Energy Supply Agreement (ESA)**, a customer can access power through a service-based commercial structure rather than owning and operating the generation assets itself. Power becomes a predictable operating expense rather than a capital project — a structure designed to **take the capital-expense burden of building a power plant off the customer's balance sheet**, while the responsibility for uptime stays with Atlas. The customer runs its plant; Atlas runs the power.

PUTTING IT TOGETHER

A few things about semiconductor manufacturing and power hold up regardless of technology or market cycle:

- **Fabs are environmental-control machines.** Most of their power maintains the cleanroom, not the chipmaking tools — and most of it runs whether or not wafers are moving.
- **The load is huge, flat, and unforgiving.** Hundreds of megawatts, 24/7/365, with near-zero tolerance for interruption or even millisecond disturbances.
- **Power is the first siting constraint.** A site without a credible path to reliable power on the construction timeline cannot host a fab, regardless of other advantages.
- **Grid vs. on-site is a real trade-off.** Grid tends to win on up-front cost; on-site tends to win on water visibility, reliability control, price certainty, and speed. Most fabs will use a blend.

> **THE IDEA IN ONE LINE**
>
> You can't separate the semiconductor from its power. Behind every chip is a small city's worth of continuous, ultra-clean electricity — and on-site, behind-the-meter generation is one of the keys to delivering it on the fab's schedule. Atlas builds, owns, and operates that power — and stands behind keeping it on — so its customers don't have to.

Your fab has a timeline. Don't let power risk it.

WWW.ATLAS.ENERGY · NYSE: AESI

About this document. *This is an educational overview intended for general information. Figures are industry ranges compiled from publicly available sources, including peer-reviewed fab energy studies, SEMI (including the SEMI F47 voltage-sag standard), the International Roadmap for Devices and Systems (IRDS), the World Economic Forum, the U.S. Department of Commerce, university engineering research, and public environmental assessments (e.g., Intel Ocotillo), and may change over time. Fab power and water figures vary by node, size, climate, and design; ranges are representative, not exact for any single facility. Charts are illustrative. Nothing here is a commitment, offer, guarantee of performance, or professional advice. Any specific project would be evaluated on its own facts.*