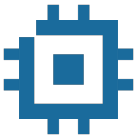

# REEL LCI Database

## Methodology Report

**Version 1.0**

21 August 2026

reellci.com

**Abstract**

This report documents the methodology for a Life Cycle Inventory (LCI) database covering a comprehensive range of electronics components. Coverage includes semiconductor integrated circuits across technology nodes from 180nm to 3nm, memory devices (DRAM, NAND flash, HBM), printed circuit boards, passive components (capacitors, resistors, inductors), IC packaging, cables and interconnects, and electromechanical components. The database provides cradle-to-gate inventory data derived exclusively from publicly available sources using a bottom-up "virtual factory" modeling approach. This methodology report follows ISO 14040/14044 guidance and aims to provide LCA practitioners with sufficient information to evaluate data quality, understand allocation procedures, and appropriately apply the inventory data in their assessments.

# Version History

This section documents major updates to the methodology report and underlying database. Interim development milestones between public versions are consolidated into each version's entry; the report and database are revised continuously between public versions, and the date on the title page is the reference date of the active version as recorded in the table below.

**Table 0.1:** Document and database version history

| Version | Date | Changes |
|---|---|---|
| 0.1 | 2026-03-23 | Initial public release. 363 unit process datasets across 18 component categories covering semiconductor wafers (3nm–180nm), complete ICs, memory (DRAM, NAND, HBM), advanced packaging (CoWoS, FOWLP, FC-BGA), PCBs, passive components, connectors, cables, and supporting infrastructure. |
| 1.0 | 2026-08-21 | Major expansion and consolidation of the database ahead of external critical review, growing coverage to 520 unit process datasets.<br>*Coverage:* current-generation memory added (HBM4 stacks, GDDR7, 1$\gamma$ DRAM, 321-layer 3D NAND, DDR5 module families), silicon photonics and optical transceivers with explicit III–V optical sub-assemblies, single-mode optical fiber and fiber connectors, tantalum capacitor vendor variants, and an expanded set of upstream specialty-chemical datasets.<br>*Variants:* datasets now ship with declared variant axes — operational scenarios for wafers, geography options for leading-edge logic and specialty wafers (with per-technology representative manufacturing regions), stack heights and production-maturity yield scenarios for high-bandwidth memory, and surface finishes for PCBs. |

*continued on next page*

Table 0.1 – continued from previous page

| Version | Date | Changes |
|---|---|---|
| | | *Methodology:* emissions-to-water extraction with chemistry-based transformation rules and abatement destruction efficiencies; a reworked facility-support model anchored to publicly disclosed reference fabs, with recipe-driven cleanroom HVAC and per-region climate adjustment; a facility-overhead layer for back-end assembly-and-test sites; a uniform good-unit functional-unit convention with explicit yield ownership between wafer, packaging, and component datasets; die-count allocation based on rectangular die packing; end-to-end uncertainty propagation with declared band provenance; per-flow source citations resolved from curated source registries; and dataset-level data-quality scoring in which authored pedigree labels are disciplined against the underlying source evidence.<br>*Corrections:* successive internal review rounds re-anchored packaging process energies, etch-gas quantities and fate, and upstream feed stoichiometry; every revision was verified dataset-by-dataset against the preceding release state.<br>*Provenance and review:* row-level citation and uncertainty-band coverage are accounted per inventory row with recorded exemptions (Chapter 3); data-quality composites use configuration-weighted scoring with one-decimal display; the intra-chain freight layer uses structured legs with a chargeable-weight air basis and documented route mixes; and an independent critical review was performed (signed statement reproduced in Appendix D). |

# Contents

# List of Tables

# List of Figures

# Chapter 1

# Introduction

## 1.1 Background and Context

Electronics manufacturing represents one of the most complex and environmentally intensive industrial sectors (Boyd et al. 2009; Williams et al. 2002). As semiconductor process complexity increases with each technology node, manufacturing energy intensity continues to grow. Meanwhile, data centers consumed approximately 1.5% of global electricity in 2024, a figure projected to double by 2030 driven largely by AI workloads (International Energy Agency 2025). This growth underscores the importance of understanding embodied carbon in electronics, as hardware manufacturing impacts may rival or exceed operational energy use for certain IT equipment. Despite this significance, the sector suffers from a lack of openly available life cycle inventory (LCI) data. This data gap presents substantial challenges for:

- **Corporate sustainability teams** attempting Scope 3 emissions accounting for purchased electronics
- **Product designers** seeking to make informed eco-design decisions
- **LCA practitioners** conducting assessments of electronics-containing products
- **Policy makers** developing environmental regulations for the electronics sector

### 1.1.1 The Electronics Data Challenge

Existing LCI data for electronics is limited by several factors:

**Proprietary barriers** Semiconductor manufacturers treat process data as highly confidential competitive intelligence. Direct measurement data from operating fabrication facilities is rarely shared publicly.

**Rapid technology evolution** Moore's Law scaling means that manufacturing processes evolve on 2–3 year cycles, making older data quickly obsolete for advanced technology nodes (IEEE International Roadmap for Devices and Systems 2024).

**Supply chain complexity** A single semiconductor product may involve 10+ distinct manufacturing facilities across multiple countries, making data collection extremely challenging.

**Aggregation obscures detail** When companies do report environmental data, it is typically aggregated at the facility or company level, obscuring the process-level detail needed for product LCA.

### 1.1.2 Project Mission

The REEL LCI Database addresses these challenges through a bottom-up "virtual factory" modeling approach that constructs detailed, auditable inventory data using only publicly avail-

able information. Rather than relying on proprietary facility measurements, this approach synthesizes data from:

- Equipment vendor technical specifications
- Peer-reviewed academic literature
- Corporate sustainability reports and disclosures
- Industry standards and technology roadmaps
- Patent filings with process parameters
- Engineering calculations based on physics and chemistry

This public data policy ensures that all values in the database are fully traceable to documented sources, enabling independent verification and continuous improvement by the user community.

## 1.2 Goal and Scope

### 1.2.1 Database Goal

The goal of this database is to provide openly documented, estimated life cycle inventory data for electronics components, enabling practitioners to conduct cradle-to-gate environmental assessments without requiring access to proprietary manufacturing data.

The database produces **inventory data only**: the quantified inputs and outputs associated with manufacturing processes. This includes energy consumption (electricity, natural gas, steam), water use (ultrapure water, process water, cooling water), material inputs (chemicals, gases, metals, substrates), wastewater generation, and air emissions (including greenhouse gases and volatile organic compounds). The database does not provide characterized impact results (such as kg $CO_2$e or kg $SO_2$e), as these require user selection of appropriate impact assessment methods.

### 1.2.2 Temporal and Geographical Scope

#### 1.2.2.1 Temporal Scope

The database represents the manufacturing landscape as of **2023–2025**, with the underlying public sources sampled between October 2025 and May 2026. Process parameters, equipment specifications, and facility characteristics reflect current production practices for each technology node. Given the rapid evolution of semiconductor manufacturing, users should consider temporal relevance when applying this data:

- **Advanced nodes** (3nm–7nm): Data reflects 2023–2025 high-volume production; assumed die yields step down toward the newest nodes (7nm highest, 3nm lowest) to reflect each node's position on the yield-maturity curve (Table 4.4)
- **Mature nodes** (14nm–65nm): Data reflects established high-volume manufacturing
- **Legacy nodes** (90nm+): Data reflects mature, stable processes

#### 1.2.2.2 Geographical Scope

The virtual factory models represent **global average** manufacturing practices rather than measurements of specific facilities. Each dataset is then placed at a **baseline representative manufacturing region**, and for the categories where regional variation matters most, **regional variant datasets** are published alongside the baseline:

**Baseline regions** Each category is assigned the region where the modeled production predominantly takes place: advanced logic and silicon photonics (Taiwan), memory (South Korea/Taiwan), packaging and assembly (Taiwan), PCBs (China). Specialty wafers carry per-technology baseline regions — for example Germany for SiC/GaN power fabrication, the United States for III–V photonics, Japan for image sensors, and China for LEDs and mature discretes.

**Regional variants** Where supply chains meaningfully diverge, the same model is published as additional geography-specific datasets: the leading-edge logic nodes ship both Taiwan and Arizona (US) datasets, and six specialty-wafer technologies ship secondary US or Japan datasets. Variant datasets differ in their regional electricity market reference and in climate-driven facility support energy; the underlying process flow is identical (see Chapter 5).

**Electricity grid mix** Electricity consumption is provided as a separate flow, so users can link it to the grid mix of their actual supply-chain geography. Grid mix is typically the single most sensitive parameter for climate impact calculations; where a geography-matched variant dataset exists, selecting it is preferable to re-gridding the baseline.

**Geography labels vs. data representativeness.** Each exported dataset carries a `geography` code (TW, KR, CN, GLO, ...) that records where the modeled production is assumed to take place: it sets the dataset's electricity-grid reference and its regional wastewater and effluent context. It does *not* mean the underlying engineering data were measured in that country — the process inventories themselves are assembled from public sources worldwide. A dataset placed in China, for example, applies Chinese grid electricity and Chinese effluent context to a globally-sourced process inventory; its documentation therefore describes the engineering data as global-average even though the dataset carries a specific country code.

### 1.2.3 Reference Flows

Following ISO 14044 terminology, the database uses **reference flows** rather than functional units for these intermediate products. A capacitor or semiconductor die does not have an intrinsic function until integrated into a final product; however, it has a defined reference flow (e.g., one piece, one wafer) that serves as the basis for inventory calculations.

### 1.2.4 Intended Applications

The database is designed to support the following use cases:

**Scope 3 emissions accounting** Companies can use component-level inventory data to esti-

**Table 1.1:** Reference flows by component category

| Category | Reference Flow | Notes |
|---|---|---|
| Semiconductor wafers | 1 wafer | 300mm diameter; 200mm for legacy nodes |
| Complete ICs | 1 IC | Combines wafer fabrication + packaging |
| Memory components | 1 component | HBM stacks, DRAM ICs |
| PCBs | 1 $m^2$ | Normalized to board area |
| Passive components | 1 piece | MLCCs, resistors, inductors |
| Packaging | 1 package | Various package types |
| Cables & connectors | 1 cable / 1 piece | Finished cable assemblies and board connectors; bulk single-mode optical fiber is published separately per km |
| Other | 1 piece / 1 component | Thermal, electromechanical, infrastructure |

mate Category 1 (purchased goods) emissions for electronics procurement.

**Product LCA** LCA practitioners can integrate electronics component data into broader product assessments, linking to background databases for upstream material production.

**Eco-design decisions** Product designers can compare environmental profiles of alternative components (e.g., different technology nodes, package types, or passive component materials).

**Hotspot identification** Organizations can identify the most environmentally intensive components in their bill of materials for prioritized improvement efforts.

**Scenario analysis** Users can model different manufacturing scenarios (e.g., regional grid mixes, abatement technologies) to understand sensitivity to key parameters.

### 1.2.5 Target Audience

This database and methodology report are intended for:

- LCA practitioners with general knowledge of life cycle assessment but not necessarily semiconductor manufacturing expertise
- Corporate sustainability professionals conducting supply chain assessments
- Product designers evaluating environmental implications of component selection
- Researchers studying electronics manufacturing environmental impacts

## 1.3 Alignment with ISO Standards

This methodology report is developed following the principles and framework of ISO 14040:2006 (International Organization for Standardization 2006a) and ISO 14044:2006 (International Organization for Standardization 2006b).

**Important clarification**: This statement describes alignment with ISO principles and structure, not a formal compliance declaration. Full ISO compliance requires external critical review and specific procedural requirements that depend on the intended application. Per ISO 14044,

external critical review is required for any LCA results intended for public disclosure. Users should arrange appropriate critical review when publicly communicating results derived from this database.

**Independent critical review**: The v1.0 database and this methodology report, dated 21 August 2026, have undergone an independent critical review, completed on 25 August 2026. The signed review statement is reproduced verbatim in Appendix D; Section 6.7.2 summarizes its scope and validity terms.

The methodology addresses the key elements specified by ISO 14044 for LCI studies:

- **Goal and scope definition** (this chapter)
- **System boundary specification** (Chapter 2)
- **Allocation procedures** (Chapter 2)
- **Data quality requirements** (Chapter 3)
- **Calculation methods** (Chapters 2 and 4)
- **Limitations documentation** (Chapter 6)

## 1.4 Database Coverage

The database provides **520 unit process datasets** organized across 18 component categories. Table 1.2 summarizes coverage.

Many datasets are published as **variant families** of a shared inventory definition: each of the 19 PCB designs ships with 6 surface finish options (HASL, ENIG, ENEPIG, ImAg, ImSn, OSP), producing 114 datasets; wafer and memory-wafer datasets include operational scenario variants (baseline, moderate-recycling, best-practice) plus regional variants where relevant (Arizona/Taiwan for the 5 nm and 3 nm nodes) and cryogenic-etch variants for high-layer-count NAND; the HBM3/HBM4 TSV-stacking packages ship in 8/12/16-high stack variants; the HBM4 stacks additionally ship ramp- and mature-yield scenario datasets alongside the baseline yield; the tantalum capacitor family spans 24 vendor-catalog SKUs across cathode chemistry ($MnO_2$, conductive polymer), CV/g grade, and EIA case size; specialty wafers ship under per-technology representative manufacturing regions, with secondary regional options published where their inventories differ from the reference region. Each variant family's reference cell is the base dataset itself, which states the configuration it represents; only variant cells whose inventories differ from that reference are published as separate datasets. Upstream material datasets provide gate-to-gate inventories for specialty chemicals not available in standard LCI databases (see Chapter 2).

**Table 1.2:** Database coverage by component category (v1.0)

| Category | Datasets | Examples |
|---|---|---|
| Semiconductor wafers | 48 | 3nm–180nm logic, BCD 180 nm; scenario + region variants |
| Memory wafers | 39 | DRAM $1\alpha/1\beta/1\gamma$, NAND 128/176/232/288/321L; scenario + cryo variants |
| Specialty wafers | 20 | SiC, GaN, MEMS, CIS, InP, GaAs, SiPh; per-technology regions, secondary region options where distinct |
| Logic ICs | 42 | MCUs, PMICs, PHYs, FPGAs, DDR5 companion ASICs, UFS controller |
| Memory components | 31 | HBM3/3e/4 (HBM4 in 3 yield scenarios), GDDR7, DRAM ICs (incl. $1\gamma$), LPDDR5X stacks, HBF |
| Packaging | 46 | FC-BGA, QFN, WLCSP, CoWoS, FOWLP, HBM TSV stacking |
| PCBs | 114 | Server, HDI, flex, backplane; 6 surface finishes each |
| Substrates | 2 | ABF, BT |
| Passive components | 57 | MLCCs, resistors, inductors, transformers; tantalum capacitors in 24 vendor SKU variants |
| Specialty semiconductors | 16 | SiC, GaN, MEMS, discrete power, transistors |
| Optoelectronics | 12 | Optical transceivers (100–400G), LEDs, fiber |
| Cables | 10 | DAC, fiber LC/MPO, USB4, SATA, power |
| Connectors | 10 | RJ45 MagJack, fiber LC/MPO, PCIe, USB4, ATX |
| Electromechanical | 4 | Axial fans, blower fans, BLDC motors, vibration |
| Thermal | 5 | Heat pipes, vapor chambers, TIM, heatsinks |
| Components (system-level) | 30 | PSUs, DIMMs (DDR5 CUDIMM/MRDIMM/CAMM2, SOCAMM), SSDs, HDDs, UFS |
| Infrastructure | 14 | MGX rack, busbars, PDUs, manifolds, fan walls |
| Upstream materials | 20 | $NF_3$, $C_4F_8$, CMP slurries, photoresist, TMAH developer solution, etc. |
| **Total** | **520** | |

## 1.5 Report Structure

This methodology report is organized as follows:

**Chapter 2: General Methodology** describes the core methodological principles including system boundaries, the virtual factory concept, allocation procedures, and data sourcing hierarchy.

**Chapter 3: Data Quality and Uncertainty** presents the data quality assessment framework, uncertainty characterization methods, and validation approach.

**Chapter 4: Sector-Specific Modeling** details the modeling approaches for each component category, including technology-specific parameters and key assumptions.

**Chapter 5: User Guidance** provides practical instructions for integrating the database with background databases and impact assessment methods.

**Chapter 6: Limitations and Future Work** documents known data gaps, limitations, and planned improvements.

**Appendix A: Glossary** defines technical terms used throughout the report.

**Appendix B: Process Inventory** lists the representative process families modeled in the database; the exported datasets are the exhaustive record.

**Appendix C: Data Sources** describes the source categories and principal anchors; the complete bibliography is given in the reference list and the per-dataset citations.

# Chapter 2

# General Methodology

This chapter describes the core methodological principles underlying the REEL LCI Database, including system boundaries, the virtual factory modeling approach, allocation procedures, and data sourcing hierarchy.

## 2.1 System Boundaries

### 2.1.1 Cradle-to-Gate Scope

The database provides cradle-to-gate life cycle inventory data, encompassing:

- **Upstream processes**: Raw material extraction and refinement (via background database linkage)
- **Manufacturing processes**: Component fabrication, assembly, and testing
- **Intra-chain freight**: Movement of the part-finished product between REEL-modelled sites — principally wafer fabrication to packaging (see below)
- **Gate boundary**: Finished component ready for integration into higher-level assemblies

Excluded from scope:

- Use phase (product operation)
- End-of-life treatment (recycling, disposal)
- Distribution and retail, including freight to the end customer
- Inbound transportation of raw materials to the manufacturing facility (included in upstream "market for" datasets when users link to background databases)
- Photomask fabrication (the mask set's embodied burden is amortized across tens of thousands of wafers at high production volumes, contributing less than 0.1 % per wafer; users assessing low-volume production should add mask fabrication separately)
- Employee transportation, facility administration, and R&D/pilot-production overhead

Silicon ingot growth and wafer slicing are *inside* the cradle-to-gate scope but are treated as upstream material inputs linked to background databases, rather than modeled as REEL processes. These exclusions and treatments are the normative boundary rules for the database; sector chapters note where they bite (e.g. Section 4.1.3), and their implications for use are summarized in Chapter 5.

Figure 2.1 illustrates the system boundary for semiconductor wafer fabrication, showing the relationship between utility inputs, support processes, process equipment, and outputs.

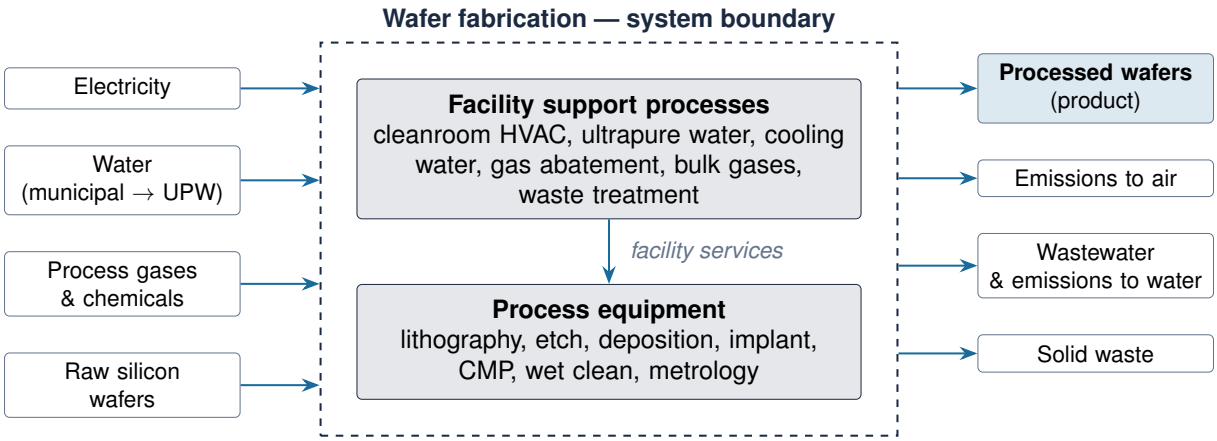


**Figure 2.1:** System boundary for semiconductor wafer fabrication. Utility inputs (electricity, water, gases, chemicals) flow through support processes and process equipment to produce processed wafers along with air emissions, wastewater, and solid waste. Support processes provide the facility infrastructure required by process equipment.

### 2.1.2 Intra-chain freight

Purchased-material freight is already carried by the upstream "market for" datasets, and freight to the end customer is out of scope. Between those two sit movements that neither covers: the part-finished product travelling between REEL-modelled manufacturing sites, and — for a small set of heavy finished assemblies such as racks, power supplies, and drives — the inbound movement of the finished good to the system-integration site. An integrated-circuit dataset is physically a wafer fraction plus a package, and where fabrication and packaging are at different sites that wafer was flown or trucked between them. Those legs are included: the published corpus carries 173 structured transport legs (161 fabrication-to-packaging and 12 inbound), of which 106 are road, 66 air, and one sea.

Each leg is authored on the consuming dataset as a tonne-kilometre input to an ecoinvent freight activity. Transported mass starts from the wafer fraction times the wafer's mass, the latter taken from the wafer dataset's own declared diameter (0.125 kg at 300 mm; 0.053 kg at 200 mm), grossed up for shipping packaging: a front-opening shipping box (SEMI M31) weighs roughly one and a half times the silicon it carries, and ships inside a secondary carton. The packaging's *production* burden is excluded because these containers are returnable, but its *mass* is carried, because freight scales with what the vehicle moves. Two mode-specific refinements apply. For **air legs**, the central mass basis is the *chargeable weight* on which low-density air cargo is actually billed — wafer shipments are volumetric-limited, so billed weight exceeds physical weight, and using physical mass would understate the freight effort the movement really commands. For **returnable-carrier logistics**, the empty return journey of the reusable containers is carried as an upper-endpoint term on the leg's uncertainty interval rather than as an uplift to the central value, since return-flow consolidation practices are not publicly evidenced per route.

Routing is assigned per dataset from the wafer's declared origin geography to the packaging site's geography, so the route follows the wafer's journey rather than the consuming product's nominal location. Advanced-node wafers ship on evidenced named routes. For the mature-

node commodity lanes out of Taiwan (28–180 nm class), no single destination is defensible, so the leg is authored as a documented *route mix*: a 55% offshore share to a composite offshore destination (45% China at 900 km, 55% Southeast Asia at approximately 2,330 km, giving a 1,700 km composite) with the remaining 45% on a 100 km domestic road lane. The mix's uncertainty interval spans a 50 km all-road lower bound to a 3,230 km all-air upper bound, the latter carrying the chargeable-weight and carrier-return factors. Where the evidence shows a vendor performs packaging on the fabrication campus — an on-campus memory back end, or a foundry advanced-packaging site inside the same science park — the band's lower bound is zero distance. An authored zero of that kind is a modelling statement about co-location, not a missing value.

Air freight is classified by ecoinvent into distance bands, and the legs modelled here fall in the short-haul band rather than long haul; the dedicated-freighter activity is used in preference to a market that blends dedicated freight with belly-hold cargo on passenger aircraft, because inter-site semiconductor freight is full-cargo. Road freight carries no distance bands in the background database and is selected by vehicle size, which is not established for these movements by any public record, so the fleet-average lorry activity is used in preference to asserting a vehicle class.

These flows are small. Across the datasets that carry them the freight legs contribute well under one percent of the dataset's screening-level greenhouse-gas total — the largest share anywhere in the corpus is just under one percent, on a complete rack assembly — and for logic and memory components it is a small fraction of one percent. They are carried because the movement is real and no other dataset accounts for it; freight between an assembly plant and a separate final-test site, where such a split exists, is not separately modelled. A larger transport contribution can lie beyond the system boundary, in outbound distribution: finished data-center equipment is commonly air-freighted from integration sites in Asia to consuming markets such as North America, and that outbound delivery leg can exceed the inter-site movements modelled here. Outbound distribution to the customer is excluded from the cradle-to-gate boundary (Section 2.1.1); users assessing delivered hardware should add that leg for their own supply routes.

### 2.1.3 Treatment of Capital Goods

Capital goods (including manufacturing equipment, cleanroom construction, and facility infrastructure) are currently **excluded** from the system boundary. This exclusion is based on:

- Lack of publicly available data on equipment embodied energy
- Significant uncertainty in equipment lifetime and allocation
- Common practice in electronics LCA to exclude capital goods
- Focus on operational inventory which dominates life cycle impacts

Future versions may include capital goods allocation when sufficient public data becomes available.

## 2.2 Modeling Framework

### 2.2.1 Attributional Approach

The database follows an attributional modeling approach, quantifying the physical flows directly associated with producing a defined functional unit. This approach:

- Describes the environmentally relevant inputs and outputs of a product system
- Uses average data for processes with multiple suppliers
- Does not model market-mediated effects or consequential changes

### 2.2.2 Virtual Factory Concept

A **virtual factory** is a modeled representation of manufacturing processes constructed using a bottom-up approach from fundamental process parameters, rather than top-down allocation from facility-level measurements (Krishnan et al. 2008; Murphy et al. 2003). Unlike traditional LCI approaches that rely on primary data collected from specific production sites, the virtual factory builds process inventories from first principles.

This bottom-up methodology is distinguished by:

**Process-level granularity** Each manufacturing step is modeled individually based on the physics and chemistry of the operation, enabling detailed understanding of inventory drivers.

**Equipment-based calculations** Energy and material consumption are derived from equipment specifications (power ratings, throughput, gas flows) rather than aggregated facility measurements.

**Stoichiometric foundations** Chemical consumption is calculated from reaction stoichiometry, film thickness targets, and utilization efficiencies rather than empirical facility averages.

**Explicit assumptions** All modeling assumptions are documented, enabling users to understand and, where appropriate, modify parameters for their specific context.

The virtual factory approach constructs process inventories from:

1. **Equipment specifications**: Power ratings, throughput, process parameters from vendor documentation
2. **Reaction stoichiometry**: Chemical consumption calculated from balanced equations and target film compositions
3. **Geometric parameters**: Material requirements derived from film thickness, wafer area, and feature dimensions
4. **Utilization factors**: Chamber efficiency, precursor incorporation rates, and waste generation based on published literature

This approach produces models that are independent of any specific facility, representing a generic manufacturing process that can be adapted through parameter adjustments for different scenarios.

### 2.2.3 Three-Tier Architecture

The calculation system uses a modular three-tier architecture that separates concerns and enables independent updates:

**Model files** Define *what* to calculate: process flow sequences, pass counts, yield parameters

**Configuration files** Define *parameters*: facility assumptions, abatement scenarios, support rates

**Process files** Define *how much*: per-operation energy, materials, emissions

## 2.3 Cut-off Criteria

### 2.3.1 Mass and Energy Thresholds

Flows are excluded from process inventories when they contribute less than 1% of the **total mass of inputs to the unit process** or less than 1% of the **total energy input to the unit process**, unless they:

- Involve substances of high environmental concern
- Are specifically regulated (e.g., PFCs, heavy metals)
- Have disproportionate impact potential (e.g., high GWP gases)

Note that the denominator is total process inputs, not product mass. In semiconductor manufacturing, input mass (water, chemicals, gases) is orders of magnitude higher than product mass, so this threshold excludes only truly minor flows.

### 2.3.2 Toxic Substance Exceptions

The following substance categories are always included regardless of mass contribution:

**Table 2.1:** Substances included regardless of cut-off criteria

| Category | Examples | Rationale |
|---|---|---|
| Perfluorocarbons | $CF_4$, $C_2F_6$, $SF_6$, $NF_3$ | $GWP_{100}$ ~7,400–24,300 (Intergovernmental Panel on Climate Change 2021); EPA regulated |
| Heavy metals | Pb, Cd, Hg, Cr(VI) | RoHS regulated; high toxicity |
| Volatile organics | Photoresist solvents, PGMEA | Air quality concerns |
| Precious metals | Au, Ag, Pd, Pt | High embodied impacts |
| Ozone depleting | CFCs, HCFCs (legacy) | Montreal Protocol |

## 2.4 Allocation Rules

### 2.4.1 Process Yield Allocation

Process yields directly affect inventory allocation. For semiconductor manufacturing:

$$I_{\text{per good die}} = \frac{I_{\text{per wafer}}}{Y_{\text{line}} \times N_{\text{geometric}} \times Y_{\text{die}}} \tag{2.1}$$

Where:

$I$ Inventory (energy, materials, emissions)

$Y_{\text{line}}$ Line yield (fraction of wafers completing all process steps)

$N_{\text{geometric}}$ Geometric dies per wafer. All allocation sites resolve this count through one shared hierarchy: an entry-authored geometric count where declared; otherwise rectangular packing of the authored die geometry (width, height, scribe/kerf allowance) onto the usable wafer area inside the edge exclusion; falling back to packing of a square-equivalent die where only the die area is known. Plain area division (usable area divided by die area) is not used — it overstates the count by $\sim$10–13 % for typical geometries

$Y_{\text{die}}$ Die yield (fraction of dies passing electrical test)

### 2.4.2 Die Yield Model

Die yield follows the Poisson defect model:

$$Y_{\text{die}} = e^{-D_0 \times A_{\text{die}}} \tag{2.2}$$

Where $D_0$ is the defect density ($1/\text{cm}^2$) and $A_{\text{die}}$ is die area ($\text{cm}^2$).

### 2.4.3 Multi-Output Process Handling

Multi-output processes are rare in electronics manufacturing. When they occur, the following allocation hierarchy applies:

1. **Physical causation**: Where outputs can be physically assigned to specific inputs (preferred)
2. **Mass allocation**: Proportional to output mass (for similar co-products)
3. **Economic allocation**: Proportional to economic value (last resort)

Current applications:

- **Bin yield**: Lower-performance IC grades receive equal inventory per unit as higher grades (no allocation by grade)
- **Scrap recovery**: Precious metal recovery credits are excluded pending data availability
- **Wafer reclaim**: Excluded; test wafers and scrap are outside system boundary

## 2.5 Data Sources Hierarchy

The database uses a tiered hierarchy of public data sources:

### 2.5.1 Source Documentation

Every data value in the database includes:

- Source citation with URL where applicable
- Extraction date
- Data quality indicators (see Chapter 3)

**Table 2.2:** Data source hierarchy by reliability

| Tier | Source Type | Description |
|---|---|---|
| 1 | Equipment specifications | Vendor datasheets with power, throughput, process parameters |
| 2 | Peer-reviewed literature | Academic papers with measured process data |
| 3 | Industry roadmaps | IRDS/ITRS public portions |
| 4 | Corporate reports | Sustainability reports, investor disclosures |
| 5 | Patents | Process descriptions with operational parameters |
| 6 | Engineering estimates | Calculated from first principles with documented assumptions |

- Uncertainty bounds (min/typical/max)

## 2.6 Calculation Methods

### 2.6.1 Equipment-Based Energy Calculation

For equipment-limited processes, energy consumption is calculated from:

$$E_{\text{per wafer}} = P_{\text{equipment}} \times t_{\text{process}} = \frac{P_{\text{equipment}}}{\text{Throughput}} \tag{2.3}$$

Where $P_{\text{equipment}}$ represents the **time-weighted average power consumption** $(\mathrm{kW})$ accounting for the duty cycle between active processing, idle, and standby states. Throughput is the effective wafers per hour achieved under normal production conditions. This formulation implicitly captures the operating state mix for a fully utilized fab; users modeling specific scenarios with different utilization rates should adjust accordingly.

**Example**: EUV lithography

- Time-weighted power: $950\,\mathrm{kW}$ exposing for 65% of the hour, $450\,\mathrm{kW}$ idle for 25%, $70\,\mathrm{kW}$ asleep for 10%, so $(950 \times 0.65) + (450 \times 0.25) + (70 \times 0.10) = 737\,\mathrm{kW}$
- Effective throughput: 152.5 wafers/hour nameplate at the same 65% exposure duty = 99.1 wafers/hour
- Energy per wafer: $737/99.1 = 7.4\,\mathrm{kW \cdot h/wafer}$, published as $7.5$

### 2.6.2 Stoichiometric Material Calculation

For chemical processes, precursor consumption is derived from:

$$M_{\text{precursor}} = \frac{M_{\text{film}}}{\eta \times r_{\text{MW}}} \tag{2.4}$$

Where:

$M_{\textbf{film}}$ Deposited film mass (from thickness $\times$ area $\times$ density)

$\eta$ Utilization efficiency (fraction incorporated into film, typically 10–50%)
$r_{\mathbf{MW}}$ Molecular weight ratio (target element / precursor)

### 2.6.3 Aggregation to Component Level

Process step inventories are aggregated to component level:

$$I_{\text{per wafer}} = \sum_{s} (I_s \times n_s) \times \frac{1}{Y_{\text{line}}} + I_{\text{support}} \tag{2.5}$$

$$I_{\text{per die}} = \frac{I_{\text{per wafer}}}{N_{\text{geometric}} \times Y_{\text{die}}} \tag{2.6}$$

$$I_{\text{per cm}^2} = \frac{I_{\text{per wafer}}}{A_{\text{wafer}}} \tag{2.7}$$

Where $I_s$ is per-pass inventory for step $s$, $n_s$ is the number of passes, and $I_{\text{support}}$ is allocated facility overhead.

For complete integrated circuits, the aggregation combines wafer fabrication and packaging inventories, applying yield allocation at each stage. Figure 2.2 illustrates this hierarchical aggregation from raw materials through to the finished IC.

### 2.6.4 Support Process Allocation

Manufacturing facilities require support processes beyond the direct production equipment. The approach to allocating these support process burdens varies by facility type and production context.

#### 2.6.4.1 Semiconductor Fabrication Facilities

For semiconductor wafer fabs, facility overhead is allocated per wafer based on fab-wide energy consumption divided by production throughput. Support processes for fabs include:

**Cleanroom HVAC** Fan filter units (FFUs), makeup air handling, exhaust systems, and temperature/humidity control for cleanroom environments (ISO Class 1–5).

**Ultrapure water (UPW)** Reverse osmosis, electrodeionization, UV treatment, and polishing systems to produce 18.2 MΩ·cm water required for wafer processing.

**Cooling water systems** Chillers, cooling towers, and heat exchangers for process tool cooling and facility climate control.

**Gas abatement** Point-of-use thermal oxidizers, plasma abatement, and wet scrubbers for treating process exhaust streams.

**Bulk gas supply** On-site generation or delivery of nitrogen, argon, oxygen, and compressed dry air.

**Waste treatment** Wastewater treatment systems for acids, bases, metals, and fluorinated compounds; solid waste handling.

**Water terminology**: Following ISO 14046 conventions (International Organization for Standardization 2014), the database distinguishes between water *withdrawal* (gross intake from the

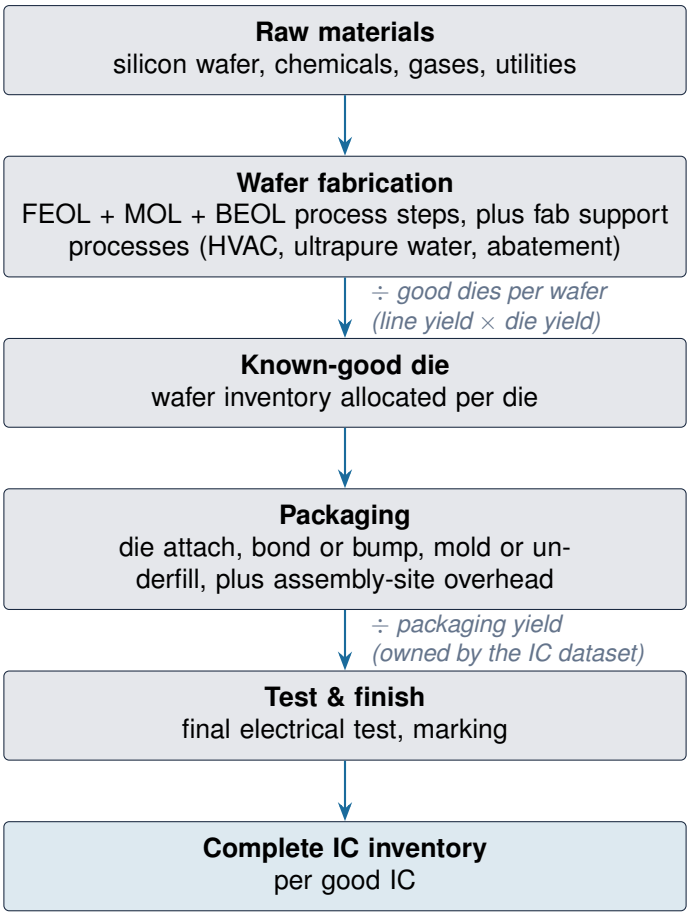


**Figure 2.2:** Aggregation of inventory for a complete integrated circuit. Wafer fabrication inventory is allocated per known-good die, then combined with packaging inventory (grossed by packaging yield) and test/finish operations to produce the complete IC inventory.

source) and water *consumption* (water not returned to the original watershed, primarily evaporative losses from cooling towers and scrubbers). Water that becomes wastewater is classified as *discharge*, not consumption, since it returns to the hydrological system after treatment. The database tracks both withdrawal and consumption metrics.

For advanced logic fabs, facility support processes account for a little over half of total energy consumption (depending on node), with the remainder consumed by process equipment. In v1.0, support energy is allocated per wafer using a fab-type-specific methodology. The cleanroom-HVAC base is modeled as an ISO-6 ballroom with per-tool minienvironments, where the tool fleet is derived from each model's own process recipe for silicon-CMOS logic and DRAM. The conversion from recipe to fleet uses Little's law, the queueing identity $L = \lambda W$: the number of wafers simultaneously in process equals the wafer arrival rate multiplied by the time each wafer spends being processed, which in turn sets the number of tool positions the line must hold — a longer or slower recipe implies a proportionally larger tool fleet. The proportionality constant, which absorbs tool utilization and batch loading, is calibrated on the 5 nm logic reference case. For 3D NAND and specialty fabs, hand-specified tool counts are retained, and LED/optoelectronic fabs take an ISO-7 ballroom override. Layered on this base are make-up air, central exhaust, ultrapure water, cooling water, gas abatement, bulk gas supply, and a sub-fab area uplift. This recipe-driven base replaces an earlier area-weighted heuristic (itself

a refinement of a capacity-scaling lookup), differentiating HVAC node-by-node where the area-weighted method assigned a single value per facility class. The area-weighted base is retained as a documented fallback. The facility-class taxonomy, reference-fab anchoring, utility conversion basis, and regional climate adjustment are defined below; their acknowledged limitations are discussed in Section 6.3.6. Figure 2.3 shows an illustrative breakdown of support process energy for an advanced logic fab.

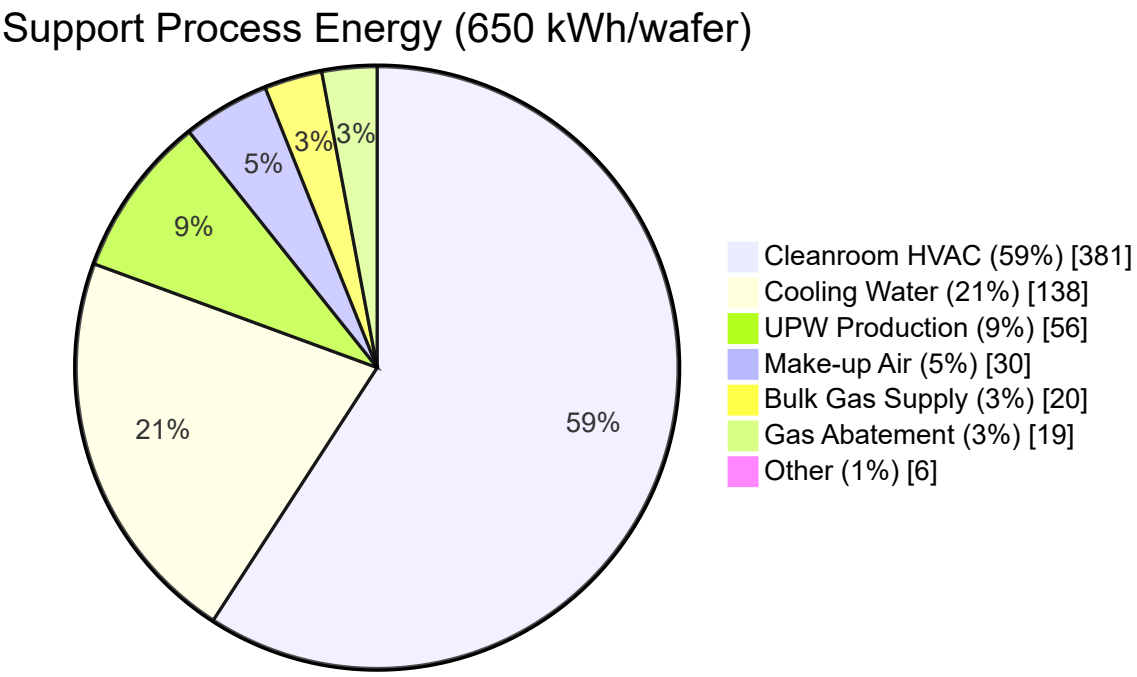


**Figure 2.3:** Illustrative breakdown of support process energy consumption for advanced logic wafer fabrication (node 7 nm, ∼650 kWh/wafer under the recipe-driven HVAC and flat-abatement methodology). Cleanroom HVAC is the largest single category, followed by cooling water, ultrapure water production, and make-up air.

#### 2.6.4.2 Facility Classes and Reference-Fab Anchoring

Support-energy allocation is organized around a fifteen-class facility taxonomy that subdivides logic fabs by node maturity, memory fabs by technology (NAND, DRAM, EUV DRAM), and specialty fabs by device class. Each class carries a cleanroom-area anchor. Seven of the fifteen classes are anchored to publicly disclosed reference-fab cleanroom areas: advanced logic (TSMC Fab 18, 58,000 $m^2$ per phase), HKMG mature logic (GlobalFoundries Fab 8 Malta, 42,735 $m^2$ at 60k wafer starts per month), DRAM (SK Hynix M15X), EUV DRAM (SK Hynix M16), SiC power (Wolfspeed Mohawk Valley), silicon power discretes (Bosch Reutlingen, shared-cleanroom allocation), and 300 mm image sensors (Sony Nagasaki). The remaining classes (NAND, two mature-logic generations, GaN power, III–V photonics, optoelectronic epitaxy, MEMS/analog, and silicon discretes) carry derived anchors, declared as derived in the dataset documentation, with ±30–50 % uncertainty bands.

The area-weighted fallback base — used for NAND and specialty facilities, and retained as the documented fallback for all classes — computes cleanroom HVAC as the per-class cleanroom area multiplied by an ISO-class-weighted power intensity: per-zone intensities from the

PG&E Cleanroom Baseline Study, weighted by the assumed distribution of cleanroom floor area across ISO cleanliness classes. The validation status of that area distribution is discussed in Section 6.3.6.

#### 2.6.4.3 Utility Conversion Basis

Utility flows follow the SEMI S23-1021 energy-conversion-factor framework, with independently anchored values retained where a better direct source exists: ultrapure water uses the IRDS 2024 baseline and compressed dry air a semiconductor-specific benchmark, while vacuum, process cooling water, high-pressure CDA, and make-up-air heat removal adopt the SEMI S23 factors. Make-up air conditioning is anchored to a published cleanroom make-up-air-handling range (8–30 kWh/m$^2$/month, central 20). Central exhaust is a fan-only figure (Sun 2017, 0.0007 kWh/m$^3$; independently corroborated by Hu et al. 2019 at 0.000723 kWh/m$^3$); scrubber and abatement energy are tracked as a separate flow. Exhaust-to-outside-air ratios are facility-class-specific (logic 0.85, memory 0.78, III–V/SiC/MEMS 0.92, image sensors 0.78), and sub-fab floor area is accounted at 1.5$\times$ the cleanroom area at 65 W/m$^2$.

No separate per-process-pass HVAC term is carried: an earlier per-pass charge was derived from the same fab benchmark the area-driven base computes, so carrying both would double-count. The zone power intensities underlying the base were measured in operating fabs running production recipes, so process-driven exhaust and bay-occupancy effects are already reflected in them.

Ultrapure-water demand is estimated from per-pass rates for wet cleaning, lithography, and CMP, with the wet-clean share set from published fab water balances.

#### 2.6.4.4 Tool Utilization

An optional tool-utilization parameter (default 1.0, i.e. full utilization) inflates the fixed facility loads — cleanroom HVAC, make-up air, central exhaust, gas-cabinet, and bulk gas supply — to reflect amortization over fewer wafer starts in ramping or low-yield fabs (mature high-volume manufacturing typically operates at $\sim$0.85–0.90; ramping fabs at $\sim$0.40–0.70). Per-pass flows scale naturally with wafer starts and are not inflated.

#### 2.6.4.5 Regional Climate Adjustment

Cleanroom HVAC support energy is climate-adjusted per manufacturing region. Make-up air conditioning is split into a climate-insensitive dehumidification floor (50 %, per Yin 2020) and a climate-modulated portion; the modulated portion carries a per-region multiplier derived from heating- and cooling-degree-day comparisons against the Taiwan baseline: Taiwan 1.00, Arizona 1.10, Oregon 0.75, Ireland 0.80, Korea 0.95, China 1.00, Germany 0.70, Japan 0.90. Every dataset's support energy therefore reflects the climate of its declared manufacturing region, and geography-variant datasets re-apply the multiplier for their own region rather than inheriting the baseline's climate energy.

#### 2.6.4.6 PCB Manufacturing Facilities

PCB facilities operate without cleanroom requirements and have different support process profiles. Key support processes include:

**Process water treatment** Deionized water for rinse tanks, typically lower purity than semiconductor UPW.

**Wastewater treatment** Significant treatment requirements for copper-laden rinse water, spent etchants, and plating bath effluents.

**Fume extraction** Local exhaust ventilation for etch, plating, and lamination processes.

**Compressed air** For pneumatic equipment and process applications.

Support energy for PCBs is allocated per square meter of board produced.

#### 2.6.4.7 Passive Component Manufacturing

Passive component facilities (MLCCs, resistors, inductors) have support requirements that vary significantly by product type:

**Cleanroom HVAC** ISO 6–7 rooms for tape casting, electrode printing, and stacking, with tight humidity control.

**Sintering furnaces** High-temperature kilns dominate energy consumption for ceramic components.

**Controlled atmospheres** Nitrogen or nitrogen/hydrogen atmospheres for sintering and soldering.

**Plating systems** Electroplating for terminations with associated water treatment and rinse water.

Facility support for passives is included in the shipped datasets as separately cited inventory rows — cleanroom HVAC, cooling water, waste treatment, and bulk gas supply electricity, plus cooling-tower make-up water — allocated per piece or per kilogram of product depending on the component type. The cleanroom-HVAC allocation factor is derived from published cleanroom air-change-rate ratios; the remaining applicability factors adapt the facility-class intensities of Section 2.6.4 to passives plants, with the basis cited on the rows themselves. Where a support process is genuinely not applicable — ultrapure water and gas abatement for MLCC production, which runs no wet chemistry before termination plating — the allocation carries an explicit zero rather than a silent omission, and where a model's sintering process file already accounts for atmosphere generation, the facility bulk-gas allocation is zeroed to prevent double counting. Compressed dry air and facility process-gas mass are declared deferrals pending an allocation driver (Section 6.2.3).

#### 2.6.4.8 Packaging and Assembly Facilities

OSAT (Outsourced Semiconductor Assembly and Test) facilities have intermediate requirements between wafer fabs and PCB facilities:

**Cleanroom (reduced)** Class 100–10,000 cleanrooms for die attach and wire bonding; less stringent than wafer fabs.

**Cooling water** For molding equipment and test systems.

**Compressed air** Pneumatic handling and packaging equipment.

Wire-bond and molded packaging processes do not use fluorinated etch or clean gases, so no

PFC abatement is modeled for them. The exception is 2.5D/3D interposer fabrication (CoWoS, EMIB, Foveros, HBM TSV): their deep-silicon Bosch etch steps consume $SF_6$ and $C_4F_8$, which appear as gas inputs on those datasets; residual F-gas emissions for these packaging-side etches are a declared gap in v1.0 (the gas-fate layer currently runs for wafer-fabrication models only) and are targeted for the next release.

#### 2.6.4.9 Production Facility Assumptions

Support process burdens are allocated based on reference facility parameters. Table 2.3 documents the production capacity assumptions used for different facility types.

**Table 2.3:** Reference facility parameters for support process allocation

| Facility Type | Capacity | Wafer Size | Default fab-type bucket |
|---|---|---|---|
| Advanced logic (3–7nm) | 50,000 wpm | 300mm | `advanced_logic` |
| Mature logic (28nm+) | 80,000 wpm | 300mm | `mature_logic_hkmg` |
| Memory (DRAM/NAND) | 120,000–145,000 wpm | 300mm | `memory_dram` / `memory_nand` |
| Legacy logic (90–180nm) | 80,000 wpm | 300mm; 200mm at 180nm | `mature_logic_90_180nm` |

These parameters represent typical high-volume manufacturing facilities operating at steady-state production. Each facility type maps to a fab-type bucket whose recipe-derived tool fleet (for logic and DRAM) or hand-specified tool count (for NAND and specialty), cleanroom envelope, and utility envelope drive the per-wafer support-energy calculation. As a representative result, total advanced-logic wafer energy in v1.0 ranges from approximately 1,150 kWh/wafer (7 nm) to approximately 1,525 kWh/wafer (3 nm), rising to approximately 1,575 kWh/wafer for the Arizona 3 nm dataset, of which facility support processes account for a little over half; the recipe-driven cleanroom-HVAC base and SEMI S23-aligned utility envelope are defined in Section 2.6.4, with their acknowledged limitations in Chapter 6. Point-of-use gas-abatement electricity is allocated as a flat per-wafer fleet load (a small, fab-type-differentiated term), not a per-process-pass cost; the fluorinated-gas treatment burden of etch-intensive processes is captured instead in the residual F-gas emission flows.

**Note on utilization**: Real fabrication facilities operate at less than 100% of rated capacity due to equipment maintenance, changeovers, and demand fluctuations. The default inventory assumes mature high-volume operation. For fabs running well below rated capacity (ramping or greenfield lines), an optional tool-utilization factor amortizes fixed facility overhead across fewer wafer-starts, increasing the per-wafer support burden; this factor defaults to unity for steady-state high-volume production (see Section 2.6.4.4).

## 2.7 Custom Upstream Material Datasets

Several semiconductor-specific materials are not available in standard LCI databases such as ecoinvent. For these materials, the REEL database provides custom upstream datasets with full gate-to-gate inventories. These datasets enable complete supply chain tracing for specialty process chemicals.

### 2.7.1 Materials Covered

Table 2.4 lists the custom upstream material datasets included in the database.

**Table 2.4:** Custom upstream material datasets

| Material | Formula | Use in Semiconductor Manufacturing |
|---|---|---|
| Nitrogen trifluoride | $NF_3$ | CVD/PVD chamber cleaning gas |
| Octafluorocyclobutane | $C_4F_8$ | Oxide/nitride etch gas |
| Hexafluoro-1,3-butadiene | $C_4F_6$ | High-aspect-ratio etch gas |
| Tungsten hexafluoride | $WF_6$ | Tungsten CVD precursor |
| Elemental fluorine | $F_2$ | Remote plasma chamber clean, base for F-gases |
| ArF photoresist | (mixture) | 193nm DUV lithography |
| TMAH | $(CH_3)_4NOH$ | Photoresist developer active compound (pure, 100% basis) |
| TMAH developer solution | 2.38% aq. | Ready-to-use 0.26 N developer as consumed by lithography steps; carries the dilution (0.0238 kg pure TMAH + 0.9762 kg deionised water per kg solution) so consumers link solution mass directly |
| CMP slurries | (mixtures) | Oxide, copper, tungsten planarization |
| Carbon tetrafluoride | $CF_4$ | Plasma etch gas (oxide, nitride, silicon) |
| Germane | $GeH_4$ | SiGe epitaxy precursor (PMOS stressor) |
| Trimethylgallium | $Ga(CH_3)_3$ | MOCVD Ga source (GaAs, GaN, InGaAs) |
| Trimethylaluminum | $Al(CH_3)_3$ | MOCVD Al source (AlGaAs, AlGaN) |
| Trimethylindium | $In(CH_3)_3$ | MOCVD In source (InP, InGaAs, InGaN) |
| Sulfur hexafluoride | $SF_6$ | Plasma etch gas, chamber clean (high GWP) |
| Ultrapure water | $H_2O$ | Wafer rinsing, wet chemistry, CMP |
| Polyimide (proxy) | PMDA/ODA | Passivation layers, flex substrates, RDL dielectrics |

The CMP slurries entry comprises three datasets (copper, oxide, and tungsten slurries) and the ultrapure water entry comprises two (process ultrapure water and a separate municipal water supply dataset), so the 17 material types listed above are exported as 20 upstream datasets, consistent with the coverage table in Chapter 1.

### 2.7.2 Methodology for Custom Datasets

Each upstream material dataset follows a consistent methodology:

1. **Stoichiometric calculation**: Material inputs are derived from balanced chemical equations for the primary production route.
2. **Literature energy values**: Process energy is sourced from academic literature, patents,

or engineering estimates based on thermodynamic requirements.

3. **ecoinvent linkage**: Where possible, input materials reference ecoinvent datasets. Materials not in ecoinvent are either proxied to similar compounds or linked to other custom datasets. See Section 5.3.2.4 for the full list of proxy assumptions and their limitations.
4. **Yield adjustment**: Production yields (typically 70–95%) account for byproduct formation and process losses.

### 2.7.3 Data Quality Considerations

These custom datasets rely on calculated values rather than measured plant data, which limits their data quality relative to inventories built from operating-plant measurements. Key limitations include:

- Energy values are often theoretical minimums or engineering estimates rather than measured industrial consumption
- Proprietary formulations (especially for photoresists and CMP slurries) require proxy compositions
- Regional variation in production processes is not captured

Users requiring higher fidelity upstream data should consult specialty chemical manufacturers directly or commission facility-specific studies.

### 2.7.4 MOCVD Precursor Production

Metal-organic chemical vapor deposition (MOCVD) precursors require specialized upstream datasets due to their absence from standard LCI databases. These pyrophoric organometallics are critical for III-V compound semiconductor fabrication (GaAs, GaN, InP epitaxy) used in LEDs, laser diodes, and RF devices.

**Production routes modeled:**

- **Trimethylgallium (TMGa)**: Alkyl exchange reaction between $GaCl_3$ and TMAl at elevated temperature. 90% yield, 5.5 kWh/kg process energy.
- **Trimethylaluminum (TMAl)**: Grosse-Mavity process—sodium reduction of dimethylaluminum chloride (DMAC). 89% yield, 7.0 kWh/kg.
- **Trimethylindium (TMIn)**: $InCl_3$ + TMAl with KF complexing agent, followed by vacuum sublimation purification. 65% yield, 14 kWh/kg (additional energy for sublimation).

All three precursors are pyrophoric and stored under inert atmosphere. Data sources include industrial patents and academic literature on III-V precursor manufacturing.

## 2.8 Elementary Flow Classification and Emission Transformations

Semiconductor manufacturing consumes a wide range of specialty gases and chemicals, many of which are not present in standard elementary flow databases such as ecoinvent. Rather than report raw process inputs as atmospheric or aqueous emissions—which would misrepresent

actual environmental releases—the database applies chemical transformation rules that reflect the abatement and treatment processes through which these substances pass before reaching the environment.

This section documents the classification logic, transformation chemistry, and cut-off decisions applied to convert process-level substance inventories into properly characterized elementary flows.

### 2.8.1 Classification of Output Flows

Output flows from manufacturing processes are classified into two categories following ISO 14044 conventions:

**Technosphere outputs** Flows directed to further treatment within the industrial system. These include wastewater sent to on-site or municipal treatment plants, and solid waste sent to disposal or recycling facilities. Technosphere outputs reference specific treatment datasets rather than direct environmental compartments.

**Elementary outputs** Flows released directly to the natural environment after all abatement and treatment. These are the emissions that enter environmental compartments (air, water) and are relevant for life cycle impact assessment.

The distinction is critical: a substance consumed in the process chamber is not necessarily an emission. It must pass through abatement systems, scrubbers, and waste treatment before reaching the environment, often undergoing chemical transformation in the process.

### 2.8.2 Non-Elementary Flow Exclusions

Several categories of flows that appear in process inventories are excluded from elementary outputs because they are not environmental emissions:

- **Inert bulk gases**: $N_2$, Ar, $O_2$, and compressed dry air (CDA) are used as carrier gases, purge gases, and process atmospheres. These are ambient atmospheric constituents and their release does not alter environmental compartment composition.
- **Wastewater volumes**: Volumetric wastewater entries (in L or $m^3$) are tracked as technosphere outputs directed to treatment. Only dissolved contaminant masses (in g or kg) are reported as elementary emissions to water.

### 2.8.3 Abatement-Destroyed Precursors

CVD and ALD processes consume reactive precursor gases that are pyrophoric, hydrolytically unstable, or thermally decomposed at relatively low temperatures. Modern semiconductor fabs employ point-of-use (POU) abatement on every process exhaust line, using thermal oxidation (burn-wet systems operating at 750–1000$^\circ$C) or plasma-based destruction.

For these reactive precursors, the destruction and removal efficiency (DRE) approaches 100%, and their combustion or hydrolysis products (metal oxides, $SiO_2$) are captured in downstream wet scrubbers as solid particulate. The raw precursor form never reaches the atmosphere.

**Table 2.5:** Reactive precursors excluded from elementary emissions (DRE ≈ 100%)

| Substance | Formula | Process Role | Abatement Product |
|---|---|---|---|
| Silane | $SiH_4$ | CVD Si/SiN deposition | $SiO_2$ (scrubber solids) |
| TEOS | $Si(OC_2H_5)_4$ | CVD oxide deposition | $SiO_2$ + $CO_2$ + $H_2O$ |
| Germane | $GeH_4$ | SiGe epitaxy | $GeO_2$ (scrubber solids) |
| OMCTS | $C_8H_{24}O_4Si_4$ | Low-k CVD | $SiO_2$ + $CO_2$ + $H_2O$ |
| TEMAH | $Hf[N(CH_3)(C_2H_5)]_4$ | ALD high-k gate | $HfO_2$ (scrubber solids) |
| TMGa | $Ga(CH_3)_3$ | MOCVD Ga source | $Ga_2O_3$ (scrubber solids) |
| TMAl | $Al(CH_3)_3$ | MOCVD Al source | $Al_2O_3$ (scrubber solids) |

Table 2.5 lists the substances excluded from elementary emissions on this basis.

This exclusion is conservative: the abatement products (metal oxide particulates) are captured in scrubber water and removed as solid waste via the technosphere waste treatment pathway, not released as elementary emissions. The associated $CO_2$ from organic ligand combustion is negligible relative to grid electricity emissions (typically $<$0.01% of total carbon footprint).

### 2.8.4 Stoichiometric Emission Proxies

Several process gases undergo chemical transformation in point-of-use abatement or in the scrubber before reaching the environment. Rather than report the parent compound (which has no ecoinvent elementary flow equivalent), the database applies stoichiometric proxy rules to report the species actually emitted, and routes each to the environmental compartment it reaches — air, water, or a cross-compartment split between the two.

#### 2.8.4.1 Reactive-Precursor Abatement and Halide Partitioning

Water-reactive and soluble halide precursors — $WF_6$, dichlorosilane (DCS), and $BCl_3$ — are hydrolyzed or combusted in point-of-use abatement, forming hydrogen halides that are then largely captured in the wet scrubber. The transformation stoichiometry is:

**Table 2.6:** Stoichiometric transformation of reactive halide precursors

| Parent | Reaction | Transformed Product | Mass Fraction |
|---|---|---|---|
| $WF_6$ | $WF_6 + 3H_2O \rightarrow WO_3 + 6HF$ | Hydrogen fluoride | 0.403 |
| DCS | $SiH_2Cl_2 + O_2 \rightarrow SiO_2 + 2HCl$ | Hydrogen chloride | 0.722 |
| $BCl_3$ | $BCl_3 + 3H_2O \rightarrow B(OH)_3 + 3HCl$ | Hydrogen chloride | 0.934 |

The mass fraction is the stoichiometric yield of transformed product per gram of parent. The v1.0 gas-fate refinement partitions that product explicitly across environmental compartments. $WF_6$ now carries an explicit abatement destruction-removal efficiency (DRE $\approx$ 0.98–0.99), so only a $\sim$1 % post-scrubber slip reaches the exhaust stack; the HF proxy is applied to that residual air slip, while the bulk fluorine is retained as the authored aqueous $F^-$ inventory (0.084 g per deposition pass for CVD tungsten) rather than a computed proxy, avoiding a double-count, and the tungsten is captured as $WO_3$ particulate in scrubber water. DCS and $BCl_3$ are instead handled as cross-compartment proxy splits: $\sim$2 % of the transformed HCl slips to air and the remaining $\sim$98 % of the chlorine is reported as aqueous $Cl^-$, which keeps the chlorine mass in

the inventory for acidification completeness. The fully-abated reactive precursors ($SiH_4$, $GeH_4$, $B_2H_6$, $PH_3$, $AsH_3$, $BF_3$) carry the same high-DRE abatement, and parent-compound name aliases are normalized so no raw reactive precursor survives to the atmospheric elementary outputs. These fate refinements are GWP-neutral (all affected species have GWP $\approx 0$); they revise the acidification and toxicity inventory only.

Note that these proxy rules apply *after* the PFC emission model (Equation 4.2 in Chapter 4), which first determines how much gas reaches the abatement system based on chamber utilization and recovery rates.

#### 2.8.4.2 Air-to-Water Transformation

Hydrogen bromide (HBr), used extensively in silicon and polysilicon plasma etching, is a highly water-soluble acid gas. In wet scrubber systems, HBr is captured with near-complete efficiency ($>$99%) and dissolves as bromide ions:

$$HBr_{(g)} \xrightarrow{\text{scrubber}} H^+ + Br^-_{(aq)} \tag{2.8}$$

The scrubber effluent containing $Br^-$ is directed to wastewater treatment, making this a cross-compartment transfer: the substance enters as an air emission but is reported as a water emission (bromide ion, mass fraction 0.988 of parent HBr).

#### 2.8.4.3 Water-to-Water Dissociation

Compound salts and ions reported in process wastewater streams dissociate in aqueous solution. The database reports the constituent ions rather than the parent compound, as these are the species that interact with aquatic ecosystems:

**Table 2.7:** Ionic dissociation proxy rules for water emissions

| Parent | Dissociation Products | Mass Fractions | |
|---|---|---|---|
| $CuCl_2$ | $Cu^{2+} + 2Cl^-$ | 0.473 (Cu) | 0.527 (Cl) |
| $K_2SO_4$ | $2K^+ + SO_4{}^{2-}$ | 0.449 (K) | 0.551 ($SO_4$) |
| $HPO_3{}^{2-}$ | $PO_4{}^{3-}$ (oxidized) | 1.462 (mass gain from oxidation) | |

For $CuCl_2$ and $K_2SO_4$, the mass fractions sum to 1.0 (complete mass conservation). For phosphite ($HPO_3{}^{2-}$), the mass fraction exceeds 1.0 because oxidation to phosphate ($PO_4{}^{3-}$) incorporates oxygen from water, increasing the mass of the emitted species.

$CuCl_2$ arises from copper etch processes and CMP waste streams. $K_2SO_4$ is a byproduct of CMP slurry neutralization. Phosphite originates from electroless plating bath chemistry and is oxidized to phosphate during wastewater treatment.

### 2.8.5 Fluorocarbon Emission Proxy

Octafluorocyclobutane ($C_4F_8$, CAS 115-25-3, GWP = 10,200) is used extensively in oxide and nitride etching but does not have a dedicated elementary flow entry in standard LCI databases. To enable database linkage, $C_4F_8$ is proxied to tetrafluoromethane ($CF_4$), the closest available

fluorocarbon flow. Exports retain the actual substance name, CAS number, and chemical formula alongside the $CF_4$ UUID. The PFC emission model (Equation 4.2) is applied to calculate the post-abatement emission quantity.

Hexafluoro-1,3-butadiene ($C_4F_6$), also used in advanced etching, has a negligible GWP due to its short atmospheric lifetime and has been cut off from the emission inventory.

# Chapter 3

# Data Quality and Uncertainty

This chapter describes the data quality assessment framework and uncertainty characterization methods used throughout the database.

## 3.1 Data Quality Indicators

### 3.1.1 Modified Pedigree Matrix

Data quality is assessed using a modified pedigree matrix adapted for virtual factory modeling (Ciroth et al. 2016; Weidema et al. 2013). Unlike traditional LCI databases that can verify primary data collection, this database relies on secondary public data, requiring adapted quality criteria.

**Table 3.1:** Data quality indicator dimensions

| Dimension | Description |
|---|---|
| Reliability | Quality of the underlying measurement or calculation method |
| Completeness | Representativeness of the sample for the intended scope |
| Temporal | Currency of the data relative to technology evolution |
| Geographic | Alignment with major production regions |
| Technological | Specificity to the exact process or technology |

### 3.1.2 Scoring Criteria

Each dimension is scored on a 1–5 scale (1 = best quality):

#### 3.1.2.1 Reliability

| Score | Criteria |
|---|---|
| 1 | Verified process data from peer-reviewed literature or detailed vendor specifications |
| 2 | Calculated from well-documented equipment specifications and stoichiometry |
| 3 | Calculated from theoretical physics with reasonable assumptions |
| 4 | Extrapolated from related processes or older technology nodes |
| 5 | Rough engineering estimate or rule of thumb |

#### 3.1.2.2 Completeness

| Score | Criteria |
|---|---|
| 1 | Data covers all known process steps and flows >1% of mass/energy |
| 2 | Major process steps covered; minor flows estimated |
| 3 | Covers major energy consumers; minor chemical flows estimated |
| 4 | Partial coverage; significant flows estimated or missing |
| 5 | Representative proxy only; substantial uncertainty |

#### 3.1.2.3 Temporal Correlation

| Score | Criteria |
|---|---|
| 1 | Data/specification <3 years old |
| 2 | Data 3–6 years old |
| 3 | Data 6–10 years old |
| 4 | Data 10–15 years old |
| 5 | Data >15 years old or of unknown vintage (significant for rapidly evolving technologies) |

**Source-year discipline** An authored temporal label records the vintage of the *evidence record* — when the underlying data were assembled and verified — which is not always the publication year of the original source. The scoring therefore disciplines authored labels against the cited sources themselves: where a source citation carries an explicit publication year, the year-derived score governs whenever it is *worse* than the authored label, and where a file's source list carries several dated citations, the median of their year-derived scores caps the file's label. The cap only ever worsens a score, so an author's own conservatism is preserved. Two source classes receive tailored year handling: equipment specifications are scored on their most recent year (a current specification for equipment still in production is current data), and secondary or aggregator sources are scored on the oldest year they cite (the vintage of the original data, not of the compilation). Citations that carry no year at all retain the authored label; deepening this discipline to registry-resolved publication years for every individual citation is a declared refinement direction (Section 6.3.9).

#### 3.1.2.4 Geographic Correlation

| Score | Criteria |
|---|---|
| 1 | Specific to major production region (Taiwan, South Korea, US, EU) |
| 2 | Global average representing production-weighted mix |
| 3 | Proxy region used with grid/efficiency adjustment |
| 4 | Unspecified region; assumed global average |
| 5 | Unknown geography; potential significant variance |

#### 3.1.2.5 Technological Correlation

The Technological dimension is composited from three sub-axes: process specificity (same chemistry / unit operation), technology class (same generation or variant), and manufacturing scale (HVM versus research/pilot). The final Technological score is taken as the **worst (highest)** of the three sub-axes (weakest-link compositing, not an average). A perfect process-chemistry match drawn from a university lab cannot score better than the scale axis allows.

**Process-specificity axis**

| Score | Criteria |
|---|---|
| 1 | Exact match to technology node/process type |
| 2 | Adjacent node data with documented scaling |
| 3 | Similar technology class used as proxy |
| 4 | Related technology with significant assumptions |
| 5 | Different technology class used as proxy |

**Manufacturing-scale axis** Many of the database's process inventories are derived from academic papers, MDPI articles, university lab measurements, or pilot lines. Downstream users need to see whether a number comes from a real fab or a research bench so they can weight uncertainty appropriately. The scale sub-axis makes this distinction explicit:

| Score | Criteria |
|---|---|
| 1 | HVM-validated for the exact process variant (vendor sustainability report, fab disclosure, IEDM/VLSI paper from a production line) |
| 2 | HVM-validated for closely-adjacent variant (e.g. 7nm data applied to 5nm; 200mm extrapolated to 300mm with documented scaling) |
| 3 | Pilot-line or vendor-claimed HVM, unverified (vendor whitepaper without independent confirmation; equipment-OEM-stated numbers) |
| 4 | Research-scale (academic paper, university lab, MDPI / preprint / arXiv; 6-inch / 100 mm lab tools where industry uses 300 mm) |
| 5 | Bench-scale or unrelated process used as proxy (single-coupon demos; analog process from a different industry) |

**Compositing rule** The composite Technological score is the maximum of the three sub-axes:

$$Te = \max(\text{process\_specificity},\ \text{technology\_class},\ \text{scale}) \tag{3.1}$$

For example, an MDPI study reporting the exact PECVD oxide chemistry on a 6-inch research-scale chamber scores process-specificity = 1, technology-class = 1, scale = 4, so $Te = \max(1, 1, 4) = 4$. The scale axis dominates and the data point is flagged as research-grade despite the chemistry match.

**Recording in process files** Each process and data file records its technological-representativeness score in its data-quality block. The weakest-link compositing is applied *at authoring time*: where the manufacturing-scale or technology-class sub-axis worsens the score, the author records the composited numeric score (1–5) directly, and the automated scorer honors a numeric entry as-is. A controlled word vocabulary ("specific process", "technology class", ...) remains supported for files annotated descriptively, each word mapping to its process-specificity score; such files are scored on that axis alone, without automated re-compositing. Files whose inventories draw on research or pilot-scale sources additionally document the scale basis in prose for human readers.

### 3.1.3 Composite DQI Reporting

Data quality indicators are reported as a **vector of individual scores** (R, C, T, G, Te) to preserve transparency, along with a **geometric mean** as the summary indicator:

$$\text{DQI}_{\text{summary}} = \sqrt[5]{R \times C \times T \times G \times Te} \tag{3.2}$$

Where $R$ = Reliability, $C$ = Completeness, $T$ = Temporal, $G$ = Geographic, $Te$ = Technological.

The geometric mean provides a balanced summary that is sensitive to poor scores in any dimension while avoiding the extreme conservatism of using the maximum. For example, data that is verified (R=1) but from a different technology class (Te=5) with other dimensions at 2 would receive a summary DQI of $\sqrt[5]{1 \times 2 \times 2 \times 2 \times 5} \approx 2.1$, indicating acceptable quality with caution advised.

**Note**: The DQI is **purely qualitative metadata** for user assessment. It is not computationally linked to the quantitative min/typical/max uncertainty bounds, which are determined independently based on source documentation and engineering judgment.

**Configuration-weighted compositing** A dataset's composite is not a flat average over its scored source inventories. Each scored inventory is weighted by how much of the dataset's calculation it actually carries — for process inventories, by the number of manufacturing operations that read the file; for aggregation datasets, by each constituent's share of the modelled cradle-to-gate energy demand. Configuration-layer files (facility assumptions, support-process rates, scenario parameters) are part of the scored population on the same basis: where a dataset's quantities are substantially driven by configuration-layer assumptions rather than process-specific evidence, the composite reflects the generic-technology grade of those assumptions instead of silently ignoring them. Each dataset discloses its weighting basis alongside the pedigree scores, so a reader can see when configuration-layer files dominate the scored weight.

**Display precision** Published composite and per-dimension pedigree scores are displayed at one decimal place. The underlying calculation carries full precision internally; one decimal is the honest display resolution for a five-way geometric mean of ordinal 1–5 grades, and it

prevents readers from ranking datasets on second-decimal differences that the scoring method cannot support.

**Table 3.2:** DQI interpretation guide

| DQI Score | Interpretation |
|---|---|
| 1 | High quality; suitable for detailed assessments |
| 2 | Acceptable quality; appropriate for most applications |
| 3 | Use with caution; consider sensitivity analysis |
| 4–5 | Screening quality only; significant uncertainty |

## 3.2 Uncertainty Characterization

### 3.2.1 Min/Typical/Max Representation

Quantitative data is characterized using three-point estimates wherever a range has been established for the value; coverage is accounted per row and is not yet complete (Sections 3.4.4 and 6.3.8):

**Minimum** Lower bound based on best-case assumptions or lowest reported values
**Typical** Central estimate representing expected or most common values
**Maximum** Upper bound based on worst-case assumptions or highest reported values

This representation captures both:

- **Variability**: Real differences between facilities, equipment generations, operating conditions
- **Uncertainty**: Lack of precise data requiring estimation

### 3.2.2 Variability vs. Uncertainty

The database distinguishes between:

**Variability** (aleatory uncertainty):

- Inherent differences between fabs, equipment, operating conditions
- Cannot be reduced with more data; reflects real-world variation
- Example: Energy consumption varies 15–25% across fabs at same node

**Uncertainty** (epistemic uncertainty):

- Lack of precise measurement or complete information
- Can be reduced with better data or more detailed models
- Example: Gas utilization efficiency not reported; estimated from physics

## 3.3 Uncertainty Propagation

### 3.3.1 Interval Arithmetic Approach

The calculation engine propagates uncertainty using interval arithmetic:

$$[a, b] + [c, d] = [a + c, b + d] \tag{3.3}$$

$$[a, b] \times [c, d] = [\min(ac, ad, bc, bd), \max(ac, ad, bc, bd)] \tag{3.4}$$

For positive intervals (typical in LCI):

$$[a, b] \times [c, d] = [a \times c, b \times d] \quad \text{where } a, c > 0 \tag{3.5}$$

**Important limitation**: Interval arithmetic is computationally simple but produces **theoretical worst-case bounds**, not probability distributions. When aggregating many process steps (as in semiconductor manufacturing with hundreds of steps), interval bounds expand conservatively and may be wider than would be observed in practice. Users requiring probabilistic uncertainty estimates should apply Monte Carlo sampling to the min/typical/max values using triangular distributions.

### 3.3.2 Propagation Rules

Uncertainty intervals are propagated through:

- Addition with interval bounds
- Multiplication and division with bound propagation (corner-product bounds, Equation 3.4)
- Aggregation across process steps
- Yield-adjusted allocation

**Deliberate exclusion of interval subtraction**: no interval subtraction is performed. Interval subtraction over-widens results and can produce negative lower bounds for physically non-negative quantities (e.g., net water balances). Wherever a subtractive relationship occurs, it is computed on central values, and the band derives from the surviving additive and multiplicative operands.

**Example**: Aggregating two process steps with uncertainty:

$$\begin{aligned}
&\text{Step 1 energy} = [5.0, 6.0, 8.0] \text{ kWh (min/typ/max)} \\
&\text{Step 2 energy} = [3.0, 4.0, 5.0] \text{ kWh} \\
&\quad \text{Total energy} = [5.0 + 3.0, 6.0 + 4.0, 8.0 + 5.0] = [8.0, 10.0, 13.0] \text{ kWh}
\end{aligned}$$

### 3.3.3 Published Band Provenance

Every published headline uncertainty band records how it was derived, so that genuinely propagated intervals are never conflated with synthetic fallbacks. A band is one of two kinds: either it was *propagated* end-to-end through the interval arithmetic above, from min/max ranges authored in the underlying process, configuration, and upstream data files; or, where no propa-

gated interval with real width was available for the field, it is a *category default* — a symmetric percentage assigned by component category.

Zero-valued flows carry no band. Beyond the headline energy and water fields, band coverage is accounted *per inventory row*: each row either carries an explicit interval, is exempt on record (it belongs to an enumerated class for which a band would restate a convention or duplicate a linked dataset's own interval), or is counted as uncovered — the same per-row accounting as for citations (Section 3.4.4). Every published band traces to a parent range or an authored basis; none is authored from nothing, and band additions are verified not to move any central value. At emission time, the writer additionally enforces that every published band brackets its own central value after output rounding, and a propagated band that collapses to zero width under rounding falls back to the category default rather than being published as a degenerate interval. Coverage limitations of this scheme are declared in Section 6.3.8.

### 3.3.4 Priority Flows for Robust Modeling

Detailed uncertainty modeling focuses on flows that typically dominate environmental profiles in electronics manufacturing. Based on literature review of published electronics LCA studies, the following flow categories warrant the most careful uncertainty characterization:

**Electricity consumption** Often the largest single contributor to manufacturing environmental footprint when characterized for climate change. Uncertainty in electricity consumption propagates directly to characterized results.

**Water consumption** Particularly significant for semiconductor manufacturing where ultrapure water production is energy-intensive and water stress varies geographically.

**Specialty gases** Perfluorocarbons ($CF_4$, $C_2F_6$, $SF_6$) and nitrogen trifluoride ($NF_3$) have high global warming potentials, making even small mass flows potentially significant.

**Precious metals** Gold and palladium in packaging have high embodied impacts per unit mass from mining and refining.

**Solvents and specialty chemicals** Volatile organic compounds in photoresists and cleaning solvents can be significant for air quality impacts.

Flows contributing less than 1% of mass and energy use simplified uncertainty factors unless they involve substances of particular environmental concern.

## 3.4 Provenance and Source Documentation

This section consolidates how sources are documented, classified, and bound to the published data: the metadata required for every value, the source taxonomy, the per-flow citation mechanism carried on exported datasets, and the guidelines for authoring uncertainty ranges. The quality scoring built on this documentation is described in Section 3.1; residual limitations of citation-binding depth are declared in Chapter 6.

### 3.4.1 Required Metadata

Every data value is documented at authoring time against the following fields, recorded in the process and data files behind it. Published datasets carry that provenance as per-flow source citations (Section 3.4.3):

**Table 3.3:** Required metadata fields for data values

| Field | Type | Description |
|---|---|---|
| `value` | Numeric | The data value |
| `unit` | String | Physical unit (SI preferred) |
| `source_type` | Enum | Category from source taxonomy |
| `source_title` | String | Document title or publication name |
| `source_url` | URL | Direct link if publicly available |
| `source_date` | Date | Publication or measurement date |
| `extraction_date` | Date | When data was extracted |
| `page_or_section` | String | Location within source document |
| `context` | String | Conditions, scope, or caveats |

### 3.4.2 Source Type Taxonomy

Data sources are classified using a standardized taxonomy:

**Equipment specifications** Vendor technical documentation, datasheets
**Academic papers** Peer-reviewed journal publications
**Conference papers** Conference proceedings (IEDM, VLSI, ECTC)
**Corporate reports** Sustainability reports, annual reports, CDP disclosures
**Industry standards** SEMI, IEEE, or other standards body documents
**Patents** Patent filings with technical specifications
**Trade publications** Trade magazines, industry news
**Government reports** EPA, DOE, EU JRC publications
**Calculated** Derived from other documented data
**Estimated** Engineering estimate with documented basis
**Assumed** Explicit assumption (requires justification)

### 3.4.3 Per-Flow Source Citations

Each flow in a published dataset lists the sources behind its value, wherever those sources have been resolved (the residual is tracked by the coverage accounting below). Where a value is aggregated from many contributing process inventories — electricity, water, and roll-up waste and emission rows — the sources of those contributors travel with the flow and are labelled as inherited. Where a source was established for a dataset or process inventory as a whole rather than for the individual quantity, it is attached with a marker saying so. A citation is therefore never presented as more specific than the evidence behind it. Coverage of citation binding across all published inventory rows is measured continuously (Section 3.4.4); residual limitations are declared in Section 6.3.9.

### 3.4.4 Coverage Accounting

Citation and uncertainty-band coverage are accounted per published inventory row: each row either carries the binding in question or is exempt from it for a recorded reason — for example, allocation co-products whose quantity restates an allocation convention rather than evidence, rows whose provenance is carried by a linked dataset, or exclusion statements with no quantity. Exemptions are enumerated as a documented policy with a recorded rationale for each class and are auditable per row, and the accounting is checked automatically as part of release validation so no row can silently fall out of it. The standard set for final publication is complete coverage separately for citations and for uncertainty bands: every published inventory row either carries at least one resolved citation or is exempt on record, and every row likewise either carries a published uncertainty interval or is exempt on record. Until that standard is met, residual uncovered rows are enumerated by named class with a declared plan for closing them (Sections 6.3.8 and 6.3.9).

### 3.4.5 Uncertainty Basis Guidelines

Default uncertainty ranges based on source quality:

**Table 3.4:** Default uncertainty ranges by data situation

| Situation | Default Range |
|---|---|
| Specification states "X ± Y%" | Use stated tolerance |
| Single source, verified | ±25% |
| Multiple sources agree | Min/max of sources |
| Multiple sources disagree | Full range, investigate |
| Engineering estimate | ±50% minimum |
| Scaling from other node | ±40% additional |
| Assumed value | ±100% or justify |

## 3.5 Data Gap Management

Data gaps (parameters for which no adequate public data could be found) are an inevitable reality when constructing LCI data from secondary sources. This section documents how gaps are identified, characterized, and addressed through workaround strategies.

### 3.5.1 Gap Identification Criteria

A data gap is formally documented when any of the following conditions apply:

**No public data available** Extensive literature search yields no quantitative data for a required inventory parameter.

**Temporal inadequacy** Available data is older than 10 years for a rapidly evolving technology, making extrapolation to current processes unreliable.

**Low reliability** Available data has significant quality concerns (e.g., values derived from outdated proxies, missing documentation of measurement conditions).

**Applicability uncertainty** Data exists but its applicability to the specific process or technology generation is questionable.

### 3.5.2 Gap Prioritization

Gaps are prioritized based on their potential impact on inventory accuracy:

**High priority** The parameter affects flows contributing more than 10% of total energy, water, or a key material category. Resolution is essential for model credibility.

**Medium priority** The parameter affects flows contributing 1–10% of totals. Resolution improves accuracy but models remain usable with documented uncertainty.

**Low priority** The parameter affects flows contributing less than 1% of totals, or relates to secondary parameters that do not directly enter inventory calculations.

### 3.5.3 Workaround Strategies

When gaps cannot be immediately resolved through additional research, the following strategies are employed:

**Table 3.5:** Data gap workaround strategies

| Strategy | Application |
|---|---|
| Proxy data | Use data from similar process/technology with adjustment |
| Stoichiometric | Calculate from reaction chemistry |
| First principles | Derive from physics (energy, mass balance) |
| Literature range | Use reported range from multiple sources |
| Expert estimate | Document assumptions and basis |
| Exclusion | Omit with documented justification (for minor flows) |

## 3.6 Validation Approach

### 3.6.1 Internal Consistency Checks

All published datasets pass automated internal consistency checks. Because the inventories are built bottom-up, flow by flow, an exact input-equals-output closure is not the acceptance criterion — product mass leaves the balance at the gate, and emissions are derived from declared fate parameters rather than by difference. The checks are instead structural: each verifies that related flows in a dataset are mutually consistent with the parameters that produced them.

- **Process-gas mass balance**: For every process gas, the same-substance emission to air must fall within the window implied by the gas's declared chamber utilization and abatement destruction efficiency (Section 4.1.6). This is enforced as a release gate across the full corpus; any deviation must either be corrected or recorded in an exceptions register with its physical reason (for example, a gas consumed stoichiometrically into a deposited film).
- **Water balance**: Wastewater discharged cannot exceed fresh water supplied; the difference is evaporative loss, which is itself modeled per facility class. The check is one-sided because evaporation legitimately reduces discharge, and it carries only a small tolerance to absorb display rounding, not a physical allowance.

- **Dimensional analysis**: Units must be consistent (zero tolerance).
- **Range checks**: Values must be within physically possible bounds, screened by automated sanity audits over the published files.

A dataset failing any of these checks is not accepted until the deviation is corrected or its cause is documented.

### 3.6.2 External Benchmarking

Calculated inventories are validated against published benchmarks:

**Table 3.6:** Benchmark sources by component type

| Component Type | Benchmark Sources |
|---|---|
| Semiconductor wafers | TSMC/Samsung sustainability reports, academic LCA |
| Memory | SK Hynix/Micron environmental reports |
| PCBs | IPC industry surveys, academic studies |
| Packaging | OSAT company reports, SEMI data |
| Passives | Academic papers, manufacturer datasheets |

### 3.6.3 Acceptance Criteria

Validation results are interpreted as:

**<30% deviation** Acceptable; document any known reasons
**30–100% deviation** Investigate; may be acceptable with justification
**>100% deviation** Likely error; must resolve before use

For parameters with high uncertainty (yield, utilization efficiency), wider acceptance ranges may apply with documentation.

## 3.7 Reporting Data Quality

### 3.7.1 Dataset-Level Summary

Each published dataset carries its data quality assessment as part of the dataset record itself: the composite DQI, the five pedigree dimension scores behind it, and the uncertainty intervals on the reported quantities. The composite DQI, the pedigree scores, and the resolved source citations for each flow can also be browsed directly in the dataset explorer on the project website (https://reellci.com), which presents each dataset's flow list and quality summary; the reported quantities and their uncertainty intervals are carried in the datasets themselves.

### 3.7.2 Flow-Level Documentation

LCI exports carry the following on each flow row:

- Uncertainty range (min/typical/max) where a band is published
- Source citations, labelled by binding depth

- Notes on limitations or caveats

The composite DQI and its five pedigree dimension scores are carried once per dataset, not on individual flow rows.

## 3.8 Database Quality Assessment

This section presents the data quality assessment results for the current database release, generated using the automated quality scoring system described in the preceding sections.

### 3.8.1 Dataset and Evidence Coverage

Table 3.7 summarizes database-wide data quality statistics across the 520 unit process datasets described in Chapter 1. Every dataset carries a composite DQI: 402 are scored directly from their own assessed source inventories, and the 118 aggregation datasets (complete ICs, system-level components, optical modules and rack assemblies) carry a composite rolled up from the independently-scored datasets they reference, weighted by each constituent's share of the modelled cradle-to-gate energy demand.

**Table 3.7:** Database-wide data quality statistics

| Metric | Value |
|---|---|
| Unit process datasets | 520 |
| Directly-scored datasets | 402 |
| Aggregation datasets (energy-share roll-up composite) | 118 |
| Mean composite DQI (all datasets) | 2.07 |
| Median composite DQI | 2.2 |
| Datasets rated Very good / Good / Fair | 195 / 323 / 2 |

**Note**: These figures reflect the configuration-weighted compositing of Section 3.1.3, under which configuration-layer files (facility assumptions, support-process rates, scenario parameters) are part of the scored population. That weighting is deliberately conservative: it grades the generic-technology character of configuration-layer assumptions into the composite instead of scoring only the best-documented process files, and it is the principal reason the database-wide mean sits near 2.0 rather than lower.

### 3.8.2 Quality by Component Category

The most actionable quality view is by component category, showing the **average composite DQI** across each category's datasets under the configuration-weighted scoring. Table 3.8 presents these results.

Quality labels match the label published on each dataset. They are assigned from the full-precision composite before display rounding, so a dataset displaying a composite of 2.0 may carry the *Very good* label; the thresholds are *Very good* $[1, 2)$, *Good* $[2, 3)$, *Fair* $[3, 4)$, *Poor* $[4, 5]$ (cf. Table 3.2).

**Table 3.8:** Data quality indicators by component category (1=best, 5=worst)

| Category | Datasets | Avg DQI | Quality |
|---|---|---|---|
| Connectors | 10 | 1.70 | Very good |
| PCBs | 114 | 1.74 | Very good |
| Upstream materials | 20 | 1.78 | Very good |
| Packaging | 46 | 1.82 | Very good |
| Thermal | 5 | 1.86 | Very good |
| Specialty semiconductors | 16 | 1.89 | Very good |
| Optoelectronics | 12 | 2.02 | Good |
| Substrates | 2 | 2.05 | Good |
| Cables | 10 | 2.12 | Good |
| Memory wafers | 39 | 2.16 | Good |
| Components (system-level) | 30 | 2.20 | Good |
| Infrastructure | 14 | 2.21 | Good |
| Logic ICs | 42 | 2.22 | Good |
| Memory components | 31 | 2.23 | Good |
| Semiconductor wafers | 48 | 2.24 | Good |
| Specialty wafers | 20 | 2.34 | Good |
| Passives | 57 | 2.53 | Good |
| Electromechanical | 4 | 2.57 | Good |

### 3.8.3 Interpretation of Quality Patterns

The quality assessment reveals systematic patterns that reflect both the underlying data landscape and the deliberate conservatism of the configuration-weighted compositing:

**Very good (DQI $<$ 2.0)** Categories whose datasets are dominated by directly-documented evidence — vendor equipment documentation, datasheet-derived masses with tight scope, and densely-cited process definitions — with comparatively little dependence on facility-level configuration assumptions. PCBs and packaging sit here because their per-operation inventories are carried largely by process-specific sources.

**Good (DQI 2.0–3.0)** Most of the semiconductor wafer families now sit in this band, and the reason is worth stating plainly: wafer datasets draw a substantial share of their quantities — facility support energy, abatement behaviour, operational scenarios — from configuration-layer assumptions that are generic-technology by nature, and the configuration-weighted composite grades that dependence honestly instead of scoring only the process files. The flagship advanced-logic wafers score 2.2–2.3 (Good) under this reading. Passives and electromechanical components sit at the lower end of the band for the more familiar reason: heavier reliance on engineering estimates and datasheet-derived masses. Two wound-passive datasets rate Fair (3.0) and are the current floor of the database.

The shift of the wafer families from the Very good band into Good relative to earlier internal assessments is a scoring-honesty change, not a data regression: the inventory quantities were untouched by the recalibration, and the movement reflects configuration-dominance and re-audited temporal grades being weighted into the composite. Composite scores should be read together with the declared uncertainty bands and the per-flow source citations rather than as a standalone grade: a low (good) composite DQI indicates well-sourced inventory values, not low

variability. Chapter 6 documents the known data gaps that persist within otherwise well-scored categories.

# Chapter 4

# Sector-Specific Modeling Approaches

This chapter details the modeling approaches for each component category, including technology-specific parameters, key assumptions, and calculation methods.

## 4.1 Semiconductor Fabrication (Front-End)

### 4.1.1 Technology Node Definitions

The database covers semiconductor nodes from 180nm to 3nm, classified into three technology generations (IEEE International Roadmap for Devices and Systems 2024):

**Table 4.1:** Technology node classification

| Class | Nodes | Architecture | Lithography |
|---|---|---|---|
| Advanced | 3nm, 5nm, 7nm | GAA/FinFET | EUV + DUV immersion |
| Mature | 10nm, 14nm, 22nm, 28nm | FinFET/Planar | DUV immersion |
| Legacy | 40nm–180nm | Planar | DUV dry/I-line |

Key differentiators by node class:

- **Transistor architecture**: Planar $\rightarrow$ FinFET $\rightarrow$ Gate-All-Around (GAA)
- **Minimum feature size**: Determines pattern complexity and mask count
- **Metal layers**: Advanced nodes require 12–15+ metal layers
- **EUV adoption**: Critical enabler for sub-7nm patterning

### 4.1.2 Process Flow Structure

Semiconductor fabrication follows a standard module structure comprising three main phases. Each phase contains multiple process steps that are modeled individually and aggregated to produce the total wafer inventory.

**FEOL (Front-End-Of-Line):** Transistor formation, including shallow trench isolation (STI), gate stack formation with high-k dielectrics and metal gates, source/drain engineering through epitaxy and ion implantation, and silicide contact formation. FEOL processes are characterized by high-temperature thermal steps and ion implantation, with significant contributions from oxidation, nitride deposition, and high-dose implants.

**MOL (Middle-Of-Line):** Local interconnect formation, including tungsten plug deposition for contacts and local interconnect layers that connect transistors to the first metal layer. MOL involves tungsten CVD, which consumes significant amounts of $WF_6$ precursor.

**BEOL (Back-End-Of-Line):** Global wiring using copper damascene metallization across 10–15+ metal layers, via formation between layers, low-k dielectric deposition and integration, and final passivation. BEOL dominates the total mask count and represents a substantial portion of total process steps and energy consumption.

The total process step count ranges from approximately 300 steps for legacy nodes (65nm and above) to over 1000 steps for advanced nodes (7nm and below), driven primarily by multi-patterning requirements and increased metal layer counts.

Figure 4.1 illustrates the overall flow from silicon wafer through the three fabrication phases to produce a processed wafer ready for dicing and packaging.

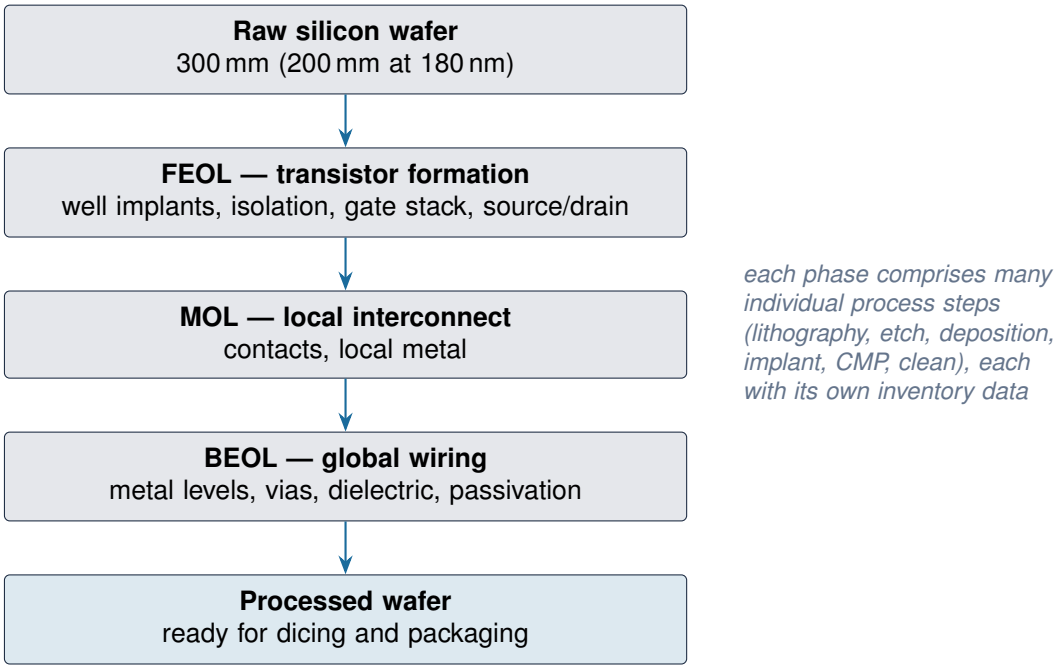


**Figure 4.1:** Semiconductor wafer process flow showing the three main fabrication phases: FEOL (transistor formation), MOL (local interconnect), and BEOL (global wiring). Each phase contains multiple process steps with individual inventory data.

### 4.1.3 Mask Layer Complexity

Mask layer count scales significantly with technology node:

**Table 4.2:** Unique reticle count by technology node

| Node | Total Masks | EUV Masks | Multi-Patterning |
|---|---|---|---|
| 3nm | 70 | 25 | SADP/SAQP |
| 5nm | 66 | 14 | SADP/SAQP |
| 7nm | 58 | 10 | LELE/SADP |
| 14nm | 62 | 0 | LELE |
| 28nm | 50 | 0 | Single |
| 65nm | 30–35 | 0 | Single |

Multi-patterning techniques:

- **LELE**: Litho-Etch-Litho-Etch (2 masks per layer)

- **SADP**: Self-Aligned Double Patterning
- **SAQP**: Self-Aligned Quadruple Patterning

**Cut-off note**: Photomask fabrication is excluded from the system boundary per the normative exclusion in Section 2.1.1 — at high production volumes the mask set amortizes to less than 0.1% of per-wafer inventory. Users assessing low-volume advanced ASIC production, where mask burden per die is significant, may need to add mask fabrication impacts separately.

### 4.1.4 Lithography Modeling

Lithography represents a significant portion of fab energy consumption:

**Table 4.3:** Lithography system energy profiles. Power is the tool-side load, excluding facility cooling, which the support model books separately. For deep-ultraviolet systems the energy column is power over throughput. For the extreme-ultraviolet system it is not: energy per wafer is the time-weighted average of the active, idle and standby states divided by the wafers actually exposed per hour, which is why 7.5 kWh/wafer is higher than the peak power divided by the nameplate throughput.

| System | Power (kW) | WPH | kWh/wafer | Nodes |
|---|---|---|---|---|
| EUV (NXE:3600) | 990 | 135–170 | 7.5 | 3nm, 5nm, 7nm (ASML Holding N.V. 2024) |
| DUV Immersion | 130 | 275–295 | 0.44 | All nodes |
| DUV Dry | 80–100 | 180–200 | 0.44–0.50 | 28nm+ |

EUV systems have multiple operating states that affect energy consumption. The active state (approximately 950 kW of tool-side load) applies during wafer exposure. Between wafers and during lot changes, systems drop to an idle state (approximately 450 kW, with vacuum maintained) or a sleep state (approximately 70 kW). For inventory modeling, energy is calculated from the time-weighted average of those states together with the wafers processed per hour.

### 4.1.5 Fab Modeling Assumptions

The virtual factory models incorporate several key assumptions about fabrication facility operations that affect how process-level inventories are aggregated and allocated.

#### 4.1.5.1 Production Capacity Basis

Per-wafer inventory values are independent of fab production capacity; they represent the resources consumed to produce one wafer through the defined process flow. However, certain facility-level allocations (such as cleanroom HVAC and bulk gas systems) are calculated based on assumed fab characteristics:

**Advanced logic (3–7nm)** Reference capacity of 50,000 wafer starts per month (WSPM), consistent with reported capacities for leading-edge facilities. This is the facility-class default: where a specific model declares its own capacity, that value governs the support-energy allocation — the 3 nm inventory, for example, assumes 40,000 WSPM to represent early production-ramp conditions.

**Mature logic (28nm+)** Reference capacity of 80,000 WSPM, reflecting higher throughput of established processes.

**Memory (DRAM/NAND)** Reference capacity of 145,000 wafer starts per month for dynamic memory and 120,000 for flash, reflecting high-volume memory production.

**Legacy nodes** Reference capacity of 80,000 WSPM; the 180nm node runs on 200mm wafers, the 90nm and 130nm nodes on 300mm.

**Capacity utilization** The models assume 100% capacity utilization, allocating all facility overhead to production output. This represents steady-state high-volume manufacturing conditions.

#### 4.1.5.2 Equipment Operating States

Process equipment cycles through multiple operating states, each with different power consumption:

**Processing** Full power consumption during active wafer processing. This is the primary state used for energy calculations.

**Idle** Reduced power when equipment is powered on but not processing wafers. Vacuum systems, temperature control, and monitoring systems remain active.

**Standby** Minimum power state while maintaining readiness for rapid return to processing.

The models calculate per-wafer energy using the formula:

$$E_{\text{wafer}} = \frac{P_{\text{effective}}}{\text{WPH}_{\text{effective}}} \tag{4.1}$$

Where $P_{\text{effective}}$ is the time-weighted average power (accounting for processing, idle, and standby states) and $\text{WPH}_{\text{effective}}$ is the achieved throughput under typical production conditions. This approach implicitly captures the operating state mix for a fully utilized fab operating at 100% capacity utilization.

#### 4.1.5.3 Cleanroom HVAC Allocation

In v1.0 the cleanroom-HVAC base is modeled as an ISO-6 ballroom with per-tool minienvironments rather than as a single area-weighted facility average. For silicon-CMOS logic and DRAM, the tool fleet that sets the minienvironment count is derived from each model's own process recipe (via Little's law — wafers in process as arrival rate times time in process, converted to tool positions as described in Chapter 2 — with the proportionality calibrated on the 5 nm logic reference case), so HVAC differentiates node-by-node; 3D NAND and specialty fabs retain hand-specified tool counts, and LED/optoelectronic fabs take an ISO-7 ballroom override. The ballroom contribution scales with the per-class cleanroom area anchored to publicly disclosed reference-fab footprints (Section 2.6.4.2), where ISO Class 1–3 zones require significantly more air changes per hour than ISO Class 5–7 zones. The earlier area-weighted base is retained as a documented fallback.

The HVAC support energy is additionally **climate-adjusted per region**: the climate-sensitive portion of make-up air conditioning carries a per-region multiplier derived from heating- and cooling-degree-day comparisons against the Taiwan baseline (spanning roughly 0.70 for temperate European sites to 1.10 for Arizona; Section 2.6.4.5). Every dataset's support en-

ergy therefore reflects the climate of its declared manufacturing region, and geography-variant datasets (for example the Arizona leading-edge logic datasets) re-apply the multiplier for their own region rather than inheriting the baseline's climate energy.

#### 4.1.5.4 Yield Factors for Wafer-to-Die Conversion

Converting wafer-level inventory to per-die inventory requires yield factors. The database applies the following default yield assumptions:

**Table 4.4:** Default yield assumptions by technology maturity

| Technology Maturity | Line Yield | Die Yield (100mm$^2$) | Notes |
|---|---|---|---|
| Mature production (>2 years) | 95% | 85–90% | High-volume, stable process |
| Volume production (1–2 years) | 90% | 75–85% | Ramping yields |
| Early production (<1 year) | 80–90% | 60–75% | Initial ramp, lower yields |

Die yield follows the Poisson model ($Y = e^{-D_0 \times A}$) where defect density $D_0$ and die area $A$ determine yield. For advanced nodes with higher defect densities or larger dies, yields decrease significantly. Users should substitute actual yield values when known, as yield is a linear multiplier on per-die inventory.

### 4.1.6 PFC Gas Handling

Perfluorinated compounds (PFCs) are used in etch and chamber cleaning with significant climate impact (U.S. Environmental Protection Agency 2006). For non-PFC process gases (CVD precursors, halide etchants) that undergo chemical transformation in abatement systems, see the emission proxy rules in Section 2.8. The PFC emission calculation follows a multi-stage model:

$$E = M_{\text{in}} \times (1 - U) \times (1 - R) \times (1 - \text{DRE}) \tag{4.2}$$

Where:

$M_{\text{in}}$ Mass of gas input (g/wafer)
$U$ Chamber utilization (fraction reacted in process)
$R$ Recovery rate (fraction captured for reuse)
**DRE** Destruction/removal efficiency of abatement

#### 4.1.6.1 Gas Recovery Scenarios

Three scenarios model gas recovery adoption:

**No Recovery** Legacy fabs (pre-2015); all unreacted gas to abatement
**Moderate Recovery** Industry average (2015–2020); recovery applies to every recoverable fluorinated gas as well as helium, hydrogen and neon, at rates from 25% to 60% – for example sulphur hexafluoride at 50% and helium at 60%.

**Table 4.5:** PFC gas parameters used in the inventory model ($GWP_{100}$ context values from IPCC AR6 (Intergovernmental Panel on Climate Change 2021); chamber utilization from the IPCC 2019 Guidelines; destruction-removal efficiencies from EPA reporting-rule and vendor evidence)

| Gas | $GWP_{100}$ | Utilization | DRE (Thermal) | DRE (Plasma) |
|---|---|---|---|---|
| $SF_6$ | 24,300 | 71% | 98% | 99% |
| $NF_3$ | 17,400 | 98% | 99% | 99% |
| $CF_4$ | 7,380 | 20% | 87% | 97% |
| $C_2F_6$ | 12,400 | 25% | 98% | 99% |
| $C_4F_8$ | 10,200 | 40% | 93% | 99% |

**Best Practice** Leading edge (2020+); $SF_6$ 90%, $CF_4$ 80%, He 95%

#### 4.1.6.2 Bulk Gas Supply

Beyond specialty gases, semiconductor fabs consume large quantities of bulk gases (nitrogen, argon, oxygen, hydrogen, compressed dry air) for purging, inerting, and process atmospheres. The energy for bulk gas supply is included in the support process allocation (see Chapter 2) and accounts for:

- On-site cryogenic air separation units (ASU) for $N_2$, $O_2$, Ar production
- Hydrogen generation (steam methane reforming or electrolysis)
- Compression and distribution systems

For advanced logic fabs, bulk gas supply typically contributes 15–25 kWh/wafer to total support energy.

### 4.1.7 Water Balance

Water flows in semiconductor manufacturing are tracked in two distinct categories that require different treatment in the inventory. Following ISO 14046 terminology:

**Water withdrawal** Total fresh water intake from the source (gross use)
**Water consumption** Water not returned to the original watershed, primarily evaporative losses
**Water discharge** Water returned to the hydrological system after treatment (wastewater)

**Fresh water withdrawal** includes ultrapure water (UPW) for wafer rinsing, cleaning, and wet process chemistry, as well as process cooling water makeup. The mass balance is:

$$W_{\text{discharge}} = W_{\text{withdrawal}} - W_{\text{consumption}} \tag{4.3}$$

Consumption (evaporative losses) occurs primarily in cooling towers (approximately 15–20% of cooling water flow), exhaust scrubbers, and wafer drying processes. An evaporation factor of 20% is applied to fresh water withdrawal to estimate consumption, with the remainder becoming wastewater discharge.

**Recirculated cooling water** flows through closed-loop systems for tool cooling, chilled water distribution, and heat rejection. This water does not become wastewater under normal operation; only makeup water to compensate for evaporative losses and blowdown enters the

fresh water balance. Recirculated cooling is tracked separately and excluded from wastewater calculations to avoid double-counting.

For process modeling, fresh water consumption is calculated from:

- UPW consumption per process step (from equipment specifications and rinse cycle counts)
- Scrubber water consumption (from abatement system specifications)
- Cooling tower makeup (from heat rejection requirements)

#### 4.1.7.1 Wastewater Contaminant Tracking

Beyond volumetric water balance, the database tracks **mass-based contaminant emissions to water** for process steps where dissolved pollutants are present in wastewater discharge. Only mass entries (g, mg, kg) are extracted; volumetric wastewater flows are tracked separately as technosphere outputs directed to treatment facilities.

Key contaminants tracked include:

**Fluoride ($F^-$)** Primary source is tungsten deposition (hydrolysis products in scrubber effluent); published values are around 0.7–1.1 g/wafer for advanced nodes, after calcium fluoride precipitation in wastewater treatment.

**Chloride ($Cl^-$)** From thermal oxidation (HCl-based clean steps) and metal etch chemistry.

**Copper ($Cu^{2+}$)** From CMP waste streams and electroplating rinse water.

**Suspended solids ($SiO_2$)** From CVD chamber cleaning (burn box particulate in scrubber water).

**Arsenic, phosphate** Trace quantities from ion implantation rinse water.

Contaminant masses are derived from mass-balance calculations where process chemistry permits, or from published treatment plant influent concentrations. These entries appear as elementary emissions to water in the unit process exports, with ecoinvent subcompartment assignments (typically “surface water”).

## 4.2 Specialty Semiconductor Wafers

The database includes models for semiconductor devices manufactured on non-silicon substrates or using specialized silicon processes that differ significantly from standard CMOS fabrication. These 12 specialty wafer models cover wide-bandgap power devices, III-V compound semiconductors, silicon photonics, image sensors, MEMS, and discrete devices.

### 4.2.1 Silicon Carbide (SiC)

SiC power devices (MOSFETs, Schottky diodes) are manufactured on 150mm SiC substrates with the following key process differences from silicon:

**Substrate cost** SiC wafers are 5–10$\times$ more expensive than silicon due to the physical vapor transport (PVT) crystal growth process operating at approximately 2200$^\circ$C.

**Epitaxy** CVD epitaxial growth of SiC drift layers at 1500–1650°C using silane and propane precursors.

**Ion implantation activation** Requires annealing at 1600–1700°C (vs. 1000°C for silicon).

**Oxidation** Thermal gate oxide growth is slower and less controllable than on silicon.

The modeled device class corresponds to the SiC power devices addressed by JEDEC's wide-bandgap qualification guidance, whose catalog of stress procedures (JEDEC Solid State Technology Association 2026b) standardizes reliability evaluation for this technology; such design-qualification testing is a development-phase activity outside the production gate-to-gate inventory boundary.

### 4.2.2 Gallium Nitride (GaN)

GaN power HEMTs are fabricated using MOCVD epitaxy on silicon substrates (GaN-on-Si), reducing substrate cost while leveraging existing 200mm silicon fab infrastructure. Key process characteristics include:

**MOCVD epitaxy** GaN/AlGaN heterostructure growth using trimethylgallium (TMGa) and ammonia at 1000–1100°C. MOCVD precursor production is covered by custom upstream datasets (Section 2.7.4).

**Fewer process steps** GaN HEMT fabrication requires significantly fewer mask layers and process steps than advanced silicon CMOS, resulting in lower per-wafer energy consumption.

**No ion implantation** Device regions are defined by epitaxial layer design rather than implantation doping.

As with SiC, the modeled GaN power-conversion device class is the one addressed by JEDEC's wide-bandgap qualification guidance—here the continuous-switching evaluation guideline (JEDEC Solid State Technology Association 2026c)—with the qualification testing itself outside the production inventory boundary.

### 4.2.3 III-V Compound Semiconductors

The database includes wafer models for III-V compound semiconductor devices used in optoelectronics and photonics:

**GaAs VCSEL** Vertical-cavity surface-emitting laser wafers for data communications and 3D sensing, fabricated by MOCVD on GaAs substrates.

**InP laser diodes** Edge-emitting laser wafers for fiber-optic telecommunications at 1310nm and 1550nm wavelengths, using InGaAsP/InP heterostructures.

**InGaAs photodiodes** Detector wafers for fiber-optic receivers, grown by MOCVD or MBE on InP substrates.

**GaN LEDs** LED wafers on 100mm sapphire or silicon substrates using MOCVD GaN/InGaN quantum well growth.

These models share the MOCVD precursor upstream datasets (TMGa, TMAl, TMIn) documented in Section 2.7.4. III-V fabrication is characterized by high substrate costs, lower wafer

throughput, and smaller wafer diameters (75–150mm) compared to silicon CMOS. Thermal characterization and modeling of compound-semiconductor FETs is standardized by JEDEC guideline JEP110A (JEDEC Solid State Technology Association 2025c), which delimits the III-V device classes these wafer models feed; as elsewhere, device characterization is outside the production inventory boundary.

### 4.2.4 Other Specialty Wafers

**CIS image sensors** CMOS image sensor wafers on 300mm silicon, combining standard CMOS transistor fabrication with specialized photodiode and color filter array processing.

**MEMS devices** Micro-electromechanical system wafers on 200mm silicon, including bulk and surface micromachining processes for accelerometers, gyroscopes, and pressure sensors.

**Discrete silicon** Standard silicon wafer processing for discrete transistors and diodes on 200mm substrates, using simplified process flows (fewer mask layers, no multi-patterning).

**Discrete power** Power device wafers (MOSFETs, IGBTs) on 200mm silicon with thick epitaxial layers and specialized termination structures.

## 4.3 Packaging and Assembly (Back-End)

### 4.3.1 Package Type Classification

Packaging technologies are classified by interconnect method:

**Table 4.6:** Package type classification

| Category | Types | Interconnect | Applications |
|---|---|---|---|
| Wire Bond | QFN, QFP, SOIC, BGA | Gold/copper wire | Consumer, automotive |
| Flip-Chip | FC-BGA, FC-CSP | Solder bumps | Processors, high-perf |
| Fan-Out | FOWLP, InFO | RDL + bumps | Mobile AP, RF |
| 2.5D/3D | CoWoS, HBM, EMIB | TSV + interposer | HPC, AI accelerators |

### 4.3.2 Die-to-Package Material Ratios

Material content varies significantly by package type:

**Table 4.7:** Typical material content by package type (relative to die mass)

| Package | Substrate | Mold | Interconnect | Total/Die |
|---|---|---|---|---|
| QFN (5×5mm) | Leadframe 3× | EMC 5× | Wire 0.1× | 8–10× |
| FC-BGA (35mm) | Organic 15× | EMC 8× | Bumps 0.5× | 20–25× |
| FOWLP | RDL 2× | EMC 3× | Bumps 0.3× | 5–7× |
| CoWoS | Si interposer 5× | Underfill 2× | $\mu$bumps 0.2× | 8–12× |

### 4.3.3 Wire Bonding Considerations

Wire bond material choice significantly affects inventory:

**Table 4.8:** Wire bonding material comparison

| Wire Type | Diameter | mg/bond | Notes |
|---|---|---|---|
| Gold (Au) | 25 $\mu$m | 0.092 | High reliability, high cost |
| Copper (Pd-coated) | 20 $\mu$m | 0.025 | Lower cost, harder process |

Mass per bond is a 3 mm wire at the stated diameter and the metal's own density, grossed threefold for the tail consumed at each bond.

Industry trend: Gold $\rightarrow$ copper transition reduces precious metal use by 60–80% per bond.

### 4.3.4 Advanced Packaging

#### 4.3.4.1 2.5D Integration (CoWoS, EMIB)

CoWoS packaging involves multiple process steps, each contributing to the total packaging inventory:

**Silicon interposer fabrication** TSV etching, liner deposition, copper fill, and redistribution layer (RDL) formation using processes similar to BEOL wafer fabrication.

**Microbump formation** Copper pillar plating with solder cap for fine-pitch interconnect.

**Die bonding** Thermocompression bonding of chiplets to the interposer with high placement accuracy requirements.

**Assembly completion** Interposer attachment to organic substrate, followed by underfill dispensing and optional overmold encapsulation.

The thinned-wafer, copper-TSV, and backside-RDL process scope modeled here corresponds to the technology platform standardized in JEDEC's reliability guidelines for 2.5D/3D silicon integration (JEDEC Solid State Technology Association 2025a), which delimit the same structures—thinned silicon with Cu TSVs and backside redistribution layers—that these packaging inventories cover.

#### 4.3.4.2 3D Stacking (HBM)

High Bandwidth Memory combines DRAM die fabrication with advanced 3D stacking:

**TSV formation** Through-silicon vias are etched and filled in DRAM dies prior to thinning, enabling vertical interconnect.

**Die thinning** Wafers are thinned to 50–100 $\mu$m to enable stacking while maintaining mechanical integrity.

**Die stacking** Multiple thinned DRAM dies (8, 12, or 16 high) are stacked using either thermocompression bonding with microbumps or direct hybrid bonding.

**Stack assembly** The completed DRAM stack is bonded to a base logic die and packaged.

Figure 4.2 shows a representative complete IC model, illustrating how wafer fabrication inventory is combined with packaging and test operations to produce the total component inventory.

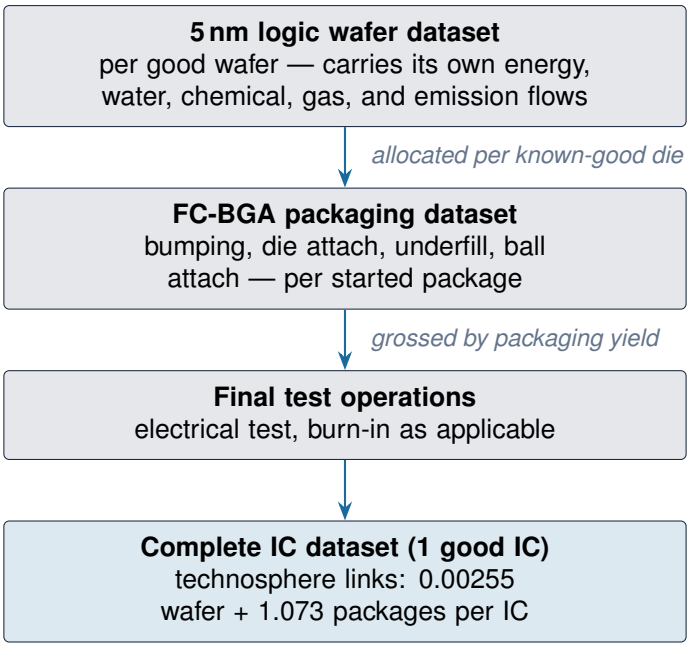


**Figure 4.2:** Aggregation structure for a complete integrated circuit (5nm logic in an FC-BGA package), with the live technosphere quantities of the published dataset. Each referenced dataset carries its own full input and output inventory (materials, energy, water, emissions, waste); this figure shows only how the datasets link, not the flows themselves — the wafer-fabrication system boundary with its complete set of inputs and outputs is given in Figure 2.1.

### 4.3.5 Packaging Facility Overhead

Just as wafer-fabrication support processes (cleanroom HVAC, ultrapure water, gas abatement) are allocated per wafer through the fab-type support model of Section 2.6.4, the back-end assembly-and-test (OSAT) datasets carry a parallel facility-overhead layer. It accounts for the site systems of an assembly-and-test plant—cleanroom air conditioning, chiller plant, compressed dry air, exhaust and dust collection, wastewater treatment, and uninterruptible power—which the direct tool-level process inventories do not themselves capture. In v1.0 this overhead is applied as an energy-only multiplier on the site-tagged direct process energy of each packaging and assembly dataset, $E_{\text{site}} = (1 + m)\, E_{\text{direct}}$ with $m = 0.80$ (band $[0.20, 1.10]$)—a total site multiplier of $\times$1.8 (band $\times$1.2–$\times$2.1). The multiplier is applied at the model level after flow aggregation and scales electricity only; material, water, and emission flows are unchanged.

**Site-tag boundary.** A process step carries the overhead when it is an assembly or test operation performed inside the OSAT house—per-package, per-die, per-bond, per-ball, per-wafer, per-panel, or per-strip. Purchased upstream materials referenced by a package model are *excluded*, because their own facility overhead is already embedded in the upstream inventory: leadframe manufacturing and plating, and the organic BGA-substrate core, are treated as upstream and receive no OSAT multiplier. In-house advanced-packaging substrate fabrication, by contrast, is site: the embedded-bridge (EMIB) bridge-formation, laminate, and drilling steps are performed within the packaging house and are tagged site, while the purchased substrate core they build upon remains upstream.

**Basis and magnitude.** The central $m = 0.80$ is triangulated from public assembly-and-test end-use disclosures—the KYEC 2018 and 2019 corporate sustainability reports (King Yuan

Electronics Co. 2019) (process equipment 44.8 %/45.9 % of site electricity; air-conditioning 35.1 %/33.5 %; compressed air and minor utilities the remainder) and the Chen and Hu survey of eleven Taiwanese assembly-and-test plants (Chen and Hu 2009) (facilities ≈53 % / tools ≈47 %), with a facility-controls-savings inversion as a lower-bound check—each evaluated on REEL's tool-side energy base and excluding lighting and office loads. Applied across the 46 packaging datasets, the layer adds a central ∼+80 % to site-tagged packaging energy; the HBM3 and HBM4 TSV-stacking datasets, for example, stand at ∼2.00 and ∼2.12 kWh per stack with the layer applied, versus ∼1.11 and ∼1.18 on direct process energy alone. Because packaging is a fraction of a fully assembled device's inventory, the effect on complete-device energy totals is far smaller—typically a few percent. The width and provenance of the multiplier, and the single-ratio-across-package-classes simplification, are discussed as residual limitations in Section 6.3.5.

### 4.3.6 Complete IC Product Coverage

The database provides complete IC datasets spanning the full breadth of semiconductor product types — 42 logic and interface ICs and 31 packaged-memory components in v1.0. Each model combines a wafer fabrication inventory (from the appropriate technology node) with a packaging inventory (from the appropriate package type), applying yield allocation at both the die and package level.

#### 4.3.6.1 Product Categories

**High-performance logic** Processors, GPUs, and NPUs at advanced nodes (3nm–14nm) in FC-BGA or CoWoS packages. These represent the highest-energy products due to large die sizes and complex packaging.

**Analog and mixed-signal** Op-amps, ADCs/DACs, voltage references, temperature sensors, and audio codecs at mature nodes (65nm–180nm) in small packages (SOIC, DFN, QFN). Lower per-die energy but very high production volumes.

**Power management** PMICs, LDO regulators, gate drivers, motor drivers, and LED drivers at 130nm–180nm in thermally-capable packages (QFN, SOIC). Energy dominated by packaging rather than wafer fabrication for small die sizes.

**Interface and networking** Ethernet PHYs and switches, CAN/RS-485 transceivers, USB hubs, and WiFi/Bluetooth SoCs across 28nm–180nm in QFN and FC-BGA packages.

**Programmable logic** FPGAs at 28nm in FC-BGA packages, representing mid-complexity devices between MCUs and custom ASICs.

**Memory ICs** Packaged DRAM ICs (including LPDDR in WLCSP), NAND flash ICs (176L, 232L and 288L in FBGA), and HBM stacks (8-, 12- and 16-high).

**RF and wireless** RF amplifiers and 5G modem SoCs, using both conventional packages and fan-out wafer-level packaging for improved RF performance.

This breadth of coverage enables bill-of-materials assessment for complete electronic products, where dozens of supporting ICs (voltage regulators, interface controllers, clock generators) may collectively contribute as much embodied energy as the primary processor.

## 4.4 Memory Components

Memory manufacturing shares many process steps with logic fabrication but has distinct characteristics that affect inventory modeling.

### 4.4.1 DRAM Wafer Fabrication

DRAM fabrication differs from logic in several important ways that affect both process complexity and per-wafer inventory:

**Capacitor formation** DRAM cells require high-aspect-ratio capacitor structures, either as deep trenches or tall stacks, using specialized deposition and etch processes.

**Technology node naming** DRAM nodes (1$\alpha$, 1$\beta$, 1$\gamma$) reflect process generations rather than minimum feature sizes, making direct comparison with logic nodes misleading. The letter convention used by some manufacturers (1a, 1b, 1c) and the Greek convention (1$\alpha$, 1$\beta$, 1$\gamma$) name the same generations; "1c" and "1$\gamma$" are the same sixth-generation node, not sequential ones.

**EUV adoption** The most advanced node modeled, the 1$\gamma$ (1c) generation, applies EUV lithography to 6 critical layers, with the two least-critical projected EUV applications reverted to DUV immersion to reflect production yield trade-offs. The 1$\gamma$ wafer is modeled at $787.8$ kWh/wafer.

**Yield characteristics** The highly regular array structure of DRAM enables higher die yields compared to logic at equivalent feature sizes.

**Process complexity** DRAM requires fewer mask layers and process steps than logic at comparable feature sizes, resulting in lower per-wafer energy consumption.

### 4.4.2 3D NAND

3D NAND flash memory achieves density scaling through vertical stacking rather than lateral feature scaling, fundamentally changing the relationship between technology generation and inventory:

**Vertical scaling** Modeled generations span 128 to 321 layers, with each generation adding layers to increase bit density. The 321-layer node is a three-deck (string-stacked) architecture with approximately 107 active wordline layers per deck, extending channel-hole depth and staircase-contact patterning relative to the 288-layer node.

**Process challenges** High-aspect-ratio etching (exceeding 50:1) and conformal deposition through deep features are critical process steps that become more demanding with increasing layer count.

**Energy scaling** Unlike lateral scaling in logic, vertical NAND scaling increases per-wafer energy consumption as layer counts grow. Each additional layer set adds to total deposition, etch, and inspection requirements. The 321-layer wafer is modeled at approximately $1{,}540$ kWh/wafer, the highest per-wafer energy of any memory wafer in the database, reflecting the deeper deposition and etch sequences of the three-deck stack. Point-of-use abatement electricity is modeled as a flat per-wafer fleet allocation rather than a per-pass

cost; the heavy fluorinated-gas load of the high-aspect-ratio etch sequence is therefore reflected in the wafer's residual F-gas emissions rather than in its abatement electricity.

**Die yield** 3D NAND die yield is modeled as a raw Poisson defect yield multiplied by a redundancy-recovery factor: production NAND tolerates array defects through bad-block and spare-block redundancy together with error-correcting codes, so a fraction of the dies a defect-only Poisson model would reject are recovered as good. For the 321-layer node this gives a die yield of approximately 0.81 ($\sim$673 good dies per wafer), raising good-die counts and correspondingly lowering the per-IC NAND inventory relative to a defect-only yield model.

**Lithography requirements** 3D NAND has less dependence on EUV lithography than advanced logic, as critical dimensions are not driven to the same extent by lateral pitch.

### 4.4.3 HBM Stacking

High Bandwidth Memory products combine DRAM wafer fabrication with advanced 3D packaging to create high-performance memory stacks. The inventory for HBM includes both the wafer fabrication component and the stacking and packaging processes:

**Die structure** HBM stacks consist of a base logic die (typically containing I/O and control circuitry) topped with 8, 12, or 16 DRAM dies.

**TSV integration** Dense through-silicon via arrays (exceeding 1000 TSVs per $mm^2$) provide vertical interconnection between stacked dies.

**Bonding technology** Either microbump bonding or direct hybrid bonding is used for die-to-die connections, with hybrid bonding offering finer pitch and better thermal performance.

The stack architectures modeled here fall within the scope of JEDEC's reliability-interaction guidelines for 3D chip stacks with through-silicon vias (JEDEC Solid State Technology Association 2026a), which—together with the 2.5D/3D thinned-wafer guidelines (JEDEC Solid State Technology Association 2025a)—standardize the TSV-and-thinned-die technology platform common to the HBM3, HBM3e, and HBM4 stacking inventories.

### 4.4.4 Packaged Memory Components

Beyond wafer-level memory models, the database provides complete packaged memory component models that combine wafer fabrication with packaging:

**DRAM ICs** Standard DRAM dies in wire-bond FBGA packages for DIMMs, and LPDDR DRAM in wafer-level chip-scale packages (WLCSP) for mobile applications. WLCSP packaging eliminates the substrate and mold compound, resulting in significantly lower packaging overhead per die.

**NAND flash ICs** 3D NAND dies (176-, 232-, and 288-layer) in wire-bond FBGA packages; the 321-layer node ships as a wafer dataset. Flash memory ICs typically use multi-die stacking within a single package (2–16 dies), with the stacking inventory modeled as repeated die-attach and wire-bond operations.

**Graphics memory ICs** GDDR7 graphics DRAM (16 Gb on a $1\beta$ die) in a 266-ball wire-bond FBGA package, modeled at approximately $0.86$ kWh per IC. This is the database's first

graphics-memory device, extending coverage beyond standard and mobile DRAM.

**HBM stacks** HBM3, HBM3e, and HBM4 products in 8-hi, 12-hi, and 16-hi configurations, combining the DRAM wafer, TSV stacking, and base die assembly inventories described above.

Two additional classes of packaged memory broaden the catalogue:

**High Bandwidth Flash (HBF)** A 16-high high-bandwidth-flash stack is provided as a **prospective** dataset, modeled at approximately $243$ kWh per component on the published architecture specification's device envelope (die area, stack composition, and a Poisson-model stacked yield). Because HBF is a pre-commercial technology (technology readiness level TRL 4–5, with first samples anticipated rather than shipping), this dataset is flagged as prospective and carries an asymmetric published uncertainty band (approximately $-24\%$ to +32%) reflecting the likelihood that production processes will be refined before volume manufacturing. It combines sixteen CBA NAND dies—each pairing a 3D NAND array wafer with its CMOS peripheral wafer by copper–copper hybrid bonding—with a base logic die, vertically stacked HBM-style via through-silicon vias and microbumps.

**System memory modules** Packaged and system-level memory now includes DDR5 form-factor variants—CUDIMM and CSODIMM (with a clock driver), MRDIMM (with registering clock driver and data buffers), and the CAMM2 compression-attached module—alongside SOCAMM and SOCAMM2 system-on-module assemblies for accelerator platforms and UFS 5.0 universal flash storage. Each is modeled as a component aggregate combining the constituent DRAM or NAND ICs, companion controller or buffer ASICs, and the module printed circuit board.

## 4.5 Optoelectronics and Silicon Photonics

The database includes models for silicon photonics, coherent optical transceivers, and the passive optical components used to interconnect them. These are modeled as inventory data—per-wafer, per-component, and per-kilometre energy, water, gas, and material flows—in the same manner as the semiconductor and packaging models.

### 4.5.1 Silicon Photonics PIC Wafer

The silicon photonics platform is modeled as a photonic integrated circuit (PIC) fabricated on a 200mm silicon-on-insulator (SOI) substrate at a CMOS-compatible 90nm node, representative of published foundry platforms. Rather than transistors, the patterned 220nm silicon device layer forms optical waveguides, carrier-depletion Mach–Zehnder modulators, germanium-on-silicon photodetectors, grating couplers, and resistive phase-tuning heaters.

**Mask complexity** The flow uses 19 mask layers—6 critical layers patterned with 193nm DUV dry lithography (three silicon etch depths, poly-silicon overlay, germanium window, contact via) and 13 non-critical layers patterned with 248nm DUV (eight implants for the modulator and germanium photodiode, one heater layer, two metal layers, one via, one passivation

opening).

**Germanium photodetector** Selective germanium epitaxy is grown in an opened oxide window over the silicon waveguide, followed by chemical-mechanical planarization and a cyclic thermal anneal to reduce threading-dislocation density.

**Metallization** Only 2 metal layers are required, since optical routing is performed in the waveguide layer rather than in metal.

The PIC wafer is modeled at approximately $402$ kWh/wafer (with an SOI substrate-production contribution and a germanium epitaxy contribution captured explicitly in the process flow), which is consistent with the expectation that a photonic process—with fewer mask and metal layers than advanced logic—sits below the per-wafer energy of comparable logic nodes.

### 4.5.2 Coherent Optical Transceivers

Optical transceiver modules are modeled as component aggregates that combine one or more semiconductor dies (DSP, silicon photonics PIC, and/or III-V laser and detector dies) with a printed circuit board, housing, fiber coupling, thermal management, and passive components.

**Direct-detect transceivers** The QSFP28 100G SR4 (short-reach multimode; four-lane NRZ with a retiming clock-and-data-recovery IC rather than a PAM4 DSP) and the QSFP-DD 400G family in DR4, FR4, and LR4 reach variants (PAM4 modulation with a 7nm DSP) use discrete optics for intra- and inter-datacenter links. Their optical sub-assemblies are modeled explicitly. Each dataset pulls its laser dies from the corresponding III–V epitaxial wafer inventory: GaAs VCSEL arrays for SR4, and InP DFB or EML lasers for the 400G reaches. Photodetector dies come from the InGaAs photodiode wafer; for the 850nm SR4 receivers this is a declared proxy, since the specification admits silicon or InGaAs detectors and the InGaAs wafer slightly overstates that variant's photodiode burden. Dedicated transmit- and receive-optical-sub-assembly (TOSA/ROSA) alignment-and-test steps complete the optical chain, and each die entry carries its own wafer-sort yield.

**Coherent transceiver** The QSFP-DD 400G ZR module integrates a silicon photonics PIC (replacing discrete transmit- and receive-optical sub-assemblies), an external InP continuous-wave laser, a 7nm coherent DSP, a thermoelectric cooler, and fiber coupling for 80 km reach with DP-16QAM modulation. The complete module is modeled at approximately $7.3$ kWh per unit, of which the coherent DSP die dominates and the silicon photonics PIC contributes a comparatively small share. The comparison across the family is instructive: the coherent ZR module ($\sim$7.3 kWh) carries a *lower* inventory than the discrete-optics FR4 and LR4 modules ($\sim$10.6 and $\sim$13.1 kWh); in the modeled architectures the difference is primarily associated with photonic integration replacing multiple individually-fabricated, yielded, and aligned III–V die channels with a single silicon wafer process.

Reach and format variants (SR4, DR4, FR4, LR4, ZR) are modeled as separate datasets because they differ in die content, laser count, fiber coupling, and test/burn-in duration.

### 4.5.3 Single-Mode Optical Fiber

Single-mode fiber (ITU-T G.652.D) is modeled per kilometre of coated fiber, covering preform fabrication by outside vapor deposition (OVD), high-speed fiber drawing, and dual-layer UV-cured acrylate coating. The model is distinct from the existing material-only cable models in that it carries a full manufacturing inventory.

The total is approximately $7.3$ kWh per good kilometre. Energy is allocated following the plant-level breakdown reported by Inakollu (2017): the draw area—furnace, in-line UV cure, cooling, and the associated production-floor infrastructure—accounts for 80–90% of plant electricity, and post-draw operations (rewinding, color coding, proof testing, packaging) for 10–20%. Preform fabrication is electrically minor because its soot-deposition step uses an oxy-hydrogen flame; that chemical energy is tracked as hydrogen and oxygen material inputs rather than as electricity. Yield is applied per manufacturing stage rather than as a single blanket factor: each stage's electricity is grossed up by its own and all downstream yield losses (preform survives the $\sim$5% preform-scrap, $\sim$5% draw-scrap, and $\sim$2% proof-test screen; the draw stage survives the draw-scrap and proof-test screen), so the per-good-kilometre inventory exceeds the raw as-drawn figure of $7.0$ kWh/km.

### 4.5.4 Fiber Connectors

Two fiber-optic connector models support the transceiver and fiber datasets:

**LC duplex** A single-fiber-pair connector with a zirconia ceramic ferrule and a glass-filled PBT housing, used on the FR4, LR4, and ZR transceiver variants.

**MPO-12** A multi-fiber push-on connector with a composite MT ferrule, used on parallel-fiber SR4/DR4 variants.

Both connectors are modeled from injection-molding, assembly, and test process files. The manufacturing of the ceramic (zirconia) ferrule itself—a precision-ground sintered component—is a documented data gap; the connector models currently account for the housing, assembly, and test inventory but not the ferrule's ceramic-forming energy.

## 4.6 Power Modules

The database includes a silicon carbide (SiC) power module model to support assessments of power-electronics assemblies for traction and renewable-energy inverters. The reference is a 600 A / 1200 V half-bridge module with six SiC MOSFET dies (three high-side and three low-side in parallel), aggregating publicly documented module designs.

The module is modeled as a component aggregate built from substrate, die-attach, interconnect, baseplate, and encapsulation process steps:

**Substrate** An active-metal-brazed (AMB) silicon-nitride ($Si_3N_4$) substrate with double-sided copper is the default, chosen for its high fracture toughness; a direct-bonded-copper (DBC) alumina or aluminium-nitride substrate is available as a variant for silicon IGBT designs.

**Die attach** Silver-sinter die attach is the default, reflecting its substantially higher thermal-cycling reliability than tin-silver solder; a legacy SAC305 solder path is available as a variant.

**Baseplate assembly** A copper baseplate is soldered to the substrate; aluminium-silicon-carbide, molybdenum, and copper–molybdenum baseplates are available as CTE-matched variants.

**Interconnect and encapsulation** Heavy aluminium wedge wire bonds carry top-side current, a silicone gel pots the wire bonds for vibration and dielectric protection, a thermal interface material is applied at the heatsink interface, and an injection-molded PBT housing completes the assembly.

The default-variant module is modeled at approximately $1.5$ kWh per module. The geography for this dataset defaults to Germany, reflecting the concentration of power-module manufacturing in Europe; this is an explicit per-model override, as the general regional manufacturing matrix does not yet cover power modules. Several inputs—the AMB substrate process energy, the silver-sinter press energy, and several specialty upstream materials (AMB brazing alloy, AlSiC, Cu–Mo, cure-in-place epoxy, phase-change material)—remain documented gaps targeted for refinement in a future release.

## 4.7 Printed Circuit Boards

### 4.7.1 Layer Count Scaling

PCB complexity scales with layer count:

**Table 4.9:** Modeled PCB energy and fresh-water withdrawal by layer count (ENIG surface-finish default; the other finish variants ship as separate datasets)

| PCB Type | Layers | Energy (kWh/m$^2$) | Water withdrawal (L/m$^2$) |
|---|---|---|---|
| Standard consumer | 4 | 56 | 1923 |
| Standard industrial | 6 | 61 | 2260 |
| HDI | 8–10 | 96–120 | 2421–2949 |
| Server motherboard | 12–16 | 96–141 | 3607–5709 |
| Networking/HPC | 18–20 | 185–201 | 7430–8126 |
| Backplane | 24–40 | 250–435 | 9896–18686 |

Values reflect the documented process steps—each step traces to an equipment specification or a documented per-operation inventory with explicit energy calculations—plus the facility support-process allocation appropriate to PCB facilities (HVAC, compressed air, water treatment; Section 2.6.4), so the totals above exceed the sum of the individual process steps by the support share. Water values are fresh-water withdrawal. Board manufacture takes no recycling deduction: a board shop runs none of the reclaim streams the deduction was derived for, so the figure published for each model is its full process and facility water demand.

Each PCB model contains a granular process flow of individually documented operations. Major process categories and their modeled per-operation energies:

**Electroplating** Copper plating for through-holes and surface layers. Modeled from rectifier current, plating time, and bath heating.

**Lamination** High-temperature pressing cycles. Each lamination press cycle modeled from equipment specifications.

**Imaging** Inner and outer layer photoresist exposure using LDI (laser direct imaging) systems at 0.071 kWh/m$^2$ per pass. Inner layers scale with layer count (e.g., 7 passes for 16-layer boards).

**Drilling** Through-hole drilling at 0.115 kWh per 1000 holes; blind/buried via drilling with estimated via densities (7000–65000 vias/m$^2$); back-drilling at 0.23 kWh per 1000 holes for stub removal.

**Desmear** $KMnO_4$ wet chemistry for smear removal at 2.5 kWh/m$^2$, modeled from conveyorized line power and throughput.

**Oxide treatment** Inner layer adhesion treatment (brown oxide or alternative oxide) at 2.0 kWh/m$^2$. Research indicates industry shifting to lower-energy alternative oxide (0.8–1.2 kWh/m$^2$).

**Surface finish** ENIG at 8.0 kWh/m$^2$; HASL, OSP, immersion silver, immersion tin, and ENEPIG are modeled separately and ship as finish variants of every board family (Section 4.7.2).

**Electrical test** Flying probe testing at 0.5 kWh/m$^2$.

**Silkscreen and routing** Legend printing plus CNC board routing at 0.35 kWh/m$^2$.

All process energy values derive from equipment power ratings divided by throughput, each documented with an explicit calculation basis and uncertainty range.

### 4.7.2 Surface Finish Variants

Surface finish significantly affects material and process inventory:

**Table 4.10:** Surface finish comparison

| Finish | Thickness | Market Share | Applications |
|---|---|---|---|
| HASL (lead-free) | 1–25 $\mu$m Sn | 20% | Low cost, general |
| ENIG | 3–6 $\mu$m Ni + 0.05 $\mu$m Au | 35% | Fine pitch, wire bond |
| OSP | 0.2–0.5 $\mu$m organic | 30% | Cost-sensitive |
| ImAg | 0.1–0.4 $\mu$m Ag | 10% | High frequency |
| ImSn | 0.8–1.2 $\mu$m Sn | 5% | Press-fit connectors |

The table lists the five volume finishes with published market-share estimates; a sixth finish, ENEPIG (electroless nickel, electroless palladium, immersion gold), is also modeled and shipped as a variant of every board family, and carries the highest process inventory of the set. Among the volume finishes, ENIG (electroless nickel immersion gold) has the highest impact due to:

- Nickel electroless plating (energy-intensive)
- Gold immersion (precious metal)
- Multiple rinse cycles

### 4.7.3 Panel Utilization

Manufacturing waste factors:

- **Panel edge trim**: 5–10% of panel area
- **Routing waste**: 3–8% (depends on board geometry)
- **Test coupons**: 2–5%
- **Yield loss**: 5–10% (layer-dependent)

Effective utilization: 70–85% of panel area becomes product.

### 4.7.4 Chemical Management

Key chemical processes:

- **Etching**: Cupric chloride or alkaline ammonia; 50–70% copper recovery
- **Electroless copper**: Formaldehyde-based reduction
- **Electroplating**: Acid copper sulfate, 85–95% current efficiency
- **Cleaning**: Alkaline cleaners, DI water rinses

### 4.7.5 Waste and Effluent Classification

Surface-finish and plating processes generate both liquid and solid waste streams whose correct classification matters for downstream impact assessment. The database applies the following conventions:

**Aqueous spent baths as wastewater** Dilute aqueous spent surface-finish baths (for example, spent electroless nickel, immersion silver, and immersion tin baths, which are 79–100% water) are classified as **wastewater** rather than solid waste. These streams are tracked as liquid waste with contaminant lists and treatment methods. Classifying them as wastewater eliminates double-counting against the treatment sludges (nickel hydroxide, copper sludge) that the same processes already report as solid waste.

**Plating-derived hydroxide sludges** Cu/Ni/Sn plating-derived metalliferous hydroxide sludges are routed to a cut-off market activity (`metalliferous_hydroxide_sludge`) rather than to a terminal underground-deposit treatment activity. Under cut-off allocation, the generator's burden ends at the market boundary, consistent with treating the sludge as a recyclable byproduct sold onward to metals recovery.

**Hazardous organic waste with energy recovery** Hazardous organic waste streams are mapped to the hazardous-waste incineration activity with energy recovery, reflecting that modern hazardous-waste incineration is predominantly waste-to-energy. Cut-off allocation means the recovered heat and electricity flow to the next user rather than crediting the waste generator.

### 4.7.6 Regional Effluent Limit Annotations

For packaged and assembled products manufactured through outsourced assembly and test (OSAT), v1.0 attaches **regional effluent discharge limits** for five OSAT-concentrated

jurisdictions—Taiwan, Malaysia, China, the Philippines, and Vietnam—as reference annotations in the exports. These limits are provided purely as documentation: REEL LCI tracks inventory as mass per functional unit, not effluent concentrations, so the limits are non-gating context that downstream users may apply to their own compliance checks; they are informational and do not alter the reported inventory flows.

### 4.7.7 Base Material Inputs

PCB models track the primary substrate materials consumed in board fabrication:

**FR-4 laminate** Glass-fiber reinforced epoxy resin, the standard rigid substrate material. Mass calculated from panel area, layer count, and laminate thickness per layer.

**Copper foil** Electrodeposited copper starting material for inner and outer layers. Typical thickness 17.5–35 $\mu$m (0.5–1.0 oz/ft$^2$).

**Prepreg** B-stage epoxy/glass sheets used as inter-layer bonding material during lamination.

**Polyimide film** Flexible substrate material for flex and rigid-flex PCBs, mapped to a custom polyimide proxy dataset.

**Coverlay** Polyimide-adhesive laminate providing solder mask functionality on flex circuits.

These materials are extracted into unit process exports as technosphere inputs with ecoinvent dataset linkage where available. Polyimide and coverlay reference the project's custom polyimide proxy dataset (see Section 2.7).

## 4.8 Passive Components

### 4.8.1 MLCC Modeling

Multilayer Ceramic Capacitors (MLCCs) are modeled by size code and dielectric type. The database provides individual models for common size codes (0201, 0402, 0603, 0805, 1206) in X7R dielectric, plus aggregate models for C0G and X5R dielectrics. Individual size-code models enable accurate bill-of-materials assessment where component sizes are known.

#### 4.8.1.1 Size Code Scaling

**Table 4.11:** MLCC size codes and typical parameters

| Size Code | L×W (mm) | Mass (mg) | Layers | Rel. Energy |
|---|---|---|---|---|
| 0201 | 0.6×0.3 | 0.35 | 100 | 0.19× |
| 0402 | 1.0×0.5 | 1.8 | 200 | 1.0× |
| 0603 | 1.6×0.8 | 7.0 | 250 | 3.9× |
| 0805 | 2.0×1.25 | 12.5 | 350 | 6.9× |
| 1206 | 3.2×1.6 | 40 | 500 | 22× |

Mass, layer count and relative energy are the values each size-code model declares; per-piece energy scales with component mass in this family.

#### 4.8.1.2 Dielectric Types

**C0G/NP0** Temperature stable (Class I); lower capacitance/volume

**X7R** General purpose (Class II); ±15% over temperature
**X5R** High capacitance (Class II); ±15%, narrower temp range
**Y5V** Maximum capacitance; +22%/-82% variation

#### 4.8.1.3 Sintering Process

Sintering is the dominant energy-consuming step in MLCC manufacturing, accounting for the majority of total production energy. The sintering process parameters vary by dielectric formulation:

**Temperature profile** Peak temperatures range from 1100 °C for Class II dielectrics (X7R, X5R) to 1350 °C for Class I (C0G/NP0), with controlled heating and cooling ramps.
**Duration** Total cycle times of 2–8 hours depending on part size and dielectric system, including ramp-up, soak, and controlled cooling phases.
**Atmosphere** Reducing atmosphere ($N_2/H_2$ mix) is required when using base metal (nickel) electrodes to prevent oxidation while achieving proper dielectric properties.

### 4.8.2 Chip Resistors

Chip resistors are fundamental passive components used in virtually all electronic circuits for current limiting, voltage division, and signal conditioning. Manufacturing processes differ significantly between thick-film and thin-film types, with corresponding differences in inventory profiles. As with MLCCs, the database provides individual size-code models (0201, 0402, 0805, 1206) to capture the significant mass and energy scaling across package sizes. Size-variant models are also available for chip inductors (0201, 0603, 0805) and ferrite beads (0402, 0805).

#### 4.8.2.1 Thick-Film Resistors

Thick-film chip resistors represent over 90% of global chip resistor production due to their low cost and adequate performance for most applications. The manufacturing process involves multiple steps:

**Substrate preparation** 96% alumina ($Al_2O_3$) ceramic substrates are scored, cleaned, and prepared for screen printing. Substrates are typically processed in large panels that are later singulated.
**Resistor printing** Ruthenium oxide ($RuO_2$) based resistive paste is screen-printed onto substrates. The paste composition determines the base resistance value, with different formulations for different resistance ranges.
**Firing** Printed substrates are fired at 850 °C peak temperature in a belt furnace, typically requiring 30–60 minutes total cycle time. This step sinters the resistive material and establishes the base resistance value.
**Laser trimming** Resistance values are adjusted to specification by laser ablation of the resistive film. This is a critical step that determines final tolerance and represents 10–15% of manufacturing energy.
**Termination** End terminations (silver, nickel, or tin) are applied by printing or plating, followed by barrier and solderable layers for board attachment.

**Protective coating** Glass or polymer overcoat protects the resistive element from environmental degradation.

Energy consumption for thick-film resistors is dominated by firing operations (approximately 60–70% of total) and laser trimming (10–15%).

#### 4.8.2.2 Thin-Film Resistors

Thin-film resistors serve precision applications requiring tight tolerances ($\pm$0.1% available) and low temperature coefficients. The manufacturing process uses vacuum deposition techniques borrowed from semiconductor fabrication:

**Substrate** High-purity alumina or silicon substrates provide stable dimensional references.

**Deposition** Resistive films (typically NiCr or TaN) are deposited by sputtering in vacuum chambers at 0.1–10 Pa pressure. Film thickness (50–500 nm) is precisely controlled to achieve target sheet resistance.

**Patterning** Photolithographic patterning and wet or dry etching define resistor geometry with high precision, enabling accurate resistance values without laser trimming for many applications.

**Passivation** Silicon nitride or other protective layers are deposited to prevent oxidation and ensure long-term stability.

Thin-film resistors require 3–5 times more energy per piece than thick-film equivalents due to vacuum processing requirements, but the higher precision reduces the need for binning and sorting operations.

### 4.8.3 Inductors

Chip inductors store energy in magnetic fields and are essential for power management, RF filtering, and signal processing applications. Manufacturing approaches differ based on required inductance values, current handling, and frequency characteristics.

#### 4.8.3.1 Multilayer Chip Inductors

Multilayer chip inductors are fabricated using ceramic processing techniques similar to MLCCs, enabling miniaturization and high-frequency performance:

**Core material** Ferrite ceramics are formulated from iron oxide combined with other metal oxides. NiZn ferrites provide high resistivity suitable for frequencies above 1 MHz, while MnZn ferrites offer higher permeability for lower-frequency power applications.

**Conductor formation** Silver or silver-palladium paste is screen-printed in a spiral pattern on unfired ferrite sheets. Multiple sheets are stacked to create 3D coil structures with the desired number of turns.

**Lamination and singulation** Stacked sheets are pressed together and cut into individual parts while still in the green (unfired) state.

**Co-firing** Parts are sintered at 900–1000 ℃ in controlled atmosphere furnaces. The co-firing process simultaneously densifies the ferrite body and forms the internal silver conductors. Cycle times of 4–8 hours are typical.

**Termination** External electrodes are applied by dipping or printing, followed by nickel barrier and tin solderable layers.

Energy consumption is dominated by the sintering process (65–75% of total), with ferrite powder preparation and termination plating contributing the remainder.

#### 4.8.3.2 Wire-Wound Inductors

Wire-wound inductors achieve higher inductance values and current ratings than multilayer types by using discrete wire windings around magnetic cores:

**Core fabrication** Ferrite cores are formed by pressing ferrite powder into drum or toroidal shapes, then sintering at 1000–1200 °C. Core geometry determines the magnetic path length and cross-sectional area, which together with turns count set the inductance.

**Winding** Copper magnet wire (typically 0.02–0.2 mm diameter) is wound onto cores using automated winding machines. Turn counts range from a few turns for high-current power inductors to hundreds of turns for high-inductance signal inductors.

**Termination** Wire ends are welded or soldered to termination pads. Surface-mount wire-wound inductors include molded or printed external electrodes for board attachment.

**Encapsulation** Many wire-wound inductors receive epoxy or molded encapsulation for mechanical protection and improved thermal performance.

Core sintering represents 50–60% of total energy consumption, with winding operations contributing 15–20% and termination/encapsulation accounting for the remainder. Wire-wound inductors typically require 2–3 times more energy per piece than equivalent-size multilayer inductors due to the discrete core firing step.

### 4.8.4 Other Capacitor Types

#### 4.8.4.1 Tantalum Capacitors

Tantalum capacitors use sintered tantalum powder anodes with anodized $Ta_2O_5$ dielectric and $MnO_2$ or conductive polymer cathodes. The database includes a tantalum capacitor model covering magnesium-thermic powder reduction, anode sintering, and cathode formation, with variant axes for cathode chemistry, charge density (CV/g), and case size. Manufacturing energy values are currently proxy estimates from analogous molten-salt systems; refinement against primary tantalum-powder metallurgy data is a priority for a future release (see Chapter 6).

#### 4.8.4.2 Aluminum Electrolytic

- Anode: Etched aluminum foil
- Dielectric: Anodized $Al_2O_3$
- Cathode: Electrolyte + aluminum foil
- Energy dominated by foil etching and forming

#### 4.8.4.3 Film Capacitors

- Dielectric: Polypropylene (PP) or polyester (PET)
- Electrodes: Metallized film or foil

- Lower energy than ceramic/tantalum
- Larger form factor

## 4.9 Complementary Rack Infrastructure

In addition to the primary focus on semiconductor components, PCBs, and passive devices, the database includes models for **complementary rack infrastructure** to support holistic assessments of IT equipment and data center deployments.

### 4.9.1 Scope and Purpose

These infrastructure models cover mechanical and structural components typically found in high-performance computing racks, including:

- Server rack frames (48RU open-frame structures)
- Compute tray assemblies and slide rails
- Power distribution components (busbars, PDUs)
- Liquid cooling manifolds and interconnects
- Board pans and cable management assemblies

The models are based primarily on published industry standards, including the **OCP Open Rack V3 Base Specification** and the **NVIDIA MGX Accelerated Computing Rack and Trays Specification**. These open standards provide dimensional requirements, material recommendations (e.g., SGC400 galvanized steel for load-bearing structures), and interface specifications that enable inventory estimation.

### 4.9.2 Data Quality Considerations

**Important**: These infrastructure models represent **approximate estimates** developed without access to primary manufacturing data. Users should be aware of the following limitations:

**No primary data** Unlike semiconductor and PCB models which draw on equipment specifications and academic literature, rack infrastructure models rely entirely on dimensional analysis from published specifications and engineering estimates.

**Mass estimation** Component masses are estimated from industry benchmarks (e.g., commercial 48U rack specifications from Eaton, RackSolutions, Cheval Group) rather than direct measurement.

**Material composition** Material breakdowns are inferred from specification recommendations and typical data center construction practices.

**Manufacturing processes** Process energy is estimated using generic sheet metal forming, welding, and coating process files rather than manufacturer-specific data.

### 4.9.3 Intended Applications

Rack infrastructure was **not the primary focus** of the REEL LCI Database but was included to:

- Enable complete bill-of-materials coverage for IT equipment assessments

- Support **hotspot analysis** where users need to understand the relative contribution of infrastructure versus active components
- Provide order-of-magnitude estimates for **data center modelers** assessing rack-level environmental profiles
- Facilitate screening-level comparisons of different rack configurations

For applications requiring higher confidence in infrastructure impacts, users should seek manufacturer-specific data or conduct primary data collection.

### 4.9.4 Model Structure

Infrastructure models follow the same three-tier architecture as other component models but with simplified process flows. Material inputs are documented with ecoinvent dataset references for upstream linkage:

- Steel sheet rolling and galvanizing (for structural frames)
- Aluminum extrusion (for heat sinks and thermal components)
- Sheet metal forming and welding operations
- Powder coating for surface finishing

Manufacturing waste (blanking scrap, coating overspray) is tracked and typically represents 15–20% of gross material input for sheet metal components, with high recyclability for steel scrap.

# Chapter 5

# User Guidance

This chapter provides guidance for practitioners applying the REEL LCI Database in their assessments.

## 5.1 Integration with Background Databases

### 5.1.1 Upstream Material Linkage

The database provides manufacturing inventories that require linkage to upstream processes for complete cradle-to-gate assessment. Users must connect the following elementary flows to background database equivalents:

#### 5.1.1.1 Raw Materials

**Table 5.1:** Key materials requiring upstream linkage

| Material Class | Examples | Background Process |
|---|---|---|
| Silicon | Wafer substrates | Silicon wafer production |
| Metals (bulk) | Copper, aluminum, nickel | Metal primary/secondary production |
| Precious metals | Gold, silver, palladium | Precious metal refining |
| Specialty metals | Tantalum, tungsten, cobalt | Specific metal supply chains |
| Ceramics | Alumina, barium titanate | Ceramic powder production |
| Polymers | Epoxy, polyimide, PET | Polymer resin production |
| Specialty chemicals | Photoresists, solvents | Chemical production |

#### 5.1.1.2 Energy Carriers

**Electricity** Connect to appropriate regional grid mix (see Section 5.4)

**Natural gas** Link to regional natural gas supply for thermal processes

**Nitrogen, oxygen, argon** Link to cryogenic air separation processes

**Hydrogen** Link to steam methane reforming or electrolysis

**Specialty gases** Many specialty gases ($NF_3$, PFCs) may require custom supply chain modeling

### 5.1.2 Electricity Grid Considerations

Electricity consumption typically dominates the environmental footprint of electronics manufacturing. The database provides electricity as a separate flow to enable user selection of appropriate grid mix:

**Note**: Carbon intensities are approximate values for illustration and vary annually. Users should

**Table 5.2:** Reference electricity carbon intensities by region

| Region | kg $CO_2$e/kWh | Production Share | Source |
|---|---|---|---|
| Taiwan | 0.50 | Advanced logic (TSMC) | Taiwan EPA (2023) |
| South Korea | 0.45 | Memory (Samsung, SK Hynix) | KEPCO (2023) |
| United States | 0.40 | Mixed (Intel, GF) | US EPA eGRID (2022) |
| European Union | 0.25 | Specialty, automotive | EEA (2023) |
| China | 0.55 | Growing capacity | IEA (2023) |
| Global average | 0.45 | Fallback for unknown supply chains | IEA World Energy Outlook |

obtain current emission factors from authoritative sources such as the IEA, national environmental agencies, or utility-specific data where available.

**Important**: Table 5.2 is provided for **screening purposes only** (quick carbon footprint estimates). For ISO-compliant LCA capturing multiple impact categories, users should link the electricity flow to a region-specific "Market for electricity" process in their background database (e.g., ecoinvent). This approach captures not only GWP but also acidification, particulate formation, and other impacts associated with the regional fuel mix.

**Recommendation**: Select the geography-matched dataset where one is published (Section 5.5), and link electricity to the regional mix of the actual supply chain where known. Fall back to the global average only for unknown supply chains.

## 5.2 Allocation and Cut-off Rules

The normative system boundary, cut-off criteria, and allocation rules are defined in Chapter 2 (Sections 2.1, 2.3, and 2.4). This section summarizes their practical implications for users applying the datasets.

### 5.2.1 What the Allocation Rules Mean in Use

**Facility support processes** Support burdens (cleanroom HVAC, ultrapure water, and related systems) are allocated per wafer within a facility class. For logic and DRAM wafers the cleanroom-HVAC share additionally differentiates by each product's own process sequence (the recipe-derived tool fleet of Section 2.6.4); the remaining support flows scale with throughput. Support shares are model-derived, not measured per production line.

**Multi-product fabs** Where a fab produces multiple technology nodes, allocation is by wafer throughput; economic allocation is not used. High-value advanced products carry the same facility share as commodity products from the same line.

**Yield allocation** All inventory burdens are assigned to good units; rejects are process waste with no burden of their own. Per-unit inventories therefore already include the resources consumed making rejected units — do not gross them up again.

**Complete ICs** Packaging and test inventories are grossed by packaging yield at the IC level; die inventory carries through from wafer fabrication with die-yield allocation. Users composing their own assemblies from wafer and packaging datasets must apply the same yield ownership (the consuming dataset owns the packaging yield).

### 5.2.2 What the Cut-offs Mean in Use

The declared exclusions (capital equipment and facility construction, photomask fabrication, administrative overheads, sub-threshold ancillary materials; Section 2.3 and Section 2.1) are calibrated to high-volume production. Users should consider whether they remain appropriate for the case at hand: for low-volume production or early-stage technologies, photomask amortization and capital intensity can be materially larger than the high-volume basis assumes.

## 5.3 LCIA Method Compatibility

### 5.3.1 Inventory-Only Scope

This database provides **inventory data only**, not characterized impact results. Users must select an appropriate Life Cycle Impact Assessment (LCIA) method:

**TRACI 2.1** US EPA method; appropriate for North American assessments
**EF 3.0/3.1** European Commission Environmental Footprint; required for EU PEF studies
**ReCiPe 2016** Global harmonized method; endpoint and midpoint indicators
**CML-IA** Widely used baseline method; midpoint indicators
**IPCC GWP** For greenhouse gas accounting; use AR6 values

The choice of LCIA method affects characterized results but not the underlying inventory data.

### 5.3.2 Elementary Flow Nomenclature

Elementary flows are mapped to ecoinvent v3.12 nomenclature for interoperability.

#### 5.3.2.1 Compartment Structure

Emissions use ecoinvent's compartment/subcompartment hierarchy with defaults appropriate for industrial semiconductor and electronics manufacturing:

**Table 5.3:** Default compartment assignments for electronics manufacturing

| Compartment | Default Subcompartment | Rationale |
|---|---|---|
| Air | non-urban air or from high stacks | Controlled point sources with abatement |
| Water | surface water | Treated wastewater to municipal/surface |
| Soil | industrial | Rare; only for direct soil emissions |

#### 5.3.2.2 Key Process Gas Emissions

Table 5.4 lists critical greenhouse gas emissions from semiconductor manufacturing with their ecoinvent mappings.

**Table 5.4:** Critical elementary flows with ecoinvent identifiers (GWP from IPCC AR6 (Intergovernmental Panel on Climate Change 2021))

| Flow | CAS | ecoinvent Name | $GWP_{100}$ | UUID (partial) |
|---|---|---|---|---|
| $CF_4$ | 75-73-0 | Tetrafluoromethane | 7,380 | b353e7d2... |
| $C_2F_6$ | 76-16-4 | Hexafluoroethane | 12,400 | 8f3e0579... |
| $SF_6$ | 2551-62-4 | Sulfur hexafluoride | 24,300 | c7c769bb... |
| $NF_3$ | 7783-54-2 | Nitrogen fluoride | 17,400 | f8cf5fd7... |
| $CHF_3$ | 75-46-7 | Trifluoromethane | 14,600 | 5e2e1740... |
| $N_2O$ | 10024-97-2 | Dinitrogen monoxide | 273 | afd6d670... |

#### 5.3.2.3 Fluorocarbon Emission Proxy

$C_4F_8$ (octafluorocyclobutane, GWP = 10,200) is proxied to $CF_4$ (tetrafluoromethane) for ecoinvent linkage, as it lacks a dedicated elementary flow entry. Exports carry the $CF_4$ UUID but retain the actual substance name, CAS number, and formula. $C_4F_6$ (hexafluoro-1,3-butadiene) has been cut off due to negligible GWP. See Section 2.8 for full details on emission transformation rules. Other process gases ($SiH_4$, $BCl_3$, $WF_6$, etc.) are either destroyed in abatement or transformed to mapped ecoinvent species.

#### 5.3.2.4 Material Input Mapping

Material inputs are mapped to ecoinvent v3.12 market datasets for upstream impact linkage, and each mapped input is published with the matched dataset's name and UUID, so the link to the background database can be made without manual matching. Where no direct ecoinvent match exists, the published record states how the link was made instead: a documented proxy (a similar dataset standing in, with its limitation stated), a composite modeled from constituent materials with explicit mass fractions, or a reference to another REEL dataset where the upstream inventory was built in this database. A small set of specialty chemicals absent from ecoinvent is mapped by name and identifier to CarbonMinds datasets; this database contains no CarbonMinds data, and users without CarbonMinds access should treat those flows as unresolved or fall back to the documented proxy where one is given. Materials excluded as negligible, and the residual data gaps, are flagged explicitly rather than dropped silently. Every technosphere input therefore either links to a background activity, references a REEL dataset, carries a documented proxy, is recorded as excluded below the cut-off threshold (Section 2.3), or carries an explicit unresolved flag. The status of every input is thus explicit; unresolved inputs remain for the practitioner to treat according to the goal and scope of their study, as do the proxy and cut-off entries. Elementary flows are handled the same way: each emission and resource flow maps to an ecoinvent elementary-flow identifier or is explicitly flagged where no equivalent exists (for example, the specialty fluorinated gases noted in Section 5.3.2.3).

**Dataset identity across versions.** Every REEL dataset carries a stable dataset UUID (and a product-level material UUID) drawn from a version-scoped registry. At a major release boundary the identifier space is deliberately reset: v1.0 identifiers are disjoint from v0.1 identifiers, so an integration that joins on UUID cannot silently mix inventories computed under different methodology versions—an upgrade is an explicit re-link, not an in-place mutation. Within a ver-

sion, identifiers are stable across regenerations, and each aggregate dataset's technosphere references carry the same UUID as the referenced component dataset (enforced by an automated referential-integrity check), so practitioners can traverse the supply chain graph by identifier alone.

**Key Proxy Assumptions** Table 5.5 documents significant proxy assumptions where ecoinvent lacks a direct equivalent. Users should be aware these proxies may underestimate upstream impacts.

**Table 5.5:** Material proxy assumptions and limitations

| Material | ecoinvent Proxy | Limitation |
|---|---|---|
| *Thermal Management* | | |
| Copper tube (OFC) | copper, cathode | Excludes tube drawing/extrusion ($\sim$0.5–1 kWh/kg) |
| Copper sheet (C1020) | copper, cathode | Excludes sheet rolling ($\sim$0.3–0.5 kWh/kg) |
| Cu-P brazing alloy | copper, cathode | Excludes phosphorus alloying (minor) |
| *Passive Components* | | |
| Barium titanate ($BaTiO_3$) | barium carbonate | Excludes $TiO_2$ input and calcination energy |
| Alumina ceramic (96%) | aluminium oxide, metallurgical | Industrial-grade proxy for electronic-grade |
| Glass frit | flat glass, uncoated | Excludes frit milling and formulation |
| *Connectors* | | |
| LCP (liquid crystal polymer) | PBT, glass fibre reinforced | Similar engineering thermoplastic |
| PA9T (high-temp polyamide) | nylon 6-6 | Lower-performance polyamide proxy |
| *Substrates* | | |
| BT resin | Upstream dataset created | See BT Substrate dataset |
| ABF (Ajinomoto Build-up Film) | Upstream dataset created | See ABF Substrate dataset |
| Solder mask (LPI) | epoxy resin, liquid | Excludes photoinitiators and fillers |
| Dry film resist | epoxy resin, liquid | Excludes photosensitive components |

**Composite Material Modeling** Several materials are modeled as composites with explicit mass fractions:

- **Mn-Zn ferrite**: 70% iron(III) oxide + 15% manganese ore (proxy) + 15% zinc oxide
- **Ni-Zn ferrite**: 70% iron(III) oxide + 15% nickel (class 1) + 15% zinc oxide
- **Enameled copper wire**: 97% copper cathode + 3% nylon 6-6 (enamel proxy)
- **Die attach epoxy (Ag-filled)**: 84.5% silver + 15.5% epoxy resin
- **SAC305 solder paste**: 82% tin + 2.5% silver + 0.4% copper + 11% flux

These compositions are derived from material safety data sheets and industry literature.

### 5.3.3 Impact Category Coverage

The inventory data supports assessment of the following impact categories:

- **Climate change**: Via electricity, fuel combustion, and direct PFC emissions
- **Acidification**: Via electricity-related $SO_x$ and $NO_x$
- **Eutrophication**: Via wastewater emissions
- **Resource depletion**: Via material inputs (metals, water)
- **Photochemical oxidation**: Via VOC emissions (solvents, photoresists)
- **Human toxicity**: Via heavy metals and specialty chemicals
- **Water use**: Via fresh water consumption

## 5.4 Geographic Adaptation

### 5.4.1 Baseline Geographic Scope and Regional Variants

Each dataset is published at a baseline representative manufacturing region; where supply chains meaningfully diverge, additional regional variant datasets are published alongside the baseline (Section 5.5):

**Table 5.6:** Baseline manufacturing regions and published regional variants by component type

| Component Type | Baseline Region | Published Variants |
|---|---|---|
| Leading-edge logic (5nm/3nm) | Taiwan | Arizona (US) |
| Other logic wafers | Taiwan | — |
| Memory (DRAM, NAND) | South Korea/Taiwan | — |
| Specialty wafers | Per technology (DE/US/JP/CN/TW) | US/JP for six technologies |
| Packaging and assembly | Taiwan | — |
| PCBs | China | — |
| Passive components | Asia-Pacific | — |

### 5.4.2 Regional Adaptation Procedure

To adapt data for specific supply chains:

1. **Identify actual supply chain geography** from supplier disclosures or industry knowledge
2. **Substitute electricity grid mix** with region-specific data
3. **Adjust water stress factors** if conducting water footprint analysis
4. **Consider regional efficiency variations** (documented in dataset metadata where applicable)

**Example**: For a product manufactured at TSMC Taiwan:

- Use Taiwan grid mix (0.50 kg $CO_2$e/kWh)

- Apply Taiwan water stress characterization factors
- No efficiency adjustment needed (datasets calibrated to leading-edge fabs)

#### 5.4.2.1 Climate Adjustment Considerations

Fab location affects facility energy consumption beyond grid carbon intensity. Climate conditions influence HVAC energy requirements for cleanroom temperature and humidity control:

**Hot and humid climates** Fabs in tropical or subtropical regions (Taiwan, Malaysia, Singapore) require additional cooling and dehumidification energy.

**Arid climates** Fabs in dry regions (Arizona, Israel) require humidification systems that consume significant water and energy, but cooling loads may be lower due to the dry-bulb temperature differential.

**Temperate climates** Fabs in temperate regions (Oregon, Ireland, Germany, parts of Japan) have lower HVAC energy requirements.

Climate is handled **explicitly**, not implicitly: every dataset's cleanroom-HVAC support energy carries a per-region climate multiplier for its declared baseline manufacturing region, derived from heating- and cooling-degree-day comparisons against Taiwan (spanning roughly 0.70 to 1.10 across the modeled regions, applied to the climate-modulated portion of make-up air conditioning; Section 2.6.4.5). Regional variant datasets re-apply the multiplier for their own region — the Taiwan and Arizona leading-edge logic datasets differ in both grid reference and climate-driven support energy. Users should select the geography-matched dataset where one exists, or refer to the regional multipliers in Section 2.6.4.5 as the governing values; older rule-of-thumb climate adjustments of $\pm$10–15% understate the spread of the explicit factors and should not be used.

## 5.5 Selecting Among Pre-Modeled Variant Datasets

A single physical product can be manufactured under several distinct configurations—different surface finishes, different operational practices, different fab geographies, or different process options. Rather than expose adjustable parameters, the REEL LCI Database ships each configuration as its own fully resolved, pre-modeled dataset with fixed assumptions. Every published dataset corresponds to exactly one product configuration with one set of assumptions, following the Explicit Dataset Principle (each export is one configuration, one assumption set).

**Important**: The published data is static. Users do not vary parameters at run time and do not re-run the model. Instead, users *select the variant dataset whose fixed assumptions match their scenario*. If no shipped variant matches a particular situation, that configuration is simply not represented in v1.0; users should pick the closest available dataset and document the residual difference, rather than attempting to adjust an existing dataset's numbers.

The v1.0 release materializes the following variant axes:

**PCB surface finish ($\times$6)** Each printed circuit board model is shipped in six surface-finish variants: HASL (hot-air solder leveling), ENIG (electroless nickel / immersion gold), ENEPIG

(electroless nickel / electroless palladium / immersion gold), ImAg (immersion silver), ImSn (immersion tin), and OSP (organic solderability preservative). The finish chemistry changes the plating-bath inputs, rinse water, spent-bath wastewater, and treatment-sludge outputs, so each finish is a separate dataset rather than a multiplier on a base board.

**Wafer operational scenarios (×3)** Semiconductor and memory wafer models are shipped in three operational-scenario variants—baseline, moderate, and best practice—that jointly set the process-gas recovery scenario and the water-recycling scenario. The baseline variant assumes no gas recovery and no water recycling; the moderate variant assumes industry-average gas recovery and moderate water recycling (the default for general assessments); the best-practice variant assumes leading-edge gas recovery together with comprehensive water reclaim (on the order of the highest publicly reported fab reclaim rates). (The abatement/destruction technology and its removal efficiencies are a property of the fab model, not of the scenario axis — scenarios vary recovery and recycling only.) These scenarios change direct fluorinated-gas emissions, fresh-water intake / wastewater discharge, and the *virgin* purchase quantities of recoverable process gases (recovered gas is netted out of the fresh hydrogen and neon inputs — that is what a recovery scenario means), but they do not change the underlying process recipe: the process flow, pass counts, and tool energy are invariant across the three scenarios.

**Fab geography (×2, leading-edge logic)** The 5 nm and 3 nm logic nodes are additionally shipped with a geography axis offering a Taiwan dataset and an Arizona dataset. The two datasets differ in the regional electricity market reference and in the climate-driven HVAC support energy (via the regional multiplier of Section 2.6.4.5); the process flow is otherwise identical. Users assessing supply from a specific region should select the matching geography dataset rather than re-grid an existing one.

**Cryogenic-etch option (×2, advanced 3D NAND)** The 232-layer, 288-layer, and 321-layer 3D NAND models are shipped with a cryogenic high-aspect-ratio etch axis, providing a conventional-etch dataset and a cryogenic-etch dataset. The cryogenic option reflects the published low-global-warming-potential gas-substitution chemistry and its associated etch-rate and energy profile; because adoption is opt-in per fab, both datasets are provided so users can select the one matching the facility being assessed.

**HBM TSV stack height (×3)** The HBM3 and HBM4 TSV-stacking packaging models are shipped in 8-high, 12-high, and 16-high stack variants. Stack height sets the number of DRAM die processed and bonded per stack, so per-stack energy and materials scale with height; users should select the variant matching the memory configuration of the device being assessed.

**HBM4 yield scenario (×3)** Each HBM4 stack ships in three core-die sort-yield scenarios reflecting production maturity: the baseline dataset (0.65 die yield, ramp-leaning central), a *ramp* dataset (0.62, early-ramp production), and a *mature* dataset (0.80, the industry mature-yield threshold). Die yield sets how much wafer fabrication is amortized into each good stack, so per-stack energy and materials scale inversely with it; users should select the scenario matching the production maturity of the supply under assessment.

**Tantalum capacitor SKUs (×24)** The tantalum capacitor model is shipped as 24 SKU datasets spanning three axes: cathode chemistry ($MnO_2$ or conductive polymer — different cathode-formation process files, materials, and failure-mode context), CV/g powder grade (5,000–50,000+ $\mu$FV/g, driving tantalum anode mass per capacitance), and EIA case size (A/B/C/D/R, driving total piece mass). Only the 24 vendor-catalog combinations are published; physically nonexistent combinations (e.g. maximum-CV powder in the largest case) are excluded by construction rather than interpolated.

Variant datasets carry explicit names that encode their fixed assumptions (for example, a 16-layer server board with an ENIG finish, or a 7 nm wafer under the moderate operational scenario), so the selected configuration is unambiguous from the dataset identifier. When the supply-chain configuration is unknown, the default datasets (moderate operational scenario; category-default geography per Table 5.6) are the recommended starting point.

## 5.6 Functional Unit Selection

### 5.6.1 Per-Component vs. Per-Function

The database provides data per physical component unit. For functional comparisons, users must define functional equivalence:

**Per component** Direct use of database reference flows (per die, per IC, per $m^2$)

**Per function** Requires user-defined equivalence (e.g., compute operations per second, storage capacity)

**Illustrative example** (round numbers, not dataset values): comparing processors by function:

- 7nm processor: 2000 kWh/wafer, 100 $mm^2$ die, 100 GFLOPS
- 14nm processor: 1200 kWh/wafer, 200 $mm^2$ die, 50 GFLOPS
- Per-die comparison favors 14nm (lower kWh/die)
- Per-GFLOPS comparison favors 7nm (higher efficiency)

### 5.6.2 Scaling to Product Level

To aggregate component inventories to product level:

1. List all components in bill of materials (BOM)
2. Match each component to appropriate LCI dataset
3. Multiply component inventory by quantity in product
4. Sum across all components
5. Add assembly and integration processes if within scope

$$I_{\text{product}} = \sum_i (I_{\text{component},i} \times Q_i) + I_{\text{assembly}} \tag{5.1}$$

## 5.7 Uncertainty Interpretation

### 5.7.1 Using Min/Typical/Max Ranges

The database provides three-point estimates (min/typical/max) wherever a range has been established for the value:

**Typical value** Use for point estimates and central tendency reporting
**Min/max bounds** Use as the variation limits when running sensitivity or scenario analyses
**Range width** Use to judge how well-constrained a value is — parameters that combine a wide relative spread with a large contribution are the ones worth investigating first

### 5.7.2 Sensitivity Analysis Recommendations

Recommended sensitivity analysis approach:

1. **Identify dominant flows**: Focus on parameters contributing $>$10% to total impact
2. **Vary key parameters**: Use min/max bounds for dominant flows
3. **Document sensitivity**: Report how results change across the range
4. **Prioritize data improvement**: High sensitivity + high uncertainty = priority for refinement

Key parameters typically requiring sensitivity analysis:

- Electricity consumption (typically 40–60% of carbon footprint)
- Die yield (strongly affects per-die allocation)
- PFC emissions (high GWP but uncertain abatement efficiency)
- Precious metal content (high upstream impact per gram)

### 5.7.3 Monte Carlo Sampling

For probabilistic assessment:

1. Assume triangular distribution with min/typical/max as a/mode/b
2. Generate $N$ samples (recommend $N \geq 1000$)
3. Propagate through calculation
4. Report median and confidence intervals (e.g., 5th/95th percentile)

## 5.8 Common Use Cases

### 5.8.1 Scope 3 Category 1 Accounting

For GHG Protocol Scope 3 Category 1 (World Resources Institute and World Business Council for Sustainable Development 2011) (purchased goods):

1. Obtain BOM with component quantities
2. Match components to LCI datasets
3. Calculate total energy consumption (kWh)
4. Apply regional electricity emission factor to obtain electricity-related GHG emissions

5. For direct PFC emissions: sum the mass flows of each PFC gas ($CF_4$, $C_2F_6$, etc.) from the inventory and apply characterization factors from the selected LCIA method (e.g., IPCC AR6 GWP values)
6. Sum for total Scope 3 Category 1 contribution

### 5.8.2 Product Carbon Footprint

For ISO 14067-compliant product carbon footprints (International Organization for Standardization 2018):

1. Define system boundary (cradle-to-gate for manufacturing)
2. Aggregate component inventories per product BOM
3. Link electricity to appropriate grid mix
4. Apply GWP characterization factors (recommend IPCC AR6)
5. Report manufacturing footprint with uncertainty

### 5.8.3 Comparative Assessment

For comparing alternative designs:

1. Ensure functional equivalence is defined
2. Use consistent methodology for all alternatives
3. Compare using same LCIA method and parameters
4. Report uncertainty ranges for both alternatives
5. Conclusions valid only when ranges do not overlap significantly

## 5.9 When Critical Review Is Required

Per ISO 14044 (International Organization for Standardization 2006b), external critical review is required for:

- Comparative assertions intended for public disclosure
- Studies supporting policy decisions
- Third-party certified environmental claims

Users intending such applications should arrange appropriate critical review for their specific study.

# Chapter 6

# Limitations and Future Work

This chapter documents known limitations and data gaps, provides a framework for future improvements, and closes with the independent critical review. Transparency about limitations is essential for appropriate use of the database.

## 6.1 Overall Data Quality Summary

The database-wide data quality assessment is presented in Tables 3.7 and 3.8 (Chapter 3). In summary, 195 datasets carry the **Very good** label, 323 **Good**, and two **Fair** (labels are assigned from the full-precision composite before display rounding; thresholds in Chapter 3), for a database-wide mean of 2.07 under the configuration-weighted compositing described there. Categories dominated by directly-documented evidence (PCBs, packaging, connectors, upstream materials) sit in the Very good band; the semiconductor wafer families sit in the Good band because their composites honestly weight the generic-technology configuration-layer assumptions (facility support, abatement, operational scenarios) that carry a substantial share of their quantities; passives and electromechanical components sit at the lower end of Good for their reliance on engineering estimates and datasheet-derived masses.

The data quality patterns reflect both the underlying data landscape and a deliberate scoring choice: a dataset is graded on everything that drives its quantities, including its configuration-layer dependence, not only on its best-documented process files. A sub-2.0 composite should be read as "well-sourced overall," not as uniformly primary-sourced (Chapter 3).

## 6.2 Known Data Gaps

Data gaps are systematically tracked and prioritized based on their impact on inventory accuracy. The following sections summarize key gaps by component category.

### 6.2.1 Semiconductor Manufacturing

The F-gas abatement-technology sensitivity is methodologically significant because destruction-removal efficiency (DRE) is strongly technology-dependent: wet scrubbers are largely ineffective for insoluble perfluorocarbons such as $CF_4$, whereas thermal and plasma-wet units achieve substantially higher DRE for the same species (U.S. Environmental Protection Agency 2006). The choice of abatement technology assumed for a given fab therefore propagates directly into the residual fluorinated-gas emission inventory for advanced nodes. Per-fluorocarbon residual-emission magnitudes are sensitive to this assumption and are being reconciled against the engine's abatement-technology resolution logic; pending that

**Table 6.1:** Key data gaps in semiconductor fabrication LCI datasets

| Gap | Priority | Impact |
|---|---|---|
| Wet-chemistry chemical specificity | Medium | Chemical inputs are populated for the major wet-process classes (CMP, plating, wet clean, developer), but per-recipe concentrations, bath lifetimes, and some file classes remain generic; limits downstream toxicity-category modeling |
| F-gas abatement-technology sensitivity | High | Abatement-technology choice (thermal vs. plasma) materially affects per-fluorocarbon destruction-removal efficiency, and hence residual fluorinated-gas emissions |
| Etch recipe specificity | High | Chemical consumption varies significantly by recipe |
| ALD precursor utilization | High | Low utilization (5–20%) means high precursor consumption |
| Emissions to water (expansion) | Medium | Mass-balance contaminant data populated for high-volume files only; broader coverage needed for eutrophication-category modeling |
| Captured-byproduct routing | Medium | The scrubber-captured share of transformed process byproducts ($\sim$0.7–1.0 kg per advanced-node wafer) is not yet routed to water or waste fates — a deliberate under-inclusion pending per-species fate assignments; the routing design is prepared for a future release |
| Wastewater volume tracking | Medium | Per-process wastewater volumes not yet systematically characterized |
| $NF_3$ chamber clean rates | Medium | Varies by tool and clean frequency |
| CMP slurry consumption | Medium | Slurry cost/impact varies by composition |
| Wet clean chemical volumes | Medium | Multiple rinse cycles affect total consumption |
| Photoresist utilization | Low | Spin coating efficiency relatively consistent |

reconciliation, advanced-node residual F-gas masses should be treated as carrying elevated uncertainty.

**Workaround approaches**:

- Generic etch recipes based on chemistry type (fluorine, chlorine, $SF_6$)
- Literature-based utilization factors with wide uncertainty ranges
- Stoichiometric estimation for thin film precursor consumption

### 6.2.2 Component Packaging and Assembly

Packaging and assembly gaps concentrate in proprietary material formulations and tool-specific energies; the process structure itself (die attach, interconnect, encapsulation, test) is well covered. Table 6.2 summarizes the known gaps.

**Table 6.2:** Key data gaps in packaging LCI datasets

| Gap | Priority | Impact |
|---|---|---|
| Advanced packaging variants | Medium | Hybrid-bonding (Cu–Cu) package assembly and next-generation bridge variants not yet modeled (CoWoS, EMIB, and Foveros are modeled) |
| Underfill composition | Medium | Proprietary formulations |
| Thermocompression bonding energy | Medium | Tool-specific energy consumption |
| Mold compound variations | Low | Relatively consistent epoxy-based |
| Marking and serialization | Low | Minor contribution |

### 6.2.3 Passive Components

Passive-component gaps stem mostly from proprietary material formulations — the dielectric, electrode, and termination chemistries that vendors treat as trade secrets. Table 6.3 summarizes the known gaps.

### 6.2.4 PCBs

PCB gaps concentrate in the wet-chemistry loop — etchant regeneration, bath lifetimes, and plating ancillaries — where operating practice varies widely between board shops. Table 6.4 summarizes the known gaps.

**Table 6.3:** Key data gaps in passive component LCI datasets

| Gap | Priority | Impact |
|---|---|---|
| MLCC dielectric formulations | High | $BaTiO_3$ dopant compositions proprietary |
| Facility-support allocation drivers | Medium | Support electricity and cooling-tower make-up water ship as separately cited rows (Chapter 2); compressed dry air and facility process-gas mass remain deferred pending an allocation driver |
| Tantalum capacitor anode energy | Medium | Anode chain integrated (see Section 6.5); residual gap is energy-value precision only — current values are proxy estimates from analogous molten-salt systems (no direct empirical measurement available). Refinement with production-scale data is a priority for a future release. |
| Sintering atmosphere details | Medium | $N_2/H_2$ ratio affects process |
| Electrode paste compositions | Medium | Nickel vs. precious metal |
| Termination plating details | Low | Standard Ni/Sn process |

**Table 6.4:** Key data gaps in PCB LCI datasets

| Gap | Priority | Impact |
|---|---|---|
| Etchant regeneration rates | High | Copper recovery and acid reuse rates |
| Laser microvia ancillary energy | Medium | Laser drilling is modeled; via fill plating energy after drilling is not yet captured |
| Via density per model type | Medium | Blind/buried via counts estimated from inline energy values; actual via counts would improve drilling calculations |
| Electroless copper bath life | Medium | Chemical replacement frequency |
| Regional efficiency variations | Medium | China vs. Taiwan vs. EU differences |

## 6.3 Methodological Limitations

### 6.3.1 Virtual Factory Approach

The virtual factory modeling approach has inherent limitations:

**No measured facility data** Models are constructed from equipment specifications and theoretical calculations, not actual facility measurements. Real fabs may operate differently than modeled.

**Idealized process assumptions** Process efficiencies represent typical or specification values, not actual operational variation. Equipment may run at different power states, utilizations, or efficiencies than modeled.

**Limited process-level validation** Published benchmarks are typically aggregated at the facility level (kWh/wafer, L water/wafer), making validation of individual process step estimates difficult.

**Recipe simplification** Actual process recipes are proprietary and highly variable. Models use generic recipes based on chemistry type rather than exact tool configurations.

### 6.3.2 Technology Currency

Electronics manufacturing evolves rapidly, creating technology currency challenges:

**Moore's Law progression** New technology nodes are introduced every 2–3 years. Models for advanced nodes (3nm, 5nm) have less validation data than mature nodes.

**Equipment generation differences** Equipment from different vendors and generations have varying efficiency. Models represent "typical" equipment rather than specific tool configurations.

**Process recipe evolution** Manufacturers continuously optimize recipes for yield and throughput. Published data may represent older process generations.

**EUV maturation** EUV lithography is rapidly evolving, with improving source power, throughput, and efficiency. Current models may not reflect latest improvements.

### 6.3.3 Geographic Representativeness

**Representative-region proxies** Default assumptions use a representative production region per category (e.g. Taiwan for logic wafers, Korea for memory) and, for specialty wafers since the July 2026 refinement, per technology (e.g. Germany for the mixed EU/US SiC fleet, United States for III–V photonics, Japan for CMOS image sensors); these are declared single-region proxies for mixed fleets, not production-weighted averages, and may not match a specific supply chain. Selected technologies additionally ship secondary regional datasets (US/JP options).

**Grid mix assumptions** Electricity carbon intensity varies significantly by region and changes over time as grids decarbonize.

**Facility-level variation** Even within a region, facilities vary in age, efficiency, and environmental management practices.

### 6.3.4 Allocation Limitations

The allocation approaches used have inherent limitations that users should consider:

**Facility overhead allocation** Support processes (HVAC, ultrapure water, waste treatment) are allocated per wafer through the facility-class methodology defined in Section 2.6.4 (fifteen classes with reference-fab anchors, recipe-driven cleanroom-HVAC base for logic and DRAM, area-weighted fallback). Limitations of both bases are detailed in Section 6.3.6.

**Multi-node facility allocation** When a fab produces multiple technology nodes, support is allocated by wafer count. Economic allocation (by revenue contribution) would yield different results for high-value advanced nodes.

**Yield allocation timing** All upstream burdens are allocated to good dies at the end of fabrication. This approach does not distinguish between early-stage failures (which could theoretically carry lower burden) and late-stage failures.

**Scrap recovery credits** Precious metal recovery from scrap (gold, copper) is excluded pending reliable recovery rate data. Including recovery credits would reduce net material inputs.

### 6.3.5 Packaging Facility Overhead

In v1.0 the back-end (OSAT) datasets carry a facility-overhead layer that applies a central multiplier of $\times$1.8 (an $m = 0.80$ adder) to the site-tagged direct process energy of each packaging and assembly dataset, paralleling the fab-type support model used for wafer fabrication; the layer, its site-tag boundary, and its evidence basis are described in Section 4.3.5. The overhead is therefore present in the reported packaging inventories—the residual limitation is the *width and provenance* of the multiplier rather than its absence.

**Multiplier uncertainty band** The overhead fraction carries an explicit band of $m \in [0.20,\ 1.10]$, i.e. a total site multiplier of $\times$1.2 to $\times$2.1. The public in-scope evidence in fact clusters near $m \approx 0.9$–$1.0$ (total $\times$1.9–$\times$2.0, with a raw un-discounted complement approaching $\times$2.13–$\times$2.23); the 0.8 central is set deliberately low, discounting once for post-2019 facility-efficiency programmes. The residual revision pressure is consequently upward rather than downward.

**Assembly-and-test end-use data sparsity** The multiplier rests on a small number of single-region (Taiwanese) assembly-and-test disclosures spanning 2009–2019 (Chen and Hu (Chen and Hu 2009); KYEC 2018/2019 CSR (King Yuan Electronics Co. 2019)). KYEC is a test-heavy operator, whose air-conditioning- and chiller-intensive test floors make its measured facility share a high-side proxy for assembly-heavy sites. Primary, recent, assembly-specific end-use data would tighten the band.

**Single ratio across package classes** A single $m$ is applied to every package class in v1.0. Because the dominant air-conditioning and chiller load co-scales with tool heat rejection, this is a defensible first-order treatment; but leadframe/wire-bond (ISO-7/8), laminate/flip-chip (ISO-6/7), and wafer- or panel-level (ISO-5/6) back-end lines differ in cleanroom class and air-change intensity. Per-class area- and recipe-based differentiation is a defined

future refinement.

### 6.3.6 Cleanroom HVAC Allocation Methodology

The cleanroom-HVAC support-energy methodology is defined in Chapter 2 (Section 2.6.4: recipe-driven ballroom-plus-minienvironment base, facility classes and reference-fab anchors, utility conversion basis, and regional climate adjustment). This section states its acknowledged limitations:

**Recipe-fleet single-node calibration** The tool fleet that sets the minienvironment count is estimated from each model's process sequence through Little's law, but the proportionality constant is calibrated on a single node (the 5 nm logic reference case) and propagated to the other logic and DRAM models. The per-tool minienvironment power intensity and tool-occupancy assumptions are documented engineering estimates rather than measured fab values; multi-node calibration and primary minienvironment-power data are refinement targets for a future release.

**Reference-fab anchor coverage** Only seven of the fifteen facility classes rest on publicly disclosed cleanroom-area anchors (Section 2.6.4.2); the remainder carry derived anchors with $\pm$30–50 % uncertainty bands, to be refined as further foundry disclosures become available.

**Single-anchor problem** Most disclosed facility classes rest on a single foundry's data. Triangulation across multiple fabs per class is methodologically preferred but blocked by the universal masking of per-fab cleanroom area in foundry sustainability disclosures (corporate aggregates only). Complementary anchors (e.g. Samsung Pyeongtaek for memory, Intel Fab 52 for advanced logic) are future candidates.

**ISO-class area distribution unvalidated** The assumed distribution of cleanroom floor area across ISO cleanliness classes is a load-bearing input to the area-weighted base. The fractions originate from a synthesis of TSMC and Micron facility disclosures and cleanroom-design literature, and have not been validated against primary measured fab data. Two dedicated literature searches for primary or facility-disclosed zone-area fractions found none publicly available: foundries treat per-zone area distribution as proprietary, and the qualifying public studies are meta-models (e.g. Lian et al. 2024) or CFD model rooms rather than surveyed distributions. The distribution is therefore retained as an explicit documented assumption.

**SEMI S23 small-fab baseline** The SEMI S23-1021 utility energy-conversion-factor table underlying the utility conversion basis (Section 2.6.4.3) is built on a 4,400 $m^2$ / 10k wafer-starts-per-month 200 mm fab (per Hu et al. 2019). Direct application of these factors to a modern 300 mm advanced-logic fab is acknowledged in that work as a known limitation; the independently anchored ultrapure-water and compressed-dry-air values partially mitigate this, but the SEMI-adopted flows inherit the small-fab basis.

**Ultrapure-water attribution overlap** The per-pass ultrapure-water bases may partially overlap tool-specific water draws (Section 2.6.4.3); an earlier construction had misallocated the fab-wide total to wet cleaning alone, and the bases were recalibrated against process anchors as an interim measure. A facility-total allocation refactor remains planned

(Section 6.6).

**Empirical foundry validation** Foundry sustainability disclosures reliably publish cleanroom area but universally mask per-fab annual electricity (a review of thirteen manufacturers' disclosures found corporate aggregates only). The single fab-level kWh/wafer figure for which both numerator and denominator are public is GlobalFoundries Fab 8 Malta (∼1,460 kWh/wafer from CDP-disclosed energy and the disclosed wafer-start capacity) — and even this is mixed-product. The bottom-up methodology consequently relies on equipment specifications, academic studies (Hu 2007/2008/2019, Sun 2017, Yin 2020, Lian 2024), and process-step composition rather than top-down calibration to fab-aggregate anchors.

### 6.3.7 Material Proxy Limitations

Several common materials in electronics manufacturing lack exact matches in ecoinvent v3.12, requiring proxy assumptions:

**Copper products** Oxygen-free copper tubes and sheets are proxied to "copper, cathode," which excludes tube drawing/extrusion energy (approximately 0.5–1 kWh/kg additional). This underestimates impacts for thermal components (heat pipes, vapor chambers).

**Ferrite ceramics** Mn-Zn and Ni-Zn ferrites are modeled as composite materials using iron oxide, manganese oxide, and zinc oxide. This captures raw material impacts but may not fully represent the high-temperature sintering process.

**Engineering polymers** Liquid crystal polymer (LCP) is proxied to glass-fiber reinforced PBT as the closest engineering thermoplastic. Processing energy differences are not captured.

**Specialty alloys** Phosphor bronze contacts are proxied to generic bronze; C7025 copper-nickel-silicon alloy is proxied to copper cathode. Alloying energy and minor element contributions are excluded.

The complete list of proxy assumptions with their documented limitations is provided in Table 5.5 (Section 5.3.2.4).

### 6.3.8 Uncertainty Band Coverage

The interval propagation described in Section 3.3 carries three declared limitations:

**Partial row-level coverage** Not every published inventory row carries an explicit uncertainty interval; the per-row coverage accounting of Section 3.4.4 tracks this continuously. The uncovered remainder consists mainly of derived output rows awaiting band propagation from their parent inputs, plus a smaller set of quantified rows for which no honest range has yet been located — each enumerated with a declared research path rather than filled with a synthetic width. Where headline fields carry synthetic category-default bands, each records that provenance so it is never conflated with a propagated interval (Section 3.3.3).

**No input-to-output band inheritance** Emission and waste output flows carry uncertainty bands only where a range was explicitly documented for the output itself. Where an output is stoichiometrically driven by a banded input (for example, a treatment byproduct

of a banded chemical input), the input's band is not automatically transferred to the output, so some published output-flow intervals may be narrower than their driving inputs justify. Automatic inheritance is a planned refinement.

**Reconciliation of malformed documented ranges** If a documented range places its central value outside its own min–max bounds (typically a stale range left behind after a central-value update), the offending bound is clipped to the central value rather than rejected, producing a one-sided band. Every published band is checked to bracket its central value at emission time, and a database-wide scan at the v1.0 release lock found no one-sided or zero-width propagated bands published. The residual limitation is that a stale documented range degrades quietly to a narrower band instead of failing loudly; stricter validation at documentation time is planned.

### 6.3.9 Citation Binding Depth

Source citations attach to individual flows at three labelled depths (Section 3.4.3): established for the specific value, inherited from the contributing process inventories, or documented at the dataset or process-inventory level. Not every published inventory row yet carries a resolved citation; the per-row coverage accounting of Section 3.4.4 tracks this continuously.

Two honest qualifications apply. First, the three binding depths are deliberately kept distinct on the published record: an inherited or dataset-level citation is labelled as such, and users should read the label — a row whose citations are labelled as dataset-level is documented at that level, not individually evidenced. Second, the uncovered remainder is enumerated by named class — principally process inventories whose source references are not yet resolved to full bibliographic entries, and waste rows whose contributing inventories are themselves not yet cited — each with a declared plan for closing it. Users should not interpret a flow whose own citation list is empty as evidence that the value is unsourced.

### 6.3.10 System Boundary Limitations

The boundary exclusions are declared normatively in Section 2.1; their principal implications are:

**Capital goods** Semiconductor fabrication is exceptionally capital-intensive, so the excluded equipment embodied energy may be a material share of full life cycle impacts over equipment lifetime — larger than the capital-goods share typical of other manufacturing sectors. Studies requiring capital goods must add them from other sources.

**Facility construction** Cleanroom construction and infrastructure are one-time investments amortized over facility lifetime; their exclusion matters most for new-build-heavy scenarios (greenfield capacity expansion).

**Wafer substrate production** Because silicon ingot growth and wafer slicing enter as background-database links rather than REEL processes, their representativeness is bounded by the linked background dataset, not by this database's methodology.

**Cradle-to-gate cut** Use phase and end-of-life are out of scope; for electronics with long, energy-intensive service lives, the gate-boundary inventory is typically a minority of the full life cycle result.

## 6.4 Appropriate Use and Cautions

### 6.4.1 Suitable Applications

This database is appropriate for:

- **Comparative assessments**: Comparing technology A vs. technology B using consistent methodology
- **Hotspot identification**: Identifying the most impactful components in a bill of materials
- **Screening-level LCA**: Initial assessments to guide detailed studies
- **Scenario analysis**: Understanding sensitivity to key parameters
- **Order-of-magnitude estimates**: Establishing reasonable bounds for electronics manufacturing impacts

### 6.4.2 Applications Requiring Caution

The following applications require additional caution:

- **Absolute environmental declarations**: Results should include uncertainty ranges and clearly state methodology limitations
- **Regulatory compliance calculations**: May require more specific or validated data
- **High-stakes procurement decisions**: Consider obtaining supplier-specific data where possible
- **Public comparative assertions**: Per ISO 14044, require external critical review

## 6.5 Resolved Since v0.1

Version 1.0 closed the major gaps declared at the initial release. In summary: emissions to water are now extracted and exported, with chemistry-based transformation rules and abatement destruction efficiencies applied to reactive process gases (a GWP-neutral correction that completes the acidification and toxicity inventory); every elementary and technosphere flow is accounted for against the background-database mapping, with explicit flags where no equivalent exists; waste streams carry refined classifications and treatment routings, including a dedicated recovery route for metal-bearing plating sludges; every exported flow carries a source-citation field resolved from curated source registries; cleanroom support energy moved to the recipe-driven ballroom-plus-minienvironment base, with abatement electricity modeled as a continuous fleet baseload and its fuel combustion surfaced as a background-database link; 3D-NAND die yield gained a redundancy-aware recovery factor; the background mapping migrated to ecoinvent 3.12; the tantalum capacitor anode chain was integrated (its energy values remain proxy estimates, Section 6.2.3); and coverage expanded across current-generation memory, silicon photonics and optical transceivers, modified semi-additive PCB processing, and the JEDEC-defined memory-module landscape. The version history at the front of this report records the consolidated change summary.

## 6.6 Future Improvements

### 6.6.1 Priority Data Needs

The known data gaps of Section 6.2 are the standing first-order priorities and are not restated here. Beyond closing those, the following improvements are planned for future versions:

1. **ALD precursor data**: Utilization factors and consumption rates for additional ALD processes beyond those currently modeled (TMA-$O_3$, HCDS, $SnCl_4$)
2. **Advanced packaging**: Hybrid bonding (Cu-Cu, oxide-oxide), EMIB v2, advanced CoWoS-S/L variants
3. **3D NAND beyond 321L**: Scaling parameters for 400+ layer architectures and additional string-stacking decks (the 321-layer three-deck node ships in v1.0)
4. **Cryogenic-etch coverage for $\geq$400-layer NAND**: A cryogenic high-aspect-ratio etch process file exists and is wired as an opt-in variant on current ($\leq$321-layer) NAND models, but the $\geq$400-layer NAND wafer models that would make cryo etch the default HARC path do not yet exist. These models are required before cryo-etch coverage can be considered complete for the highest-layer-count architectures.
5. **Connector manufacturing depth**: Process-level detail for connector fabrication (stamping, plating, ferrule manufacturing for fiber connectors) remains thinner than for semiconductor categories; ceramic ferrule manufacturing is an explicit open gap.
6. **Specialty-wafer process completeness**: A small number of planned specialty-wafer operations do not yet have complete inventories; they are explicitly marked as gaps in the affected models rather than filled with placeholder values.
7. **Yield-scenario variants for logic wafers**: The HBM4 stacks ship ramp- and mature-yield scenario datasets alongside their baseline (Section 5.5 of Chapter 5). Extending the same published-scenario pattern to logic-wafer die yields — per node, distinguishing ramp from mature production — would let users match wafer datasets to the production maturity of their supply without re-deriving yield assumptions.
8. **Two still-planned models**: A smartphone application-processor (SoC) model and an SSD-controller model remain planned but unauthored as of v1.0.
9. **EUV high-NA evolution**: Power, throughput, and consumable parameters for ASML EXE:5000-class systems
10. **$CaF_2$ sludge custom dataset**: Currently routed to underground deposit (conservative); >90 % of Taiwan/Korea $CaF_2$ sludge is recovered to steel-industry flux. Custom dataset planned for a future release.
11. **Solder dross custom dataset**: Currently routed to the hazardous-inorganic disposal route (conservative); proper Sn-refining cut-off market activity needed.
12. **Regional electricity for non-fab categories**: PCBs and passives currently use China (CN) defaults; Taiwan/Vietnam/Mexico variants would improve geographic specificity.
13. **Per-ISO-class central chilled-water-plant intensities**: v1.0 adds make-up air and central exhaust as explicit flows and adopts the recipe-driven ballroom + minienvironment HVAC base, but the central chilled-water-plant contribution is still folded into the zone intensities rather than resolved per ISO class. A future release should resolve it per ISO

class from semiconductor-specific peer-reviewed sources and validate the recipe-fleet calibration across multiple nodes.

14. **Capacity-dependent HVAC scaling**: Constant per-bucket HVAC anchors are appropriate for the mid-capacity range but lose accuracy at $<$10k or $>$200k WPM extremes. A small residual economy-of-scale factor or sub-tier buckets would refine a future release.
15. **Ultrapure-water per-pass rates**: Ultrapure-water demand is currently estimated from per-pass rates for wet cleaning, lithography, and CMP together with tool-specific water inputs; these bases may partially overlap and were recalibrated against process anchors as an interim measure. A future release will replace them with a facility-total allocation method parallel to the cleanroom-HVAC base.

### 6.6.2 Methodology Enhancements

Planned methodology improvements:

**Distribution-based uncertainty propagation** Version 1.0 propagates uncertainty end-to-end as min/typical/max intervals (Section 3.3). A future enhancement is sampling-based (Monte Carlo) propagation over full probability distributions, including correlations between parameters, which would yield confidence intervals rather than the bounding intervals published today.

**Automatic output-flow band inheritance** Uncertainty bands on emission and waste outputs are published only where a range is established for the output itself; where an output is driven by a banded input, the output's band is closed by explicit authoring. A mechanism that would derive such output bands automatically from their driving inputs was designed during v1.0 development but not enabled (Section 6.3.8).

**Capital goods inclusion** Equipment embodied energy allocation when sufficient data becomes available.

### 6.6.3 Coverage Expansion

Planned additions to component coverage:

- Display components (LCD, OLED driver ICs)
- Advanced memory (next-generation HBM)
- Chiplet and heterogeneous integration packages
- Advanced cables and interconnects

## 6.7 Critical Review

### 6.7.1 Internal Review Process

The database has undergone internal technical review including:

- Consistency checks across LCI datasets (process-gas mass balance, water balance, unit and range checks; Section 3.6.1)
- Benchmarking against published facility intensities
- Cross-verification of data extraction from sources

- Peer review of calculation methodology

### 6.7.2 External Review Statement

The v1.0 database and this methodology report, dated 21 August 2026, have undergone an independent critical review, completed on 25 August 2026. The review assessed whether the methods underlying the inventory modeling and data collection are consistent with recognized guidance and best practices; whether the described methods and assumptions are scientifically and technically valid; whether the collected data are appropriate and reasonable for the goals of the database; and whether the methodology report and additional documentation are transparent and consistent. It was conducted in a series of one-to-one sessions during July and August 2026, together with offline review of the requested documentation and custom reports extracted from the database. The review did not include a verification of individual numerical values.

The signed review statement is reproduced verbatim in Appendix D. Its own terms apply as written: the statement is valid only for the specific report named in its title, dated 21 August 2026, and information received and discussed up to that date; the changelog in the Version History records revisions made after that date.

# Appendix A

# Glossary

## Semiconductor Manufacturing Terms

**ALD (Atomic Layer Deposition)**
: A thin-film deposition technique using sequential self-limiting surface reactions to achieve atomic-level thickness control.

**BEOL (Back-End-Of-Line)**
: The portion of IC fabrication where individual devices are interconnected with metal wiring. Includes contact, via, and metal layer formation.

**CMP (Chemical-Mechanical Planarization)**
: A polishing process combining chemical etching and mechanical abrasion to create flat wafer surfaces.

**CVD (Chemical Vapor Deposition)**
: A deposition process where gaseous precursors react on a heated substrate to form thin films.

**DUV (Deep Ultraviolet)**
: Lithography using 248nm (KrF) or 193nm (ArF) light sources. Includes dry and immersion variants.

**EUV (Extreme Ultraviolet)**
: Lithography using 13.5nm light, enabling sub-7nm feature patterning with fewer masks than DUV multi-patterning.

**FEOL (Front-End-Of-Line)**
: The portion of IC fabrication where individual transistors are formed. Includes STI, gate, and source/drain formation.

**FinFET**
: A 3D transistor architecture with a fin-shaped channel, used at 22nm and below for improved electrostatic control.

**GAA (Gate-All-Around)**
: A transistor architecture where the gate surrounds the channel on all sides, used at 3nm and below.

**Little's Law**
: The queueing identity $L = \lambda W$: the average number of items in a system equals their arrival rate multiplied by the average time each spends in it. Used in this database to derive fab tool fleets from process recipes (Chapter 2).

**Low-k Dielectric**
: Insulating materials with dielectric constant below $SiO_2$ ($k < 3.9$), used to reduce RC

delay in BEOL.

**MOL (Middle-Of-Line)**
The transition region between FEOL and BEOL, including local interconnects and contacts.

**PFC (Perfluorocarbon)**
Fluorinated gases ($CF_4$, $C_2F_6$, $SF_6$, $NF_3$) used in etching and chamber cleaning with high global warming potential.

**PVD (Physical Vapor Deposition)**
Deposition processes using physical mechanisms (sputtering, evaporation) rather than chemical reactions.

**Reticle**
The patterned mask used in lithography to project circuit patterns onto the wafer.

**STI (Shallow Trench Isolation)**
Isolation structures formed by etching trenches and filling with oxide to separate transistors.

**UPW (Ultrapure Water)**
Highly purified water (18.2 M$\Omega$-cm resistivity) used for wafer cleaning and rinsing.

## Packaging Terms

**BGA (Ball Grid Array)**
Package type with solder balls on the bottom surface for board attachment.

**CoWoS (Chip-on-Wafer-on-Substrate)**
TSMC's 2.5D packaging technology using a silicon interposer.

**Flip-chip**
Interconnection method where the die is flipped and bonded face-down using solder bumps.

**HBM (High Bandwidth Memory)**
3D-stacked DRAM using TSVs for high bandwidth memory applications.

**Interposer**
A substrate (silicon or organic) providing electrical connections between dies in advanced packages.

**Leadframe**
Metal frame providing external connections in traditional packages (QFN, SOIC).

**TSV (Through-Silicon Via)**
Vertical electrical connections passing through silicon dies for 3D integration.

**Wire Bond**
Traditional interconnection using thin wires (gold or copper) between die pads and package leads.

## LCA Terms

**Attributional LCA**
: LCA approach describing the environmentally relevant flows associated with a product system.

**Cradle-to-Gate**
: System boundary from raw material extraction through manufacturing, excluding use and end-of-life.

**DQI (Data Quality Indicator)**
: Metric assessing the quality of LCI data across multiple dimensions.

**Elementary Flow**
: Flow from or to the environment (e.g., $CO_2$ emission, ore extraction).

**Functional Unit**
: Quantified performance of a product system for use as a reference unit.

**GWP (Global Warming Potential)**
: Impact category measuring climate change contribution relative to $CO_2$.

**LCI (Life Cycle Inventory)**
: Phase of LCA involving compilation of inputs and outputs of a product system.

**LCIA (Life Cycle Impact Assessment)**
: Phase of LCA aimed at understanding environmental significance of LCI results.

**Pedigree Matrix**
: Framework for assessing data quality across multiple dimensions (reliability, completeness, etc.).

## Units and Abbreviations

**kWh**
: Kilowatt-hour (energy)

**MJ** Megajoule (energy)

**L** Liter (volume)

**kg** Kilogram (mass)

**g** Gram (mass)

**mg** Milligram (mass)

**sccm**
: Standard cubic centimeters per minute (gas flow)

**slm** Standard liters per minute (gas flow)

**WPH**
: Wafers per hour (throughput)

**nm** Nanometer (length, technology node)

**mm** Millimeter (length)

**$\mu$m** Micrometer (length)

# Appendix B

# Process Inventory List

This appendix provides a representative list of the process inventories included in the database, organized by manufacturing category. It is not exhaustive: the full inventory comprises more than 270 process files, and the tables below highlight the principal process families across each sector. Some processes listed below are modeled as sub-steps within broader process files rather than having dedicated standalone inventories (e.g., copper seed deposition is included within the PVD barrier process inventory).

## B.1 Semiconductor Processes

### B.1.1 Lithography

**Table B.1:** Lithography process inventories

| Process | Description |
|---|---|
| EUV lithography | Extreme ultraviolet exposure at 13.5nm wavelength for advanced patterning |
| DUV immersion lithography | ArF immersion lithography at 193nm for high-resolution patterning |
| DUV dry lithography | ArF/KrF dry lithography for mature nodes |
| I-line lithography | 365nm lithography for legacy and non-critical layers |

### B.1.2 Etch

**Table B.2:** Etch process inventories

| Process | Description |
|---|---|
| Silicon plasma etch (shallow) | Reactive ion etching for silicon and polysilicon features ($HBr/Cl_2$-led chemistry; gate, spacer, STI-class steps, and shallow photonics waveguide/grating etches in the same depth class) |
| Silicon plasma etch (deep Bosch) | Time-multiplexed $SF_6/C_4F_8$ deep reactive ion etching for MEMS and TSV-class structures |
| Dielectric etch | Plasma etching for $SiO_2$ and low-k dielectrics |
| Metal etch | Metal layer patterning by plasma or wet etch |

*Table B.2 — continued from previous page*

| Process | Description |
|---|---|
| High aspect ratio etch | Deep contact and via etching for advanced interconnects |

### B.1.3 Deposition

**Table B.3:** Deposition process inventories

| Process | Description |
|---|---|
| Thermal oxidation | High-temperature oxide growth for gate and isolation |
| PECVD oxide | Plasma-enhanced CVD for interlayer dielectrics |
| PECVD nitride | Silicon nitride deposition for barriers and spacers |
| ALD high-k | Atomic layer deposition for gate dielectrics |
| ALD metal gate | Work function metal deposition by ALD |
| CVD tungsten | Chemical vapor deposition for contact fill |
| PVD barrier | Physical vapor deposition for diffusion barriers |
| PVD copper seed | Copper seed layer for electroplating |
| Epitaxy | Epitaxial Si/SiGe growth for strain engineering |
| SiC epitaxy CVD | CVD growth of SiC drift layers at 1500–1650$^{\circ}\mathrm{C}$ for power devices |
| SiC wafer growth | Physical vapor transport (PVT) bulk SiC boule growth and wafer slicing for power substrates |
| Ge photodetector epitaxy | Selective two-step germanium-on-silicon epitaxy for silicon photonics photodetectors |

### B.1.4 Planarization and Cleaning

**Table B.4:** CMP and cleaning process inventories

| Process | Description |
|---|---|
| Oxide CMP | Chemical mechanical polishing for dielectric planarization |
| Metal CMP | Copper and tungsten CMP for damascene metallization |
| Wet cleaning | Multi-step wet chemical cleaning sequences |
| Post-CMP cleaning | Slurry residue removal and surface preparation |

### B.1.5 Other Front-End Processes

**Table B.5:** Other semiconductor process inventories

| Process | Description |
|---|---|
| Ion implantation | Dopant introduction for source/drain and well formation |
| Laser anneal | High-temperature activation and defect annealing |
| Copper electroplating | Electrochemical deposition for damascene metallization |
| Chamber clean | In-situ chamber cleaning using $NF_3$ or other gases |

## B.2 Packaging Processes

**Table B.6:** Packaging process inventories

| Process | Description |
|---|---|
| Wafer backgrind | Wafer thinning for package integration |
| Wafer backgrind (ultrathin) | Thinning to 50–100 $\mu$m for HBM die stacking |
| Wafer dicing | Die singulation by blade or laser sawing |
| Epoxy die attach | Die bonding using epoxy adhesive |
| Solder die attach | Die bonding using solder preform |
| Gold wire bonding | Interconnection using gold wire |
| Copper wire bonding | Interconnection using copper wire |
| Flip-chip bumping | Solder bump formation for flip-chip assembly |
| Flip-chip reflow | Solder joint formation by reflow |
| Underfill | Epoxy dispensing for mechanical reinforcement |
| Mold encapsulation | Epoxy mold compound encapsulation |
| Post-mold cure | Thermal cure of mold compound for cross-linking |
| LED encapsulation | Silicone/epoxy encapsulation with phosphor for LED packages |
| Package singulation | Individual package separation |
| Package marking | Laser marking for identification |
| Final test | Electrical parametric and functional testing |

## B.3 Advanced Packaging Processes

**Table B.7:** Advanced packaging process inventories (2.5D/3D integration)

| Process | Description |
|---|---|
| EMIB bridge placement | Embedded multi-die interconnect bridge assembly |
| Thermocompression bonding | Fine-pitch die bonding for 2.5D/3D integration |
| HBM die stacking | High bandwidth memory die stack assembly |
| HBM TSV stacking | Through-silicon via layer interconnection |
| HBM4 TSV stacking | TSV formation and multi-die stacking for HBM4 (2048-bit, 32-channel interface) |
| CBA wafer bonding | CMOS-bonded-array face-to-face wafer bonding for advanced 3D NAND |
| UBM deposition | Under bump metallization for flip-chip |
| Solder ball placement | Ball grid array attachment |
| Flux application | Flux dispensing for soldering |
| Passivation opening | Via opening in passivation layers |
| Through-mold via | Vertical interconnect through mold compound |
| Lid attachment | Heat spreader attachment and sealing |
| Carrier debonding | Temporary carrier removal after processing |
| Burn-in test | Accelerated reliability testing |

## B.4 Power Module Assembly Processes

**Table B.8:** Power module assembly process inventories

| Process | Description |
|---|---|
| DBC substrate | Direct bonded copper substrate via Cu–O eutectic bonding to ceramic |
| AMB substrate | Active metal brazing of copper to $Si_3N_4$/AlN ceramic substrate |
| Silver-sinter die attach | Pressure-assisted Ag-sinter die attach for SiC/GaN power semiconductors |
| Baseplate assembly | Solder bonding of substrate to power module baseplate |
| TIM application | Thermal interface material application at module-to-heatsink interface |

## B.5 Silicon Photonics and Optical Fiber Processes

**Table B.9:** Silicon photonics and optical fiber process inventories

| Process | Description |
|---|---|
| Ge photodetector epitaxy | Selective germanium-on-silicon epitaxy for silicon photonics photodetectors |
| Photonic fiber coupling | Active-alignment fiber-array attachment (edge/grating coupler) to a silicon photonics PIC |
| TOSA assembly | Transmit optical sub-assembly: EML laser-die bonding, lens and optical-isolator placement, high-precision active alignment, and hermetic sealing (per 1000 assemblies) |
| ROSA assembly | Receive optical sub-assembly: InGaAs photodiode and transimpedance-amplifier die integration with optical interface (per 1000 assemblies; no isolator, lower alignment precision than TOSA) |
| VCSEL TOSA assembly | Multimode transmit sub-assembly for GaAs VCSEL arrays (per 1000 assemblies; no isolator or hermetic seal, direct chip-on-board mounting) |
| Fiber preform OVD | Outside vapor deposition of doped silica optical-fiber preform |
| Fiber draw | High-speed fiber drawing from preform in a graphite resistance furnace ($\sim$2000°C) |
| Fiber UV coating | In-line dual-layer UV-cured acrylate coating of drawn optical fiber |

## B.6 PCB Processes

**Table B.10:** PCB process inventories

| Process | Description |
|---|---|
| Imaging (LDI) | Laser direct imaging of photoresist; 0.071 kWh/m$^2$ per pass from equipment specs |
| Desmear ($KMnO_4$) | Chemical smear removal after drilling; 2.5 kWh/m$^2$ conveyorized wet chemistry |
| Oxide treatment | Inner layer adhesion (brown oxide CuO); 2.0 kWh/m$^2$ heated vertical dip line from vendor data |
| Lamination | High-temperature pressing of multilayer stacks |
| Through-hole drilling | Mechanical CNC drilling; 0.115 kWh per 1000 holes from spindle power and throughput |

*Table B.10 — continued from previous page*

| Process | Description |
|---|---|
| Laser via drilling | $CO_2$/UV laser microvia formation; 0.0017 kWh per 1000 vias |
| Back-drilling | Controlled-depth stub removal; 0.23 kWh per 1000 holes |
| Electroless copper | Seed layer deposition for plating; formaldehyde-based reduction |
| mSAP/SAP electroless copper | Ultra-thin electroless Cu seed deposition for modified/full semi-additive fine-line processing |
| mSAP/SAP pattern plating | Selective acid-copper electroplating into trace trenches for fine-pitch (semi-additive) builds |
| Pattern plating | Selective electrolytic copper plating for traces and vias |
| Etching | Copper pattern definition; cupric chloride or alkaline ammonia |
| Solder mask | LPI solder mask application; 5.1 kWh/$m^2$ from equipment specs |
| HASL finish | Hot air solder leveling surface finish |
| ENIG finish | Electroless nickel/immersion gold; 8.0 kWh/$m^2$ |
| OSP finish | Organic solderability preservative coating |
| Hard gold plating | Electrolytic gold for edge connectors; 2.5 kWh/$m^2$ board area |
| Silkscreen and routing | Legend printing (inkjet) plus CNC board singulation; 0.35 kWh/$m^2$ |
| Electrical test | Flying probe continuity/isolation testing; 0.5 kWh/$m^2$ |
| Panel consumables | Edge trim, test coupons, and routing waste material tracking |

## B.7 Passive Component Processes

**Table B.11:** Passive component process inventories

| Process | Description |
|---|---|
| Ceramic tape casting | Green sheet preparation for MLCC |
| Electrode printing | Internal electrode screen printing |
| Layer lamination | Multilayer stacking and pressing |
| Sintering | High-temperature ceramic densification |
| Termination plating | External electrode formation |
| Thick-film resistor | Screen-printed resistive element processing |
| Thin-film resistor | Vacuum-deposited precision resistor processing |
| Inductor winding | Wire winding for magnetic components |
| Transformer winding | Wound power transformer manufacturing for switching power supplies |
| Ferrite processing | Ferrite core sintering and preparation |
| Tantalum powder reduction | Magnesiothermic (vapor-phase) reduction of $Ta_2O_5$ to capacitor-grade Ta powder |
| Tantalum anode sintering | Powder agglomeration, pressing around riser wire, and high-vacuum sintering of porous anode |
| Tantalum cathode formation | Anodization of $Ta_2O_5$ dielectric and cathode ($MnO_2$/polymer) formation with encapsulation |

## B.8 Cable and Connector Processes

**Table B.12:** Cable and connector process inventories

| Process | Description |
|---|---|
| Connector contact stamping | Progressive-die stamping of copper-alloy strip into connector contacts |
| Connector plating (Sn/Ni/Au) | Tin, nickel, and gold electroplating of connector contacts |
| Connector latch fabrication | Latch and retention-feature manufacturing for connector housings |
| Connector assembly | Contact insertion and housing assembly |
| Connector test and packaging | Electrical/mechanical test and packaging of finished connectors |
| LCP injection molding | Glass-filled liquid crystal polymer molding of high-speed connector housings |
| PA9T injection molding | PA9T (polyamide) molding of connector housings |

*Table B.12 — continued from previous page*

| Process | Description |
|---|---|
| Cable termination | Jacket stripping, contact crimp/IDC, and connector attachment |
| Twisted-pair formation | Conductor twisting for differential signal pairs |
| Insulation extrusion | FEP/PE/PVC/TPE insulation extrusion over conductors |
| Cable shielding | Foil and braid shielding application |
| Copper wire drawing | Multi-pass drawing of copper conductor wire |
| Cable assembly testing | Continuity, isolation, and signal-integrity testing of cable assemblies |

## B.9 Thermal Management Processes

**Table B.13:** Thermal management and liquid-cooling infrastructure process inventories

| Process | Description |
|---|---|
| Heat pipe fabrication | Copper heat pipe manufacturing and charging |
| Vapor chamber fabrication | Two-phase thermal spreader manufacturing |
| Furnace brazing | High-temperature (vacuum default) joining for thermal assemblies |
| Tube fabrication | Stainless steel tube forming for cooling systems |
| Stainless tube fabrication | ERW/seamless stainless steel tube fabrication for liquid-cooling manifolds |
| Orbital TIG welding | Automated autogenous orbital GTAW of thin-wall stainless tube joints (ASME BPE) |
| Citric-acid passivation | Citric-acid passivation of stainless steel surfaces (ASTM A967) |
| TIM preparation | Thermal interface material mixing and filling |
| TIM pad manufacturing | Solid thermal pad production |
| TIM application | Thermal interface material application at the module-to-heatsink interface |
| Heatsink extrusion | Aluminum heatsink profile extrusion |

## B.10 Facility Support Processes

**Table B.14:** Facility support process inventories

| Process | Description |
|---|---|
| Cleanroom HVAC | Air handling for cleanroom temperature and particle control |
| UPW production | Ultrapure water generation and distribution |
| Bulk gas supply | Nitrogen, argon, oxygen supply systems |
| Gas abatement | PFC and VOC emission treatment |
| Wastewater treatment | Industrial effluent treatment and discharge |

# Appendix C

# Data Sources Summary

This appendix summarizes the primary data sources used in constructing the database.

## C.1 Source Categories

### C.1.1 Equipment Vendor Documentation

| Vendor | Coverage | Source Type |
|---|---|---|
| ASML | EUV/DUV lithography systems | Technical specs, sustainability reports |
| Applied Materials | Etch, CVD, PVD systems | Technical papers, press releases |
| Lam Research | Etch, deposition, CMP | Technical documentation |
| Tokyo Electron (TEL) | Coating, thermal processing | Equipment specifications |
| KLA | Metrology, inspection | Technical papers |
| Edwards Vacuum | Abatement systems | Product documentation |

**Table C.1:** Equipment vendor sources

### C.1.2 Academic and Research Publications

Primary sources of peer-reviewed data:

- IEEE International Electron Devices Meeting (IEDM)
- VLSI Technology Symposium
- Journal of the Electrochemical Society
- Solid State Technology
- Semiconductor Engineering

### C.1.3 Industry Roadmaps

- IEEE International Roadmap for Devices and Systems (IRDS) – public portions
- SEMI Standards and technology roadmaps

### C.1.4 Corporate Sustainability Reports

Manufacturing facility-level data from:

- TSMC Environmental Report

- Samsung Electronics Sustainability Report
- Intel Corporate Responsibility Report
- GlobalFoundries Environmental Reports

### C.1.5 Custom Upstream Dataset Sources

For semiconductor-specific materials not available in standard databases, custom datasets were developed from:

- **Patent literature**: Production process descriptions and stoichiometry (e.g., US 2006/0075959 A1 for TMGa synthesis, CA 2054511A1 for TMAl, US 5756786A for TMIn)
- **Academic LCA studies**: Published life cycle assessments of specialty chemicals (e.g., Smith et al. 2018 on III-V precursors)
- **Industry factsheets**: Critical raw materials data from SCRREEN2 project (gallium, indium)
- **Thermodynamic calculations**: Energy requirements estimated from reaction enthalpies and process conditions

## C.2 Per-Flow Source Attribution

Beginning with v1.0, source attribution can be carried at the level of the individual inventory flow rather than only at the dataset level. Published flows list the sources behind their values — equipment specifications, peer-reviewed papers, patents, government reports, and corporate disclosures — either established for the specific value or inherited from the process inventory or dataset they derive from; not every flow yet carries a resolved citation, and Section 6.3.9 explains both the distinction and the residual. When inventories are aggregated, these references are preserved, so an aggregated dataset (for example, an integrated circuit assembled from a wafer and a package) links to the source lineage of its upstream contributors.

The resolved per-flow citations travel with the published datasets themselves and can be browsed alongside each dataset's flow list in the dataset explorer on the project website (`https://reellci.com`); this appendix lists the principal source *categories*.

## C.3 Reference Fabrication Facilities (Fab-Overhead Anchoring)

The v1.0 wafer-fab overhead methodology anchors cleanroom area and facility-energy intensities to publicly disclosed reference fabrication facilities. Each anchor is expressed as cleanroom area per unit of wafer-start capacity and is used only to scale facility overhead; no proprietary operating data is used. The facilities below were selected because their cleanroom dimensions are available in public disclosures or, where noted, independent industry reporting.

| Facility | Segment / Node | Disclosure basis |
|---|---|---|
| TSMC Fab 18 | Advanced logic (3/5nm), Taiwan | Company press release; per-phase cleanroom area (TSMC 2022) |
| GlobalFoundries Fab 8 | Mature HKMG logic, Malta NY, USA | Independent industry reporting (facility tours, trade press); area not company-published |
| SK Hynix M16 | Leading-edge DRAM (1c), Cheongju, Korea | Public facility disclosure |
| SK Hynix M15X | Leading-edge DRAM, Cheongju, Korea | Public facility disclosure |
| Wolfspeed Mohawk Valley | 200mm SiC power, NY, USA | Owner's-representative documentation (Aubertine & Currier 2024) |
| Sony Nagasaki (Fab 1, Fab 2) | CMOS image sensors, Japan | Foundry-disclosed cleanroom area |
| Bosch Reutlingen | Power discrete / MEMS & analog, Germany | Company press disclosure of cleanroom expansion (Bosch 2022) |

**Table C.2:** v1.0 reference fabrication facilities used for fab-overhead anchoring

## C.4 Methodology and Standards Sources

### C.4.1 Energy Conversion and Roadmap Standards

- **SEMI S23-1021** (SEMI 2021) — Guide for conservation of energy, utilities, and materials used by semiconductor manufacturing equipment. Provides the energy-conversion-factor table used to map facility utility consumption (compressed dry air, process cooling water, exhaust, etc.) to an electricity-equivalent basis.
- **IRDS 2024 ESHS/ESSF tables** (IEEE IRDS 2024) — The Environment, Safety, Health and Sustainability (ESSF) tables of the IEEE International Roadmap for Devices and Systems, used for fab-blended utility baselines (e.g., ultrapure-water production intensity).

### C.4.2 Device-Technology and Reliability Standards (JEDEC)

JEDEC JEP-series guidance documents are cited where they standardize the device classes and technology platforms that the inventory models cover. These documents delimit process and reliability scopes only; none supplies inventory data.

- **JEP375** (JEDEC Solid State Technology Association 2025a) — reliability guidelines for 2.5D/3D silicon thinned wafers with Cu TSVs and backside RDL; matches the process scope of the CoWoS, EMIB, Foveros, and TSV-stacking inventories.

- **JEP158A** (JEDEC Solid State Technology Association 2026a) — reliability-interaction guidelines for 3D chip stacks with through-silicon vias; covers the HBM3/3e/4 stacking platform.
- **JEP204** (JEDEC Solid State Technology Association 2026b) — catalog of stress procedures for silicon carbide devices (PECS); delimits the SiC power-device class.
- **JEP182** (JEDEC Solid State Technology Association 2026c) — continuous-switching evaluation of GaN power-conversion devices; delimits the GaN HEMT device class.
- **JEP110A** (JEDEC Solid State Technology Association 2025c) — thermal characterization and modeling of compound-semiconductor (III-V) FETs.
- **JEP30G.01** (JEDEC Solid State Technology Association 2025b) — part-model XML guidelines for electronic-device packages; monitored as a potential future venue for chiplet-level sustainability metadata (Supply Chain sub-schema JEP30-S101).

### C.4.3 Cleanroom-Energy Literature

Facility-overhead intensities (recirculation/air-handling, make-up air, exhaust, and cooling subsystems) are anchored to the peer-reviewed cleanroom-energy literature:

- Hu et al. (2019) (Hu, Lin, et al. 2019) — measured subsystem energy-conversion factors for high-tech fabrication plants.
- Sun & Chang (2017) (Sun and Chang 2007) — power-consumption benchmark for a semiconductor cleanroom facility system.
- Yin et al. (2020) (Yin et al. 2020) — cleanroom environmental-control performance and cold–heat offset, basis for the make-up-air baseline.
- LBNL cleanroom field studies (Xu & Tschudi, 2002) (Xu and Tschudi 2002) and the associated PG&E cleanroom baseline model (Pacific Gas and Electric Company 2010) — per-cleanliness-class recirculation energy intensities used to cross-validate the modeled zone intensities.
- Hu & Chuah (2003) and Hu & Tsao (2007) (Hu and Chuah 2003; Hu and Tsao 2007) — measured fab power densities and HVAC subsystem decomposition for Taiwanese high-tech fabs.

## C.5 Sector-Specific Source Domains (v1.0)

v1.0 extends inventory coverage into additional sectors, each with its own primary source domain:

- **Silicon photonics and optical fiber**: Energy and material inventories for optical-fiber production (preform deposition, draw, coating) draw on the published fiber-optics carbon-footprint case study of Inakollu, Morin & Keefe (2017, *Sustainability* 9(5):865, OFS Fitel), used for the fiber-energy allocation between flame-deposition and electrical loads. Silicon photonic device flows reuse the logic/specialty wafer process inventories with photonic-specific additions (e.g., germanium photodetector wet-etch emissions).
- **Power modules and discrete semiconductors**: Packaging and assembly inventories reference the relevant JEDEC/JEP solder, finish, and reliability standards (e.g., JEDEC

Pb-free solder classifications) for material specification and bill-of-materials definition. These standards define material composition and test conditions only; no proprietary process data is used.

# Appendix D

# Critical Review Statement

This appendix reproduces, verbatim, the signed statement of the independent critical review of the REEL LCI Database, Version 1.0, dated 21 August 2026. The review was completed on 25 August 2026 and did not include a verification of individual numerical values. The statement's own scope, validity, and signature terms apply as written on the following pages.

## - Critical Review Statement -

## Review of "REEL LCI Database; Version 1.0; 21 August 2026"

Reviewer: Arpad Horvath; Consultant; Berkeley, California

Date of completion of review: August 25, 2026

### Scope of the Critical Review

The Critical Review focused on the "REEL LCI Methodology Report" and custom reports extracted from the REEL LCI database upon the Reviewer's request.

The goals of the Critical Review were to assess whether:

- the methods underlying the life-cycle inventory (LCI) and data collection are consistent with the guidance in ISO 14040 and ISO 14044 (to the extent they are applicable) and the best practices,
- the described methods and assumptions are scientifically and technically valid,
- the collected data are appropriate and reasonable for the goals of the database,
- the methodology report and additional documentation are transparent and consistent.

### Critical Review Process

The review was conducted in 6 one-to-one sessions (carried out in July and August 2026) that reviewed all aspects of the REEL LCI database and additional, offline reviews of the requested and provided documentation of REEL LCI. The Reviewer received all the information he requested, and additional information was provided by the database author.

### General Evaluation

The REEL LCI database provides an inventory of cradle-to-gate data (energy, water, and material inputs and air emissions and wastes) for 18 categories of electronics components: semiconductor-integrated circuits across technology nodes from 180 nm to 3 nm, memory devices (DRAM, NAND flash, HBM), printed circuit boards, passive components (capacitors, resistors, inductors), integrated-circuit packaging, cables and interconnects, and electromechanical components. Of the life cycle stages, manufacturing (including processing, facility overhead, and testing) and transportation are included.

The way the database was set up includes several distinctive features. Methods and data are derived exclusively from publicly available sources. It is using a bottom-up "virtual factory" modeling approach. Most values in the database are calculated rather than copied from a single reference. A handful of specialty inputs have no public dataset, so the database author estimated them or used a documented stand-in. The approach to

methodology is, from best to last resort: (1) build a transparent dataset from public chemistry and energy data; (2) use a documented same-family stand-in; (3) cut it off with a written reason if it is below materiality. It reports physical units for elementary and technosphere flows (kWh, liters, kilograms), not environmental scores, so anyone can apply their own impact method. Water use is reported net of recycling, not as the total demand of a facility.

The REEL LCI database has several distinctive strengths.

The range of included electronics components is significant and representative. The components are modeled at a sufficiently granular process (and some on flow) level, with important and necessary details provided, rather than reporting just one number for a technology (e.g., 40 nm wafer) that then one uses for reporting or extrapolation. Gases and precursors are shown for fab processes. Allocation calculations are documented. There are hundreds of datasets in the database.

There is necessary and expected transparency present in the documentation. There are no silent omissions. It is stated what is included in the REEL LCI database versus what is not. Cut-off decisions are described. Every assumption is documented and independently verifiable.

Every numerical value traces to a publicly available source (equipment specifications, peer-reviewed articles, patents, government data, corporate sustainability reports) or an engineering estimate and is documented on process or flow level. No licensed commercial databases are used for inputs or outputs values. The review did not include a verification of individual numerical values.

For every source, the basis of how numerical values were derived from it are stated (evidence-based derivations, declared estimates, or “flagged conflicts” when the evidence is inconclusive).

Detailed citation maps for components’ wafer, printed circuit board, packaging, upstream materials, and facility overhead stages were made available to the Reviewer. Each specific value (or method) in the database connects to one or more sources, with a link to the source and a short note on how the source was actually used, as well as where to find the information. The links go to the original documents. The database also holds archived copies of nearly all source documents in case a link becomes dead.

Uncertainty is accounted for and discussed. The Reviewer was provided with “pedigree matrix” scores for data quality and uncertainty ranges for the data, with recorded exemptions.

The Reviewer was able to review two detailed examples of how a numerical value is created and cited.

The datasets and their documentation can be viewed in the Data Explorer, accessible at reellci.com.

### Conclusions

The methods underlying the life-cycle inventory (LCI) modeling and data collection are consistent with the guidance in ISO 14040 and ISO 14044 (to the extent they are applicable) and the best practices.

The described methods and assumptions are scientifically and technically valid.

The used data are appropriate and reasonable in this moment in time with respect to the goal of the database and public availability of data.

The methodology report is sufficiently transparent and consistent. It documents every aspect of the database: system boundaries, allocation procedures, data quality indicators, and sector-specific modeling approaches.

This review statement is only valid for the specific report named in the title, dated August 21, 2026, and information received and discussed up to the same date, thus, it does not apply to any other report versions, derivative reports, excerpts, press releases, and similar documents.

The Critical Reviewer signs this review statement as an independent expert. No statement should be construed as comparative. The Reviewer's signature does not constitute an endorsement of the REEL LCI's potential uses or results.

Arpad Horvath

# References


ASML Holding N.V. (2024). *EUV Lithography Technology Overview*. ASML. URL: https://www.asml.com/.

Aubertine & Currier (2024). *Wolfspeed Mohawk Valley Fab Owner's Representative Documentation*. "Three level, 200 mm SiC cleanroom with approximately 100,000 to 135,000 SF of cleanroom space" = 9,290–12,540 $m^2$. Independent secondary cite for Wolfspeed Mohawk Valley cleanroom area; corroborates REEL value 11,148 $m^2$ at midpoint.

Bosch (2022). *Reutlingen Wafer Fab Cleanroom Expansion*. "Cleanroom space in Reutlingen is set to grow from around 35,000 $m^2$ at present to over 44,000 $m^2$ by the end of 2025" (semiconductor-today.com via Bosch press July 2022). Used for REEL specialty_si_power_discrete and specialty_mems_analog anchors with shared-cleanroom allocation caveat.

Boyd, S. B., A. Horvath, and D. Dornfeld (2009). "Life-Cycle Energy Demand and Global Warming Potential of Computational Logic". In: *Environmental Science & Technology* 43.19, pp. 7303–7309. DOI: 10.1021/es901514n.

Chen and S.-C. Hu (2009). "Specific Energy Consumption (SEC) for the Integrated Circuit Assembly and Testing (IC A/T) Industry in Taiwan". In: *ASHRAE Transactions* 115.2. Survey of eleven Taiwanese IC assembly-and-test plants; facility systems approximately 53 percent of site electricity. URL: https://www.researchgate.net/publication/237312155_Specific_Energy_Consumption_SEC_for_the_Integrated_Circuit_Assembly_and_Testing_IC_AT_Industry_in_Taiwan.

Ciroth, A., S. Muller, B. Weidema, and P. Lesage (2016). "Empirically based uncertainty factors for the pedigree matrix in ecoinvent". In: *International Journal of Life Cycle Assessment* 21.9, pp. 1338–1348. DOI: 10.1007/s11367-013-0670-5.

Hu, S.-C. and Y. K. Chuah (2003). "Power consumption of semiconductor fabs in Taiwan". In: *Energy* 28.8. Measured survey of nine representative Taiwan semiconductor fabs: average total power consumption 2.18 kW/$m^2$, average cooling load 0.434 RT/$m^2$; facility systems 56.6 % of total power, process tools 40.4 %. Primary source for REEL's 2.18 kW/$m^2$ fab power density. Monolithic per-area fab average — no ISO-class or logic/memory breakdown., pp. 895–907. URL: https://ideas.repec.org/a/eee/energy/v28y2003i8p895-907.html.

Hu, S.-C., T. Lin, B.-R. Fu, C.-K. Chang, and I.-Y. Cheng (2019). "Analysis of energy efficiency improvement of high-tech fabrication plants". In: *International Journal of Low-Carbon Technologies* 14.4. Measured ECFs for Taiwanese semiconductor fab subsystems. Exhaust (excluding MAU) 0.000723 kWh/$m^3$ corroborates Sun 2017 fan-only ECF 0.0007 within 3%., pp. 508–515. DOI: 10.1093/ijlct/ctz031.

Hu, S.-C. and J.-M. Tsao (2007). "A comparative study on energy consumption for HVAC systems of high-tech FABs". In: *Applied Thermal Engineering* 27.17–18. MAU pressure drop and decomposition. Documents the internal composition of the cleanroom HVAC baseline;

basis for REEL's removal of a separate per-process-pass HVAC term., pp. 2758–2766. DOI: 10.1016/j.applthermaleng.2007.03.011.

IEEE International Roadmap for Devices and Systems (2024). *IRDS 2024 Edition*. Tech. rep. IEEE. URL: https://irds.ieee.org/.

IEEE IRDS (2024). *International Roadmap for Devices and Systems — Environment, Safety, Health and Sustainability (ESHS-ESSF) Tables*. Tech. rep. Fab-blended UPW production baseline 18.1 kWh/m$^3$; supersedes academic-paper averages. URL: https://irds.ieee.org/images/files/pdf/2024/2024IRDS_ESHS-ESSF_Tables.xlsx.

Intergovernmental Panel on Climate Change (2021). *Climate Change 2021: The Physical Science Basis*. Tech. rep. Contribution of Working Group I to the Sixth Assessment Report. Cambridge University Press.

International Energy Agency (2025). *Energy and AI*. Tech. rep. Data centers accounted for around 1.5% of global electricity in 2024, projected to double by 2030. IEA. URL: https://www.iea.org/reports/energy-and-ai.

International Organization for Standardization (2006a). *Environmental management – Life cycle assessment – Principles and framework*. Geneva, Switzerland: ISO.

International Organization for Standardization (2006b). *Environmental management – Life cycle assessment – Requirements and guidelines*. Geneva, Switzerland: ISO.

International Organization for Standardization (2014). *Environmental management – Water footprint – Principles, requirements and guidelines*. Geneva, Switzerland: ISO.

International Organization for Standardization (2018). *Greenhouse gases – Carbon footprint of products – Requirements and guidelines for quantification*. Geneva, Switzerland: ISO.

JEDEC Solid State Technology Association (2025a). *2.5D/3D Silicon Thinned Wafers with Through-Silicon Vias (Cu TSVs) and Backside RDL: Reliability Guidelines*. Arlington, VA, USA: JEDEC.

JEDEC Solid State Technology Association (2025b). *Part Model Guidelines for Electronic-Device Packages — XML Requirements*. Arlington, VA, USA: JEDEC.

JEDEC Solid State Technology Association (2025c). *Thermal Characterization and Modeling of Compound Semiconductor FETs (III-V)*. Arlington, VA, USA: JEDEC.

JEDEC Solid State Technology Association (2026a). *3D Chip Stack with Through-Silicon Vias: Reliability Interactions Guidelines*. Arlington, VA, USA: JEDEC.

JEDEC Solid State Technology Association (2026b). *Catalog of Stress Procedures for Silicon Carbide Devices (PECS)*. Arlington, VA, USA: JEDEC.

JEDEC Solid State Technology Association (2026c). *Continuous-Switching Evaluation of GaN Power Conversion Devices*. Arlington, VA, USA: JEDEC.

King Yuan Electronics Co. (2019). *Corporate Social Responsibility Report (2018 and 2019 editions)*. Tech. rep. Plant electricity end-use disclosure: process equipment 44.8/45.9 percent, air-conditioning 35.1/33.5 percent, compressed air approximately 9 percent. King Yuan Electronics Co.

Krishnan, N., S. Boyd, A. Somani, S. Raoux, D. Clark, and D. Dornfeld (2008). "A Hybrid Life Cycle Inventory of Nano-Scale Semiconductor Manufacturing". In: *Environmental Science & Technology* 42.8, pp. 3069–3075. DOI: 10.1021/es071174k.

Murphy, C. F., G. A. Kenig, D. T. Allen, J.-P. Laurent, and D. E. Dyer (2003). “Development of Parametric Material, Energy, and Emission Inventories for Wafer Fabrication in the Semiconductor Industry”. In: *Environmental Science & Technology*. Vol. 37. 23, pp. 5373–5382. DOI: 10.1021/es034434g.

Pacific Gas and Electric Company (2010). *Cleanroom Baseline Energy Model*. Tech. rep. Per-ISO-class recirculation/FFU energy intensity by cleanliness class (ISO 3 = 366 W/m$^2$, ISO 5 = 161 W/m$^2$, ISO 7 = 54 W/m$^2$). Load-bearing source for REEL’s per-ISO-class zone intensities; cross-validated against Xu & Tschudi (2002), LBNL-49106. Pacific Gas and Electric Company. URL: https://www.pge.com/assets/pge/docs/save-energy-and-money/rebate-and-incentives/clean_room_baseline.pdf.

SEMI (2021). *Guide for Conservation of Energy, Utilities and Materials Used by Semiconductor Manufacturing Equipment*. Tech. rep. SEMI S23-1021. Energy Conversion Factor table for utility-to-electricity mapping. Baseline fab: 4,400 m$^2$ cleanroom, 10k 200mm WPM (per Hu et al. 2019 IJLCT documentation).

Sun, Y.-C. and F. H. Chang (2007). “Power consumption benchmark for a semiconductor cleanroom facility system”. In: *Energy and Buildings* 39.12. 8” DRAM benchmark: MAU 0.0042 kWh/m$^3$, general exhaust air 0.0007 kWh/m$^3$ (fan-only), recirculation 0.00018 kWh/m$^3$. ScienceDirect S0378778808000662., pp. 1290–1296. DOI: 10.1016/j.enbuild.2007.01.009.

TSMC (2022). *Fab 18 Phase Disclosure — 3nm Production Ceremony Press Release*. “Phases 1 through 8 of TSMC Fab 18 each have cleanroom area of 58,000 square meters, approximately double the size of a standard logic fab.” Primary anchor for REEL advanced_logic 1.50 m$^2$/WPM. URL: https://pr.tsmc.com/english/news/2986.

U.S. Environmental Protection Agency (2006). *Semiconductor Industry PFC/HFC Emissions: Review of U.S. and International Estimates*. Tech. rep. EPA 430-R-06-002. EPA.

Weidema, B. P., C. Bauer, R. Hischier, C. Mutel, T. Nemecek, J. Reinhard, C. O. Vadenbo, and G. Wernet (2013). “Overview and methodology: Data quality guideline for the ecoinvent database version 3”. In: *Ecoinvent Report* 1.v3.

Williams, E. D., R. U. Ayres, and M. Heller (2002). “The 1.7 Kilogram Microchip: Energy and Material Use in the Production of Semiconductor Devices”. In: *Environmental Science & Technology* 36.24, pp. 5504–5510. DOI: 10.1021/es025643o.

World Resources Institute and World Business Council for Sustainable Development (2011). *Corporate Value Chain (Scope 3) Accounting and Reporting Standard*. Tech. rep. GHG Protocol. URL: https://ghgprotocol.org/standards/scope-3-standard.

Xu, T. and W. F. Tschudi (2002). *Energy Performance of Cleanroom Environmental Systems*. Tech. rep. LBNL-49106, HT-447. Presented at ESTECH 2002 (report dated Nov 2001); administered by PG&E. Field measurements of 13 cleanrooms (ISO Class 4, 5, and above) across various industries. Reports recirculation-fan power intensity 16–38 W/ft$^2$ (ISO 4) and 3–24 W/ft$^2$ (ISO 5). Cross-validates REEL ISO 4–5 zone intensities only; contains no ISO 1–3 data. Lawrence Berkeley National Laboratory. URL: https://eta-publications.lbl.gov/sites/default/files/lbnl-49106.pdf.

Yin, J., X. Liu, Y. Guan, Y. Zhao, and T. Zhang (2020). “Performance and improvement of cleanroom environment control system related to cold-heat offset in clean semiconductor fabs”. In: *Energy and Buildings* 224. Indoor cooling load 457–484 W/m$^2$ constant across seasons; MAU heating reduction 52.8–94.3 % with heat recovery in transition/winter. Basis for REEL’s make-up-air split into a climate-insensitive baseline and a climate-modulated portion., p. 110243. DOI: 10.1016/j.enbuild.2020.110243.