Liquid Chemical Deposition (LCD)

A Non-Technical Primer on the Underlying Physics

Prepared for: Non-technical investors, board members, counsel, and general-audience readers interested in the underlying physics of the LCD platform

On behalf of: Cross Roads IP Holdings

Date: April 2026

> This primer is a companion to the LCD technical white paper. It is written for readers who want to understand, at a conceptual level, why LCD's manufacturing discipline produces the device-level "wow" that the white paper claims — better capacitors, better thermoelectrics, better magnetics — without requiring a physics or engineering background.
>
> We will keep the physics honest but presented at an intuitive level, with diagrams to carry most of the load. The goal is not to make you a physicist; it is to leave you with a clear mental model of what LCD actually does and why the same underlying process enables such different-sounding devices.

The foundational insight: why "getting the atoms right" matters

A reasonable first intuition about materials is that steel is steel, glass is glass, and the atoms are the atoms. That intuition works well enough for structural things. But for the materials we are talking about — high-performance capacitors, thermoelectrics, soft magnetics — the "magic" of the device comes from a very specific recipe of *which* atoms sit *where*, in *what ratios*. Shift the recipe by even a few percent and the material stops doing its job.

Take barium titanate ($BaTiO_3$), a classic high-performance capacitor material. One barium, one titanium, three oxygens. The titanium atom sits in a little cage of oxygen atoms, and when you apply an electric field, the titanium shifts slightly off-center inside that cage. That tiny off-center displacement is what makes the material polarizable — what makes it a good capacitor dielectric. If you have too much of a neighboring element mixed in instead of barium, the cage is the wrong size; the titanium cannot shift the same way; your capacitor's performance collapses.

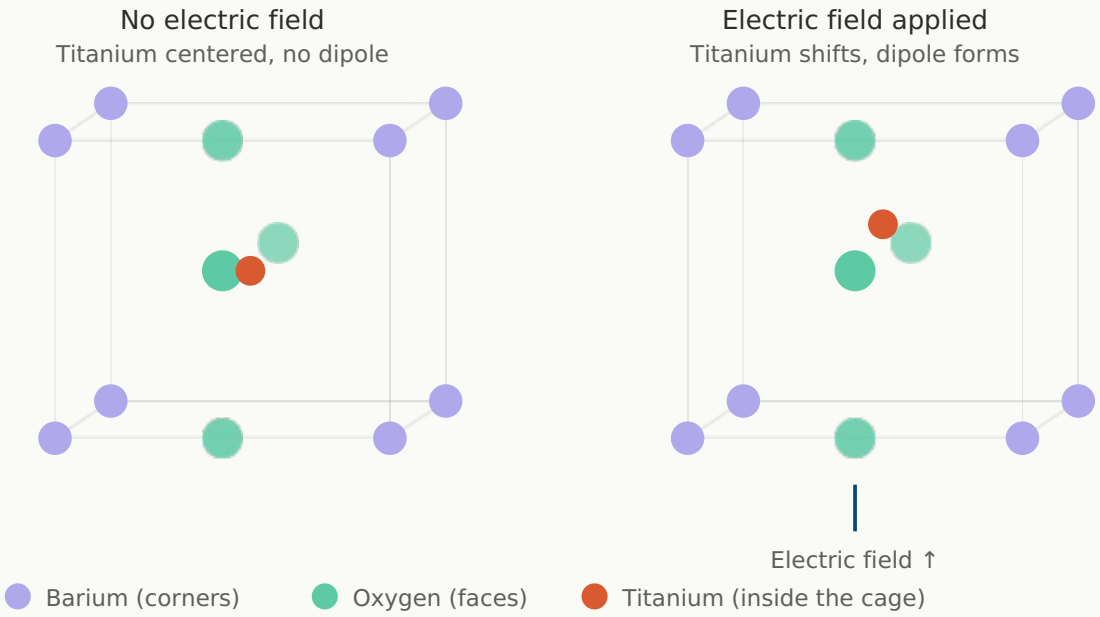

Figure 1. The titanium atom sits inside an oxygen cage. Under an electric field, it shifts off-center — the microscopic origin of polarization in a capacitor dielectric.

The crucial thing is that the off-center shift on the right is what stores energy. Billions of these little unit cells all shifting together is what makes the material a usable capacitor. If you got the recipe wrong — too much of one atom, too little of another — the cage is the wrong shape, and nothing shifts.

This is why LCD's central claim matters. When the process dissolves all the metal atoms in a single liquid beforehand, every droplet that hits the substrate carries the exact same recipe. By contrast, if you evaporate barium and titanium from two separate sources (the conventional vacuum approach), they travel separately, condense separately, and if one evaporates slightly faster on Tuesday than on Monday, your whole batch has the wrong ratio. In that world, getting the atoms right is fighting the process. In LCD, it is built in.

The hidden villain: grain boundaries

Now comes the twist that most people outside materials science never encounter. Even if you get every atom in the right place at the unit-cell level, real materials are not one giant perfect crystal. They are made of **grains** — little crystallites, each one a tiny well-ordered region, stuck together like crazy-paving. The places where grains meet are called **grain boundaries**, and those boundaries are where most materials go to die.

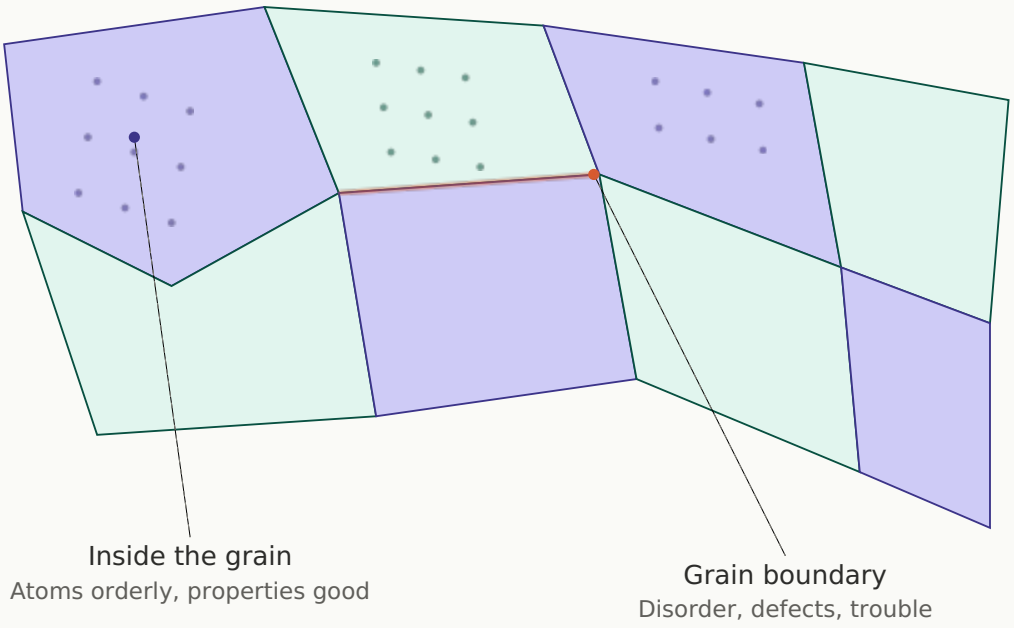

Figure 2. A polycrystalline ceramic. Inside each colored grain the atoms are well-ordered. At the boundaries between grains, disorder, defects, and impurities accumulate.

Each colored region is a grain — a little crystal where the atoms are lined up beautifully. But at the boundaries between grains, atoms do not know which crystal to belong to. They sit in awkward positions, form weird bonds, attract impurities, and create places where electrons can leak through, where cracks can start, where magnetic domains get stuck, where heat cannot flow.

Here is the LCD connection. A process that gets composition right but leaves the material full of ragged grain boundaries still produces a lousy device. What LCD's low-temperature-plus-stoichiometrically-correct combination enables is separate, controlled post-processing — you can cook the grain boundaries into submission (densify them, clean them up, align them) without disturbing the compositional recipe, because the recipe is already locked in.

With that foundation, we can now walk through each application. Each story has the same shape: get the atoms right at the unit-cell level (LCD's core trick), then engineer the grain structure (LCD's decoupled post-processing step). Different physics, same two-step.

Story 1: High-ε capacitors — why the cage has to be right

The capacitor story picks up where the barium titanate picture above left off. A capacitor's performance (the dielectric constant, written as ε) depends on how easily the titanium-in-the-cage shifts under an electric field. Three things can kill that performance:

1. **Wrong recipe.** Cage the wrong size, titanium cannot shift.

2. **Bad grain boundaries.** The disordered boundary regions have terrible dielectric properties. A capacitor is only as good as its weakest slice, so too many boundaries drags down the whole.
3. **Dielectric breakdown.** If there is a leaky path through the grain boundaries, electrons tunnel through and short-circuit the capacitor. This is where **quantum tunneling** becomes the enemy: at high field strengths, electrons do not need to climb over the energy barrier to cross the dielectric — they can tunnel *through* it. Grain boundaries provide shortcut paths.

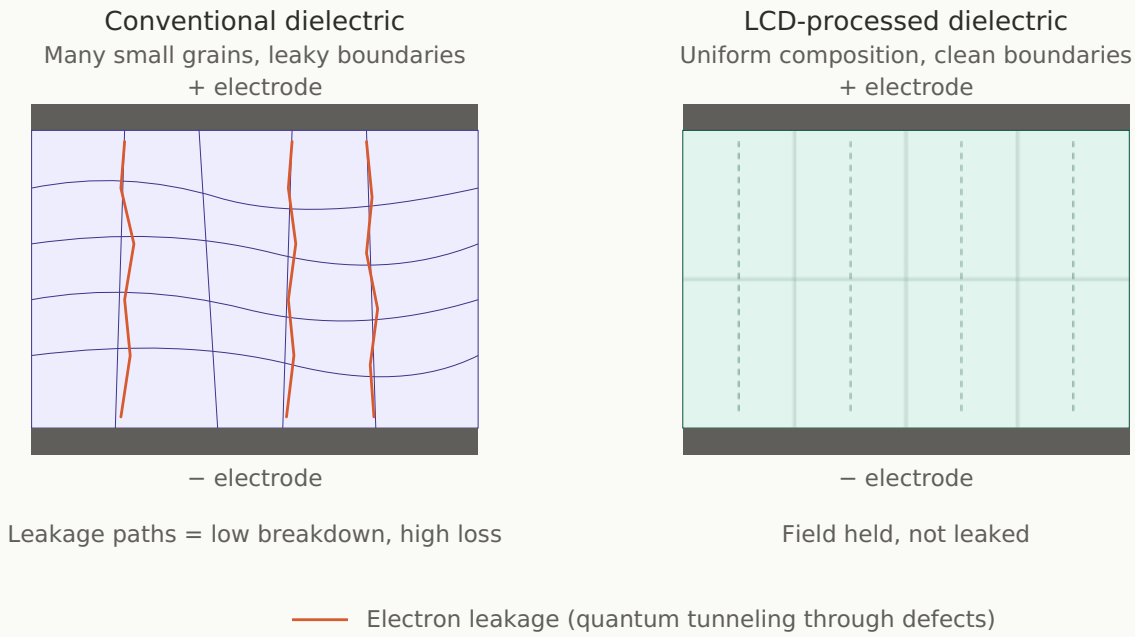

Figure 3. On the left, a conventional dielectric with many leakage paths through defective grain boundaries. On the right, an LCD-processed dielectric with uniform composition and clean boundaries.

The left-side picture is why conventional high-ε capacitors struggle. Each orange arrow is an electron quantum-tunneling from one side to the other through the defective grain boundary network. Enough of those and the capacitor leaks, heats up, and eventually breaks down. You can compensate by making the dielectric thicker — but then you lose capacitance density, because capacitance is inversely proportional to thickness. The engineer's impossible triangle: high ε, thin, and non-leaky — pick two.

LCD's contribution: stoichiometric exactness (correct cage, correct ε) *plus* low-temperature processing that lets you densify the grain boundaries without baking the film at temperatures that destroy the underlying silicon. That is why the qualification program targets high dielectric constant and high breakdown field at the same time. Those two targets together are what is hard.

Story 2: Thermoelectrics — the phonon-electron magic trick

Thermoelectrics are the most beautiful physics in this whole primer. The idea: put one side of a block of material hot, the other side cold, and — if the material is right — you get a voltage across it. Free electricity

from a temperature difference. Great for recovering waste heat from AI data centers, which is exactly what the qualification program is targeting.

The catch: you want the material to conduct **electricity well** (so the electrons move freely) and conduct **heat badly** (so the hot side stays hot and the cold side stays cold). In most materials, these two properties are coupled — good electrical conductors are usually good thermal conductors too, because the same free electrons carry both.

The trick, roughly: heat in solids is carried by two things — **electrons** (which we want) and **phonons** (lattice vibrations, which we *do not* want carrying heat). If you can engineer a material that blocks phonons while letting electrons through, you have a thermoelectric miracle.

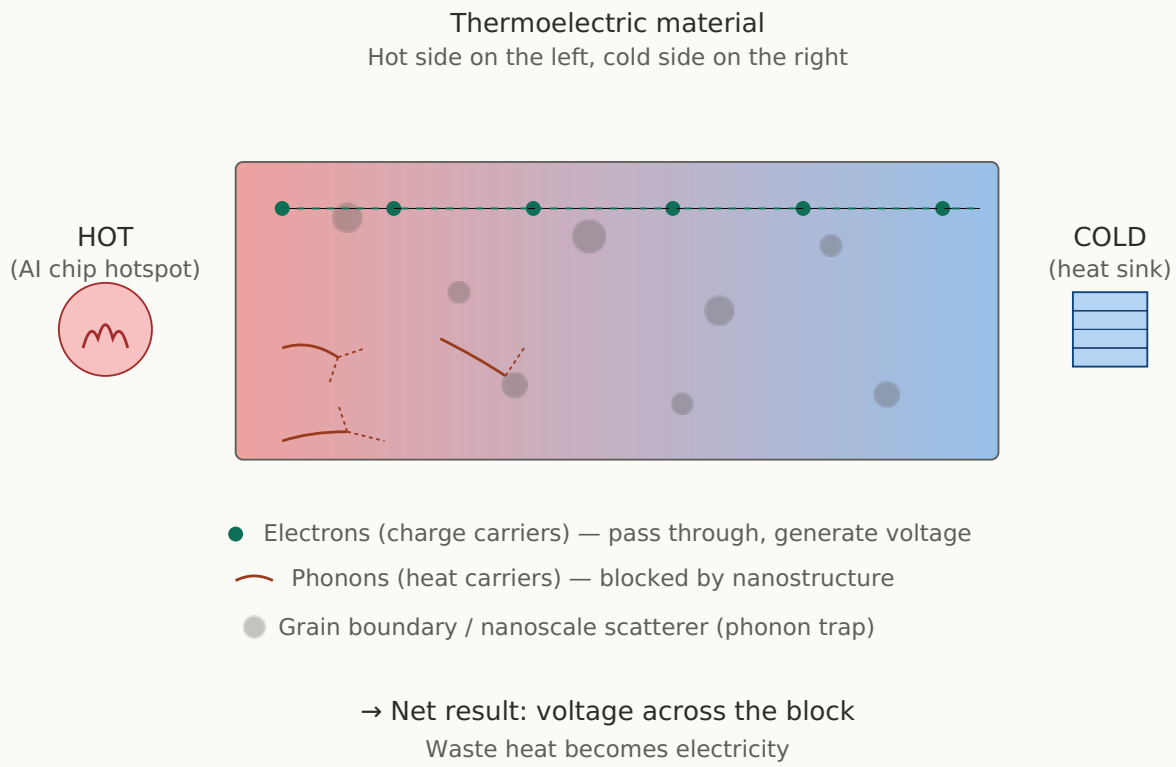

Figure 4. A well-engineered thermoelectric material. Electrons (green) pass smoothly through, carrying charge from hot to cold side. Phonons (orange) are scattered by nanoscale features, so heat cannot flow easily — the hot side stays hot.

The green dashed line shows electrons flowing smoothly from hot to cold side — they are small, they are charged, and electric fields guide them straight through. The orange wiggly paths are phonons — packets of lattice vibration, which is how heat physically moves through a solid. Phonons have wavelengths (typically a few nanometers to tens of nanometers in most ceramics), and they scatter off anything that is the same size as their wavelength.

This is where the LCD magic really kicks in. If you can engineer a material with exact stoichiometry (so the electron-conducting properties are optimal) *and* deliberately-placed nanoscale features (small grains, embedded nanoparticles, compositional variations) sized to scatter phonons but not electrons, you get a

thermoelectric with high performance. This is extremely hard to do with conventional routes, because they struggle to control composition and nanostructure independently. LCD does composition in the liquid phase and nanostructure in post-processing — the two-step decoupling from earlier — which is exactly what this application needs.

A small note on the gray dots in the picture: in modern thermoelectrics, those are often deliberately engineered nanostructures — little inclusions of a second phase, maybe 5–20 nanometers in size. LCD's ability to deposit multi-cation compositions means you can design the chemistry so these phases naturally separate out during post-processing annealing, at the right size, in the right density. That level of microstructural engineering is the frontier of thermoelectric research.

Story 3: Low-loss magnetics — the domain wall story

The magnetics story is different again, and it is delightful.

Inside a magnetic material like a ferrite, the atoms behave as tiny magnets. At any given moment, large regions of atoms have their tiny magnets all pointing the same direction — these regions are called **magnetic domains**. Between domains, there are **domain walls** — thin zones where the magnetic orientation gradually rotates from one domain's direction to the next.

When you put a magnetic field across a ferrite (as happens roughly ten million times per second in an inductor), these domain walls sweep back and forth across the material. Each sweep loses a little energy — to friction against defects, to eddy currents, to heat. That energy loss is what “core loss” means, and it is one of the primary specifications for an integrated magnetic component.

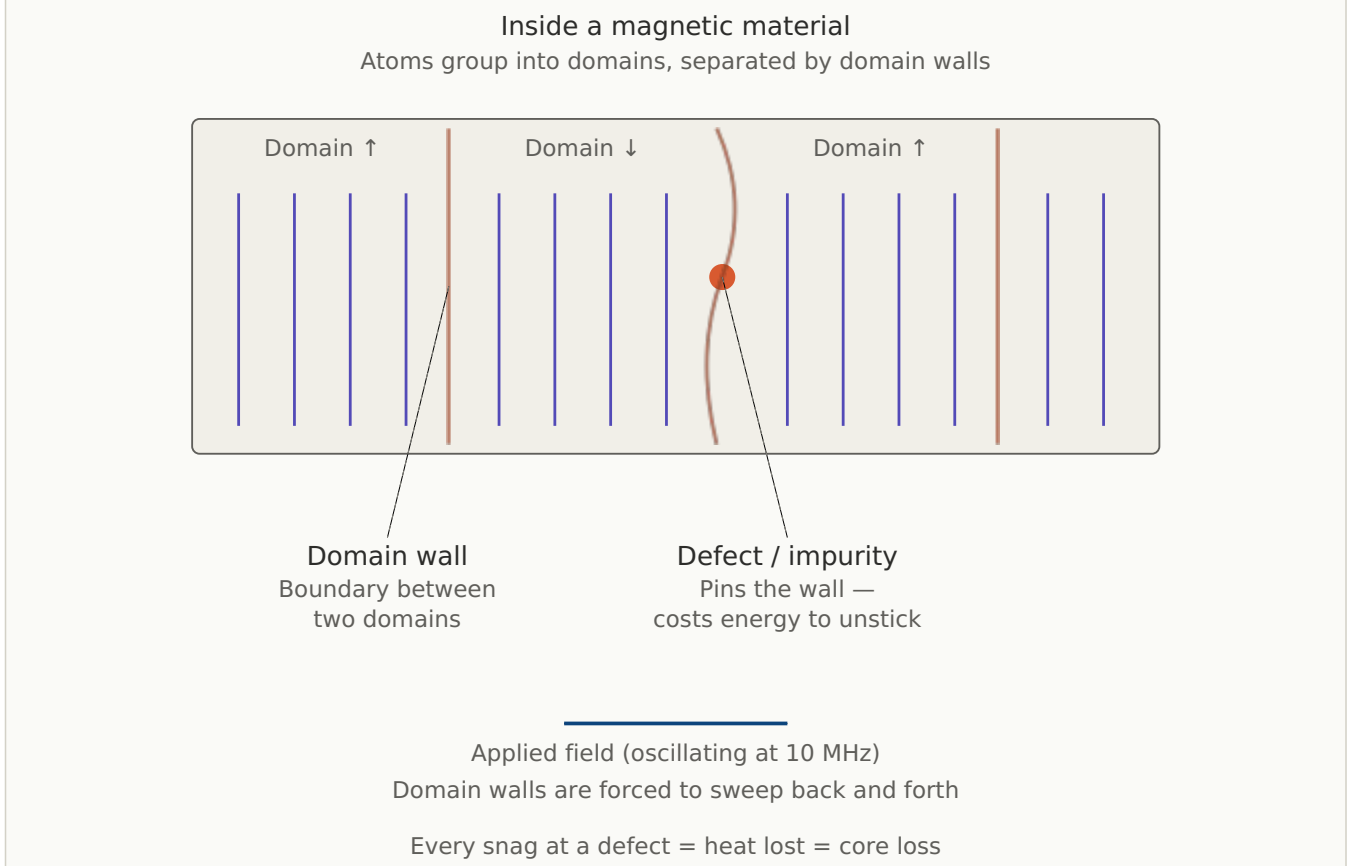

Figure 5. Inside a magnetic material. The material is divided into domains with their magnetic arrows pointing in different directions. The boundaries between them (domain walls) sweep under an applied field, and snag at defects — each snag dissipates energy as heat.

Look at the middle domain wall. It is bent around that orange defect — the defect is pinning it. When the external magnetic field reverses and the wall tries to sweep the other way, it has to physically unstick from the defect, which costs energy. That energy comes out as heat. Do that ten million times a second in a power-conversion circuit, and you have a hot, inefficient inductor.

Low-loss magnetics are fundamentally about making the material free of defects that pin domain walls. You want big clean grains, smooth grain boundaries, and few impurities. Which is exactly what LCD does well: composition correct at deposition (no stoichiometric impurities), and low-temperature post-processing that can grow clean grains without high-temperature sintering that introduces diffusion-based defects.

The qualification program targets high permeability (easy domain wall motion), high saturation (the material does not give up at high fields), and low core loss (clean wall motion) simultaneously. That combination is the magnetics engineer's version of the capacitor's impossible triangle. It is achievable only when you can engineer composition and microstructure as independent design variables. That is the LCD value proposition again, in a different guise.

The common thread

You have now seen three stories. Notice that they are structurally identical:

1. **Get the atoms right at the unit-cell level** — so the fundamental quantum-mechanical or electromagnetic property (dipole shift, electron transport, magnetic alignment) is available.
2. **Get the grain structure right** — so defects, leakage paths, and pinning sites do not destroy the property you just engineered.
3. **Do it at a low enough temperature** — so you have not cooked the chip, substrate, or surrounding layers.

LCD's unique claim is that it handles all three simultaneously — because it separates the composition problem (solved in the liquid phase, molecularly, at room temperature) from the microstructure problem (solved in post-processing, at modest temperatures). Every other route forces a trade-off among them.

That is why the physics community is excited about this platform, and why the same underlying technology enables capacitors, thermoelectrics, *and* magnetics — three very different physical phenomena, but all downstream of the same atoms-to-grains manufacturing discipline. The device-level “wow” you read about in the technical white paper is not a coincidence of three separate innovations; it is a single manufacturing insight expressed across three different physics problems.