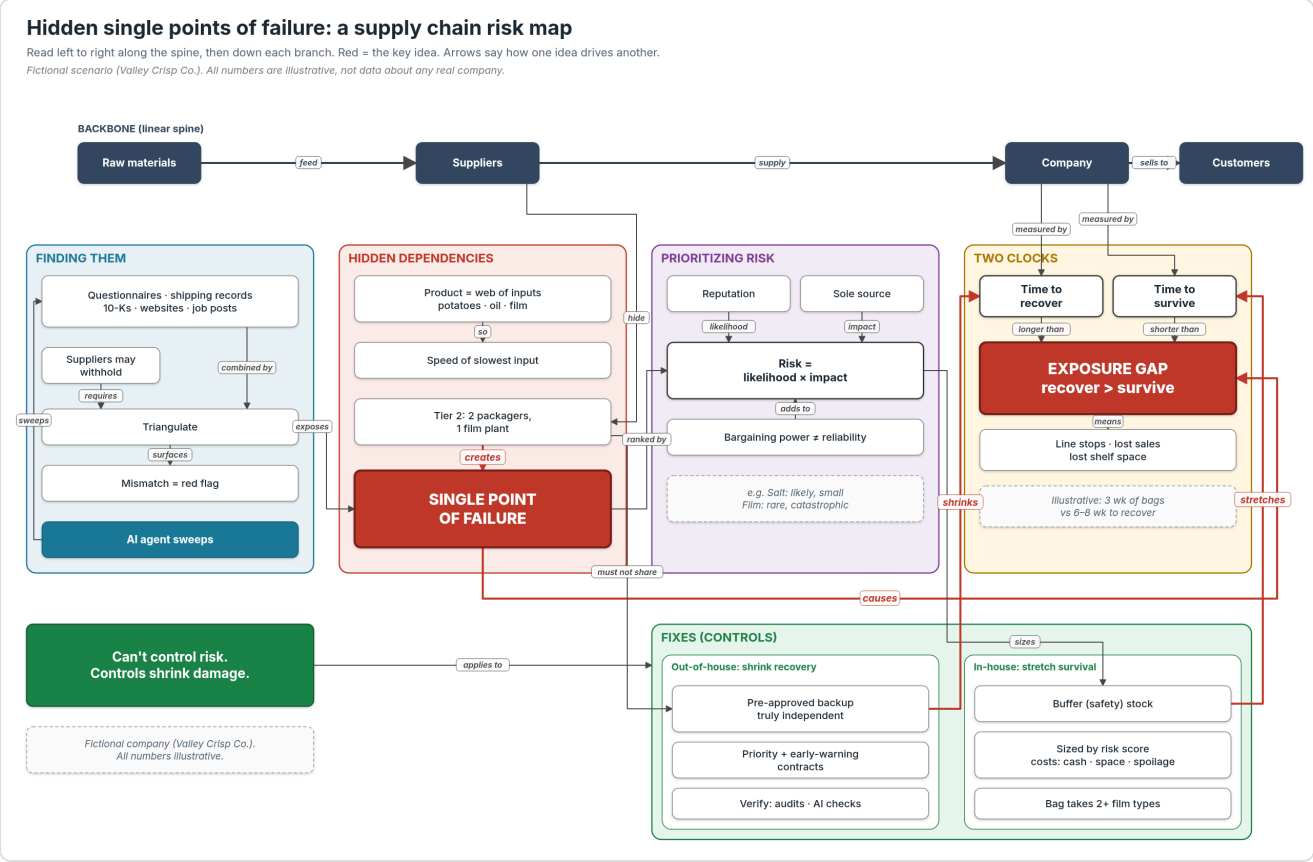
Two Suppliers, One Factory

How I find the hidden single points of failure in a supply chain, measure the damage in weeks, and where AI agents fit (and where they don't).



> **Hypothetical scenario.** Everything below uses **Valley Crisp Co., a fictional Shenandoah Valley chip maker**. The company, its suppliers, and every number are made up to illustrate the reasoning. Nothing here describes any real company's suppliers or operations.

A bag of chips is a web, not a line

Ask someone to draw a supply chain and you get a line: raw materials, suppliers, the company, customers. I started there too. Then I tried it on a bag of chips.

Valley Crisp needs potatoes, frying oil, and printed plastic bags, all at the same time. Two out of three doesn't make a product you can sell. So the chain is really a web of required inputs, and it moves at the speed of its slowest one. A late shipment of bags stops the line just as surely as a bad potato harvest.

The backup that wasn't

Say Valley Crisp buys bags from two packaging suppliers. On paper that's redundancy. If one fails, the other covers.

Now say both of those suppliers buy their plastic film from the same factory. The redundancy is gone. If that film plant goes down, both packaging suppliers fail together and the line stops anyway. The film plant is a **tier 2 supplier** (a supplier's supplier), and it's the real **single point of failure**. Valley Crisp never signed a contract with it, and it isn't on their vendor list. That's exactly why it's dangerous.

One more signal I'd check: competitors' sourcing. If every regional chip maker turns out to use the same film plant, it's an industry-wide bottleneck, and in a shortage everyone is fighting over the same output.

Finding what's hidden

Tier 2 dependencies don't announce themselves, so you go looking. The sources I'd use:

- **Supplier questionnaires.** The most common method: ask.
- **U.S. import shipping records.** These show who ships to whom, which is a direct trace for goods coming from overseas.
- **Supplier websites, press releases, and local news.** Companies like to announce partnerships and new plants.
- **10-K filings.** Public companies disclose material risks, and sometimes key suppliers.
- **Job postings and LinkedIn.** A supplier hiring for a specific material or machine tells you what it runs on.

Questionnaires have an incentive problem. A supplier has good reason to stay quiet about its own weak spots, and you lose time waiting on answers that may be incomplete. So I don't treat any one source as the truth. I **triangulate**: ask the supplier, then check the answer against filings and shipping records. Where they disagree, that's a red flag worth a follow-up call.

Every source has blind spots. Shipping records miss goods trucked domestically and importers who request confidentiality. Filings only exist for public companies. Triangulation narrows the gaps but doesn't close them, and I'd rather say that up front.

Prioritizing: likelihood × impact

You can't investigate every supplier deeply, so you rank them. My first instinct was reputation: on-time delivery history, quality complaints and recalls, what other buyers say, financial health. That's useful, but it only tells you how *likely* a failure is.

Risk is **likelihood × impact**. Compare two suppliers:

- A shaky salt supplier. Fairly likely to slip, but salt is easy to replace, so the impact is small.
- The sole film plant. Unlikely to fail in any given month, but if it does, production stops. Catastrophic.

The film plant is the bigger risk even though it looks more reliable. A dominant **single source** also brings a second problem, which is bargaining power. It can set prices and decide who gets product first in a shortage. Power over you isn't reliability, and it's a risk of its own.

Two clocks

"It would be bad" isn't a measurement, so I put impact in units of time:

- **Time to survive:** how long the business keeps running on what it already has.
- **Time to recover:** how long until supply flows again, whether the plant restarts or a new supplier gets qualified.

(Both terms are associated with David Simchi-Levi's supply chain risk research at MIT.)

Illustrative numbers for Valley Crisp: 3 weeks of bags in the warehouse, 6 weeks for the film plant to restart, 8 weeks to approve a new supplier. Even in the faster case, recovery outruns survival by 3 weeks. That window is the **exposure gap**: the line stops, sales are lost, workers sit idle, and competitors take the shelf space. Every fix I care about either stretches the first clock or shrinks the second.

Fixes: in-house and out-of-house

I group mitigations by who actually controls them.

In-house. Safety stock of bags stretches time to survive. It isn't free, since it ties up cash and warehouse space and some materials degrade on the shelf. So I size buffers by risk score: a deep buffer for the sole-source film, none for salt. Designing the bag to run on more than one film type is another in-house move that loosens the dependency.

Out-of-house. I can't run a supplier's processes, so I manage my exposure to them instead.

- *Reduce dependence.* Complete **supplier qualification** on a backup before you need it, ideally giving it a small share of regular orders so it can scale up fast. That shrinks time to recover. The catch from earlier still applies: the backup has to be truly independent. If it buys film from the same plant, it isn't a backup.
- *Put it in writing.* Priority-in-shortage clauses and early-warning requirements.
- *Verify, don't trust.* Audits, certifications, and public-data checks.

You can't control risk

The thing I had to unlearn was the language of control. A disruption is a probability. Valley Crisp controls its own processes, not its suppliers', and believing otherwise is an illusion of control. What a company *can* change is how much a failure hurts: stock levels, a ready backup, contract terms.

That's also what the word means in risk and compliance. A **control** is a noun, a specific safeguard like buffer stock or a pre-approved backup. It doesn't control the risk. It limits the damage.

How AI fits

Most of the work above is reading. One analyst can't comb shipping records, filings, websites, and job postings for hundreds of suppliers by hand. An AI agent can sweep them and flag when two "independent" suppliers point back to the same company.

What it shouldn't do is decide. Whether a match is real, how to score impact, how much buffer is worth the cash: those calls stay with a person. That's the design behind my open-source project, supply-web-audit, an 8-stage supply risk audit framework that AI agents walk one stage at a time, with a human reading and approving each stage before the next one starts. Its "narrow points" stage looks for single points of failure, and its "test moves" stage asks what happens if one fails and how long recovery takes. The reasoning in this piece is what that framework is built to automate. (Its current worked example is cyber risk; the layers are domain-neutral.)

What I'd build next

This is future work, not something that exists yet. An agent that takes a company's supplier list and:

1. Sweeps shipping records, filings, and supplier websites for each supplier.
2. Maps shared tier 2 dependencies across the list.
3. Flags where a supplier's questionnaire answers don't match the public evidence.
4. Hands a person a ranked list to review, with sources attached.

It would come with honest limits. It'll miss whatever never shows up in public records, and every flag needs a human to confirm it before anyone acts on it.

So here's the question I'd start any supplier review with: which of your "backup" suppliers buys from the same factory as your primary? If nobody knows, that's the first gap.

Disclaimer

Valley Crisp Co. is fictional. Its suppliers, inventory levels, restart times, and qualification times are hypothetical numbers chosen to illustrate the reasoning. This piece makes no claims about any real company's suppliers or operations, and it isn't legal, financial, or procurement advice.

Rashon Alsberry · Supply chain risk + AI automation