THE WON'T GET FOOLED AGAIN ACT

The Emergency Version: Frequently Asked Questions

How the plan works when the crash has already begun

November 2026 · Beta v0.2.0

Published by New Consensus · newconsensus.com

Contents

The Basics **3**

What is the emergency version? 3

What is the main difference from the plan-ahead version? 3

How is this different from what happened in 2008? 3

Does ordinary borrowing from the Federal Reserve count as asking for help? 3

Why does the rule apply to the Federal Reserve? 3

Would it have been better to pass this ahead of time? 4

How far along is this proposal? 4

Who Pays **4**

How is this different from a bailout? 4

What stops the government from rescuing a firm anyway? 4

Does any of the cost fall on taxpayers? 4

How much could it cost? 4

Will the levies raise prices for customers? 5

The Lenders **5**

Why do lenders receive anything when shareholders receive nothing? 5

Will this make it more expensive to borrow for AI infrastructure? 5

Could losses for lenders spread to the rest of the financial system? 5

How are retirees and pension funds protected? 6

Who gets paid ahead of the lenders? 6

Public Ownership **6**

Is public ownership like this unusual in the United States? 6

Who would run a converted bank or data center? 6

What protects the public institutions from politics? 7

Why does the plan apply only to companies that have failed or asked for help? 7

Why is some of the hardware sold to private buyers? 7

What are the Community Innovation Hubs? 7

What happens to the workers? 8

What does this mean for AI safety? 8

What happens to a failed lab's models and data? 8

Monopoly and the Chip Rules **8**

How does this keep the biggest technology companies from getting bigger? 8

How does this relate to antitrust? 8

What keeps the public utility from becoming too dominant? 9

What is the chip registry, and why is it needed? 9

Who has to register or get a license? 9

How do the chip rules stay fast and predictable? 9

How does this fit with existing export controls? 9

Deposits and Banks **10**

Why insure every deposit? ... 10
What did the FDIC conclude when it studied this? ... 10
Which banks are converted? ... 10
Who owns a converted bank? ... 10
What happens to my account if my bank is converted? ... 11
What would a Public Benefit Bank do differently? ... 11
Markets, Innovation, and China ... 11
What does this mean for AI innovation? ... 11
What does this mean for competition with China? ... 11
Why isn't ordinary bankruptcy enough for these companies? ... 11
How does the plan avoid picking winners and losers? ... 12
The Law and the Limits on Government ... 12
What is the constitutional basis for converting a failed company? ... 12
Is the price the public pays fair to owners and lenders? ... 12
Who decides that a company is "systemically important"? ... 12
Who decides that a company has failed? ... 13
Does the chip registry give the government control over computing? ... 13
What keeps the agency's role from expanding over time? ... 13
How is this different from the resolution authority in Dodd-Frank? ... 13
Acting in a Crisis ... 13
Is it fair to change the rules after lenders have already made their loans? ... 13
How can the government bar the big technology companies from buying assets if there is no chip registry yet? ... 14
Why do the rules take effect when the bill is introduced, and not when it is signed? ... 14
Can sales that have already happened really be undone? ... 14
What happens to companies that are already in bankruptcy court? ... 14
Why does the FDIC handle failed technology companies? ... 14
How do you keep a company in receivership running when nothing was prepared? ... 14
Why does waiting cost more? ... 15
Why is this a full-length plan and not a short emergency bill? ... 15
Politics ... 15
What would it take to pass this? ... 15
Who is this for? ... 15
What happens after the emergency is over? ... 15

The Basics

What is the emergency version?

It is the Won't Get Fooled Again Act rewritten for a crash that has already begun. The Act was designed to be passed ahead of time, so that the funds, the rules, and the preparations would be in place when institutions started to fail. This version is for the case where that did not happen. The goals are the same: convert what fails into public institutions, have the people who financed the bubble bear its cost, protect workers and depositors, and keep the largest firms from buying up the wreckage.

What is the main difference from the plan-ahead version?

Preparation. The central rule is the same in both versions: any bank, and any systemically important technology company, that needs public money to survive has failed and receives help only on the terms of this Act. What the plan-ahead version has, and this one does not, is time: years of premiums in the funds, a registry of who owns the chips, firms that have filed plans for their own conversion, and lenders who knew the rules when they lent. So in this version the money is borrowed first, the FDIC acts as receiver from the first day, the rules take effect when the bill is introduced, and the aid rule does most of the work, because in a crisis the important firms do not wait to fail. They ask for help.

How is this different from what happened in 2008?

In 2008 the Treasury bought shares in the banks with few conditions. It took no control, left management and shareholders in place, and bondholders were paid in full. The largest failures were folded into even larger banks with government support. Under this Act, public money buys public ownership, the people who financed the bubble take the losses, and the largest firms are barred from acquiring the failures. There is an older American precedent for the tougher approach: in 1933 the Reconstruction Finance Corporation recapitalized thousands of banks by taking shares with voting rights, and used them to replace managers and cap salaries.

Does ordinary borrowing from the Federal Reserve count as asking for help?

No. A sound bank borrowing from the Fed against good collateral, at the posted rate, has not failed. The test is whether the support is something a solvent firm could get on the same terms. A capital injection, a guarantee of a firm's debts, a loan on terms the market would not offer, or a purchase of its assets above their market value all count as aid.

Why does the rule apply to the Federal Reserve?

Because the Fed's emergency lending is the quiet route by which losses can be moved onto the public without a vote of Congress. Under the Act the Fed keeps its proper job, which is lending to solvent institutions against good collateral through broadly available programs. It may not create a program whose purpose or effect is to support the debts of failing AI companies, or the funds that lent to them, outside the terms of the Act.

Would it have been better to pass this ahead of time?

Yes, and the plan-ahead version explains how. Passed in advance, the industry pays for its own insurance, lenders know the rules before they lend, a registry shows who owns what, and conversions run on rails. This version exists because the choice in a crisis is not between this plan and a better-prepared one. It is between this plan and another bailout.

How far along is this proposal?

It is a beta. The mechanisms are worked out, but many of the specific numbers are our best current judgment and are marked for further work, and the legal design needs review by bankruptcy and constitutional lawyers. The rule that aid means conversion, and the limits on the Federal Reserve, need the closest look.

Who Pays

How is this different from a bailout?

A bailout rescues the people who made the bets. It keeps shareholders in place, makes lenders whole, and usually leaves management in their jobs. The Act does the opposite on all three counts. Shareholders lose everything. Lenders get the current value of their collateral and lose the rest. Executives are removed and their pay is clawed back. What is rescued is the service, the jobs, and the customers, not the investors.

What stops the government from rescuing a firm anyway?

The Act itself. Any bank, and any systemically important technology company, that takes emergency federal support is treated as having failed: its shareholders are wiped out, its lenders are paid what their collateral is worth, and it is converted. The rule binds every federal agency and the Federal Reserve by name. Ordinary lending to sound banks against good collateral is not affected.

Does any of the cost fall on taxpayers?

In a crisis the public has to advance the money, because there are no accumulated premiums to draw on. The emergency version sets up a $250 billion Treasury credit line to pay for resolutions up front, and lets the bank fund borrow up to $75 billion more to capitalize converted banks. The industry repays it over 10 years or more through a tax on high-end chips and large training runs, premiums from the largest technology companies, and a surcharge on the largest banks, and the rates rise automatically if repayment falls behind. So public money is at risk for longer than it would have been under the plan-ahead version, and we say so plainly. What the public buys, it buys at the market's price on the day of failure, and owns.

How much could it cost?

Nobody can know in advance. Our rough estimate is that in a serious crash the debts of failed AI-related companies could come to $150 billion to $285 billion. Under the Act the public does not pay those debts.

It pays what the underlying assets are worth and receives assets of that value. If used chips lose most of their value in the crash, that might be $50 billion to $100 billion. If they hold their value better, it could be roughly double. The public cannot choose the lower figure: it pays what the market shows the assets to be worth. For scale, Congress authorized $700 billion for the 2008 bank rescue, later reduced to $475 billion. These figures are estimates and the full plan marks them as such. On the banking side, the FDIC's records give the scale: the failures of 2023 would have needed about $43 billion to $53 billion of capital, and those of 2008 through 2013 about $55 billion to $69 billion over six years. That capital stays inside working banks, and if a giant fails, its own loss-absorbing debt comes first.

Will the levies raise prices for customers?

Partly, yes. That is true of deposit insurance premiums too, and it is how insurance is supposed to work: the price of a risky activity should include the cost of the risk. Small-scale research and educational computing are exempt from the compute tax. And the alternative is not free. It is a bailout paid for by everyone, including people who never used the product.

The Lenders

Why do lenders receive anything when shareholders receive nothing?

Two reasons. The first is legal. A lender with a lien holds a property right, and the Constitution does not let the government take property without paying for it. The second is that the law asks for less than people assume. What a secured lender is owed is the value of the collateral on the day of failure, not the amount of the loan. A fund that lent $10 billion against chips now worth $3 billion gets $3 billion. The other $7 billion becomes an unsecured claim that will recover little or nothing. That is a real loss, and it is the point.

Will this make it more expensive to borrow for AI infrastructure?

Yes, somewhat. Loans for speculative build-outs will carry higher interest and stricter terms once lenders know that nobody will make them whole. That is intended. Cheap credit with an implied government guarantee is how bubbles are inflated. Companies with real revenue will still be able to borrow.

Could losses for lenders spread to the rest of the financial system?

This is a more serious concern in a crisis than it would have been in advance. If the rules had been in place for years, lenders would have priced in their losses gradually. Imposed in the middle of a panic, the same losses arrive as a shock. The emergency version handles this by protecting the channels the whole economy depends on, not the bets. Every bank deposit is guaranteed at once. If banks lose access to short-term borrowing, the FDIC can guarantee their new borrowing for a fee, as it did in 2008. A bank that fails because of its losses is converted. Pensioners and insurance policyholders are protected directly. What the Act does not do is guarantee the bonds and loans investors already hold. Ireland did that in 2008, and it nearly bankrupted the country.

How are retirees and pension funds protected?

This is one of the most important questions, and the plan addresses it directly. Pension funds rarely lend directly. They invest through private credit funds that also hold money from sovereign wealth funds, wealthy individuals, and insurers, so rescuing the loans rescues everyone. Instead, the Act creates a Pension Protection Facility, paid for by the industry levies, that makes long-term loans or grants to pension plans with documented losses. It protects retirees' benefits, not fund balances, and comes with conditions on future risk-taking. It is aimed mainly at state and local plans, which have no federal insurance. Congress did something similar for failing union pension plans in 2021.

Who gets paid ahead of the lenders?

A lender with collateral is entitled to that collateral's value, because a claim on collateral is a property right that the Constitution protects. It receives that value as the market sets it on the day of failure, which may be far less than it lent, and nothing more. The government does not cover lenders' losses. For everything else a failed firm owes, the Act sets the order. Employees come first, for everything they have earned. Then small suppliers, up to a cap. Customers, whose contracts and prepaid credits carry over to the new public entity. Then local governments and utilities. Banks, bondholders, and investment funds come after all of them for whatever their collateral did not cover. In a crisis the order cannot be written before the loans were made. It is written into the law before any firm is resolved, so that the same order applies to every failure and nobody can accuse the government of picking favorites after the fact.

Public Ownership

Is public ownership like this unusual in the United States?

No. The plan calls for public ownership of institutions that have failed and that the country cannot do without, and the United States has done this many times: the Tennessee Valley Authority, the conversion of the bankrupt northeastern railroads into Conrail, the Bank of North Dakota, which has operated since 1919, and every weekend that the FDIC takes over a failed bank. Ownership by depositors is older still: mutual savings banks and credit unions have belonged to the people who bank with them for generations. The aim is a system in which the people who take the gains also take the losses.

Who would run a converted bank or data center?

It would not be run by a federal office. A converted bank keeps its staff and branches. A local or regional bank is governed by a board its depositors elect or, where a state or city takes ownership, by a board drawn from the places it serves, on the model of Germany's public savings banks, which kept lending through 2008 while private banks pulled back. A large national bank gets a board accountable to the public at large, with seats for its employees, its customers, and the regions where it operates. A converted technology firm keeps its engineers. In a crisis there is no roster of outside chief executives waiting, as the FDIC has for banks, so the receiver relies on the firm's own people and on the Department of Energy's national

laboratories, which already operate the largest computers in the country and can supply experienced managers. Public enterprises can be run well or badly, as private ones can. The plan keeps the people who know the operation and changes who they answer to.

What protects the public institutions from politics?

This is a real risk, and the plan is designed around it. Public Benefit Banks are barred from speculation and limited to ordinary lending. Most converted banks belong to their depositors, who elect the board, so no official appoints it and no administration can remove it. Where boards are publicly appointed, at the largest banks and the technology institutions, they mix public appointees, employees, customers, and technical experts, so that no single interest controls them. The computing utility must offer non-discriminatory access at regulated prices and publish what it does. No design makes political influence impossible. Since the Supreme Court's 2026 ruling in Trump v. Slaughter, Congress may not be able to stop a President from replacing the directors of the federal institutions, so the plan does not rely on their tenure. It relies on what the law fixes: narrow mandates, a bar on privatizing without an Act of Congress, ordinary supervision, and public reporting. A hostile administration could still run these institutions badly. The plan is written for one that wants them to work.

Why does the plan apply only to companies that have failed or asked for help?

Some people we respect would go further and take the banks and the AI labs into public ownership outright. We limit conversion to firms that have failed, or that need public money to survive, for three reasons. It is on firm constitutional ground, because the owners of a failed firm have already lost their investment, and a firm that accepts terms in exchange for help has made a bargain. It is fair to the public, which pays only what the assets are worth on the day of failure and owns what it pays for. And a crisis is the one moment when large changes can be made. A healthy company that asks for nothing is left alone.

Why is some of the hardware sold to private buyers?

There is no ceiling on what the public may take, and it gets the first pick of everything. But a crash means there are more chips than customers, and idle machines burn electricity and age quickly. What the public cannot put to use within 30 days must be offered to smaller private buyers. Cheap hardware in the hands of mid-size companies, universities, and startups builds competitors to the dominant firms, which serves the same goal as the public utility.

What are the Community Innovation Hubs?

They are buildings: community workspaces where residents who meet simple criteria get a free membership and can come to work, take classes, get advice, and meet other people who are starting things. Places like this already exist across the country. The efactory in Springfield, Missouri, is one example, and its clients have created roughly 3,000 jobs. The Act funds many more of them, and funds existing ones to grow, with at least one in every congressional district. They are paid for by a 1% fee on the largest AI and cloud companies, on the model of the fees that cable companies paid to fund public access television. They start when the Act passes and do not depend on a crash or on any public computing facility.

What happens to the workers?

They keep their jobs. Converted institutions are prohibited from mass layoffs, workers hold seats on the boards, and wages and benefits owed are paid first. Because much of the pay at AI companies is stock that a failure makes worthless, the agency pays retention bonuses from the first day so that the people who keep services running do not leave.

What does this mean for AI safety?

A public institute is free of the pressure to ship products before they are tested, and the plan directs it toward safety and interpretability research. That is an advantage, not a guarantee. The Act is a plan for resolving failures. It is not a safety regime for what AI systems do, and it says so. That is a separate and larger question that deserves its own law.

What happens to a failed lab's models and data?

They are frozen the moment the receiver is appointed and kept in public hands. The plan gives access to academic and nonprofit researchers under strict safety protocols. It does not publish model weights to the world. The freeze also means a failed firm's models and data cannot be deleted, sold to data brokers, or bought by a foreign government in a bankruptcy auction, which is what could happen today.

Monopoly and the Chip Rules

How does this keep the biggest technology companies from getting bigger?

In a crash, the natural buyers of failed companies' chips and data centers are the three or four firms that already dominate computing. The Act bars any company that controls more than 5% of the nation's AI computing capacity from buying those assets out of a bankruptcy or foreclosure. It closes the two obvious ways around that: such a company may not lease chips that someone else bought in a distressed sale, and may not buy a firm that has warned regulators it is close to failing. The rule is not a limit on size. A capped company can still buy as many new chips as it likes. The public gets the first option on the distressed assets, and everything else goes to smaller buyers. The banking equivalent already exists, a rule barring any bank from buying another if that would leave it with more than 10% of the nation's deposits, but it has an exception for buying failed banks, which is exactly when it matters. Ours has no exception.

How does this relate to antitrust?

The two work together, and nothing here stands in the way of antitrust enforcement. This Act does a narrower job. It stops a crisis from concentrating the industry further, and it builds a public competitor. Antitrust cases take years. A fire sale takes weeks.

What keeps the public utility from becoming too dominant?

It must offer access to everyone on equal terms at regulated prices, which is what separates a utility from a monopolist. The rule against hoarding forces idle capacity back onto the market. And it will be competing with some of the largest and best-funded companies in the world, which are left intact.

What is the chip registry, and why is it needed?

It is a record of who owns each large cluster of advanced AI chips, where it is, and who has lent against it. The model is the federal aircraft registry, which has recorded the owners of and the liens on every plane in the country for decades without slowing the sale of one. Without it the government cannot know who controls the nation's computing power, cannot enforce the 5% rule, and cannot tell where chips go when a failed company is liquidated. Today, once a chip is sold inside the United States, no law requires anyone to know where it is.

Who has to register or get a license?

Only the owners and operators of very large clusters: 1,000 or more advanced AI accelerators, defined by computing power rather than by maker, so that chips a company designs for itself count the same as chips it buys. Below that line nothing applies. Researchers and small companies need no permission to do their work, and gaming cards, workstations, university labs, and small company clusters are never touched. Above it, the license must be issued to anyone who meets published security standards, and the agency does not review what is run on the chips.

How do the chip rules stay fast and predictable?

The rules are designed to be simple. They are bright lines published in advance. Transfers go through automatically after 30 days unless the agency objects in writing. The agency is funded by the compute levy, not by hourly fees, so it gains nothing from a long review. Its mandate includes widening access to computing, not only preventing harm. Real discretion exists in one situation only, the sale of a failed firm's assets.

How does this fit with existing export controls?

The two do different jobs. Export controls already stop legal sales of advanced chips to China and other countries of concern, and the Act adds no new export rules. What export controls do not stop is smuggling, which already happens through shell companies, false paperwork, and shipment through third countries. A fire sale, with thousands of chips sold in small lots by liquidators who only want cash, would be the ideal market for it. Under the Act, chips from registered clusters are tracked by serial number through every resale. A paper trail will not stop a determined smuggler, but it will show whose chips went missing, which today nobody could say.

Deposits and Banks

Why insure every deposit?

The main reason is that a limit on deposit insurance does not work the way it is supposed to. Everyone knows the government will never let depositors at the largest banks lose money. So at the first sign of trouble, large depositors move their money from smaller banks to the giants. In the week after Silicon Valley Bank failed, about $120 billion left smaller banks, and tens of billions of dollars arrived at the largest ones. A limit therefore rescues the biggest banks by accident, whatever their condition, and drains their sounder competitors. Insuring every deposit at every bank ends that. It also means a converted public bank keeps its depositors, which it needs in order to function.

On risk-taking: a common concern is that full insurance removes depositors' reason to watch their bank. In practice depositors have not been good at policing banks. Silicon Valley Bank's customers included some of the most sophisticated investors in the country, and they noticed nothing until they all ran on the same day. Under the Act, discipline falls on shareholders, bondholders, and executives, who are paid to bear risk and lose everything when a bank fails. Banks pay risk-based premiums for the new coverage, and the largest pay the most.

What did the FDIC conclude when it studied this?

After the 2023 failures the FDIC studied three options. It found that unlimited coverage would deal most directly with bank runs, but it preferred a narrower option, covering business payment accounts, because of the cost to the insurance fund and the effect on risk-taking. We read the same evidence and come down differently, for the reason given above: anything short of full coverage keeps pushing money toward the largest banks. We agree with the FDIC that full coverage has to be paid for, and the plan has banks pay for it.

Which banks are converted?

Every failed bank that is still a functioning bank, whatever its size. Converting only the largest banks would leave small and mid-size failures to be sold in the usual way, usually to a bigger bank. A small town's bank matters as much to its customers as a giant does to the country, and the giants should not get better treatment. The FDIC can decline to convert a bank only by finding, in writing and in public, that it has no real base of customer deposits, no real lending business, or failed mainly because of fraud. In that case its accounts and sound loans go to a nearby converted bank, credit union, or community bank. They are never sold to one of the largest banks.

Who owns a converted bank?

It depends on the bank's size. A local or regional bank becomes the property of its depositors: its business moves into a newly chartered mutual bank, each depositor has one vote, and the depositors elect the board. Mutual savings banks and credit unions have worked this way in America for generations. A state or a city can take ownership instead, as North Dakota owns its bank, or its accounts and loans can go to

a local credit union. Only the largest banks, those with more than $250 billion in assets, become federal public corporations, because no single community can speak for a bank with customers in every state. In every case there are no shareholders.

What happens to my account if my bank is converted?

Nothing you would notice. The FDIC already resolves failed banks over a weekend. On Monday your account, your card, and your loan work as they did on Friday, and the bank keeps its name. What changes over time is what the bank does with its money. And if it is a local or regional bank, you become one of its owners, with a vote for its board.

What would a Public Benefit Bank do differently?

It takes deposits and makes loans to local businesses, housing, and infrastructure. It is barred from proprietary trading, derivatives speculation, and lending to shadow banks. Its leadership answers to its depositors and the communities it serves, not to shareholders. The Bank of North Dakota has worked this way for more than a century, has never needed a rescue, earns a higher return on its assets than Goldman Sachs, and returns its profits to the state.

Markets, Innovation, and China

What does this mean for AI innovation?

Nothing in the Act limits what healthy companies may research or sell. Speculative borrowing will cost more, so some projects that depend on cheap credit may not be built. After a crash, the choice is not between this plan and a thriving private industry. It is between this plan and a fire sale in which the technology is broken up, shut down, or absorbed by a few giants. The Act keeps the researchers employed, the machines running, and access to computing power wider than it was before.

What does this mean for competition with China?

In a crash, the danger to the American position is that capacity goes dark, research teams scatter, and models or chips are sold to the highest bidder, who may be foreign. The Act prevents each of those. It keeps capacity running, retains the staff, freezes models and data against sale, puts every distressed sale through national security review, and tracks chips through resale. Keeping that infrastructure intact through a downturn is what protects the country's position.

Why isn't ordinary bankruptcy enough for these companies?

For most companies that is the right answer, and the Act leaves it alone. Bankruptcy has one job, which is to get creditors as much money as possible. For a few institutions that is the wrong goal. A liquidation trustee has no duty to keep a hospital's AI systems running, no reason to care whether model weights go to a foreign buyer or customer data to a broker, and every reason to sell the hardware to whichever giant

bids most. Congress has long recognized this. It is why banks do not go through ordinary bankruptcy either.

How does the plan avoid picking winners and losers?

The rules are written into the bill, and they are the same for every firm. The thresholds for which firms count as systemically important are numbers in the statute, not names. The order in which creditors are paid is fixed in the statute. The price the public pays is set by open bidding, with a published formula as the floor. What the plan rules out is deciding, one firm at a time, which creditors should be saved.

The Law and the Limits on Government

What is the constitutional basis for converting a failed company?

Congress's power to write special rules for failed institutions, which it has used for banks for 90 years. Converting a failed company is constitutional under four conditions, and the plan is built to meet them. The firm must actually have failed. Secured lenders must receive the value of their collateral. Every creditor must get at least what a liquidation would have paid. And courts must be able to review both the receivership and the price. In the 1970s Congress transferred the assets of the bankrupt northeastern railroads to a government-created company, and the Supreme Court upheld it. The plan also deliberately does not rely on the Defense Production Act. When President Truman seized the steel mills without clear authority from Congress, the Court reversed him. The power here comes from the statute itself.

Is the price the public pays fair to owners and lenders?

The owners of a failed company hold shares that the market has already made worthless. Nothing is taken from them. Lenders are paid what their collateral would actually fetch: the assets are offered to every eligible buyer, the public can take them only by matching the best bid, and it never pays less than a published formula price. Even that is not the last word: a creditor who believes a liquidation would have paid more can go to court, and the courts decide. The railroad precedent shows this is a real right. After years of litigation the government paid the Penn Central estate about $2.1 billion.

Who decides that a company is "systemically important"?

In a crisis there is no time for a council to build a record on each firm, so the bill does it directly. It sets numeric thresholds: a share of the nation's computing capacity, revenue and user counts, the number of hospitals, utilities, and agencies that depend on the firm, and how much it owes to the financial system. A firm that meets any one of them is covered. The FDIC and the Treasury apply the thresholds within days and publish the list. A firm can contest its place on the list in court afterward, but that does not delay anything. A permanent council takes over once the emergency has passed.

Who decides that a company has failed?

For banks, nothing changes: the bank's regulator closes it and appoints the FDIC, as happens today. For a technology company, three events count as failure automatically, with no vote: it files for bankruptcy, it leaves a payment of $100 million or more unpaid when the grace period in its own contract runs out (30 days at most), or it takes emergency federal support. Regulators can also act earlier, when a firm is about to fail but has not yet crossed one of those lines. That is a judgment, so no single official can make it. It takes a written recommendation from two-thirds of the FDIC's board and two-thirds of the Federal Reserve Board, and a decision by the Treasury Secretary. Either way the company can go to court, which has 24 hours to rule. If it does not, the receivership goes ahead.

Does the chip registry give the government control over computing?

It records who owns very large clusters of export-controlled chips and who has lent against them. It does not record what anyone computes, and the agency has no authority to review it. Ownership of aircraft, ships, broadcast licenses, and nuclear facilities has been recorded this way for generations. We would rather have a narrow rule in place, with bright lines and hard limits on discretion, than have Congress write a sweeping one in the middle of a panic.

What keeps the agency's role from expanding over time?

That is a reasonable concern, and the plan includes limits for it. The conversion power applies only to failed firms, with a court looking over the agency's shoulder. The registry has a fixed size threshold. Transfers clear automatically on a deadline. The agency has no fee income to protect. Beyond that, the honest answer is that every law depends on Congress continuing to watch how it is used.

How is this different from the resolution authority in Dodd-Frank?

It borrows from it: the 24-hour court review, the rule that shareholders and creditors bear the losses, and the Treasury credit line repaid by the industry. There are two differences. Dodd-Frank covers only financial companies, and this extends to the technology companies the economy now depends on. And Dodd-Frank aims to wind a failed firm down and sell the pieces. This Act keeps the institution and converts it to public purpose. Dodd-Frank's authority has gone unused in part because regulators in a crisis reach for a rescue or a merger instead. This Act takes those options away.

Acting in a Crisis

Is it fair to change the rules after lenders have already made their loans?

The constitutional protections do not change: a secured lender must receive the value of its collateral, and every creditor must receive at least what a liquidation would have paid. In a collapsing market both amounts are low, and the Act pays them. It is true that lenders who lent before the Act existed have a stronger argument than they would have had if it had been passed in advance, and some of their lawsuits will be close. The Act prepares for that. Wherever it can, it works through conditions that firms accept in

exchange for help, which are much harder to challenge than terms imposed on them. Every claim goes to a single court, and the remedy is money, so no lawsuit can stop a conversion. And the credit line is sized on the assumption that some cases will be lost.

How can the government bar the big technology companies from buying assets if there is no chip registry yet?

Until a registry can measure market shares, the barred buyers are set by a formula anyone can apply: the five largest providers of cloud computing in the United States by revenue, and any company they control. An emergency census requires every owner of 1,000 or more advanced AI chips to report what it holds within 14 days, which becomes the first edition of the registry. When the registry publishes its first measurement, the formula lapses and the 5% rule takes over.

Why do the rules take effect when the bill is introduced, and not when it is signed?

Because otherwise the weeks of debate become a deadline. Every lender would race to seize and sell its collateral, and every dominant firm would race to buy, before the law took effect. Tax bills have long used the date of introduction for the same reason.

Can sales that have already happened really be undone?

Within limits, yes. Bankruptcy law already lets a trustee reverse payments and transfers made in the 90 days before a failure. The Act applies the same idea to distressed computing assets sold to the dominant firms after the crisis began, from a date fixed in the bill. The buyer gets back what it paid, and the assets go through the Act's process, with the public given the first option.

What happens to companies that are already in bankruptcy court?

Their cases are moved to the receiver, and the freeze in the Act replaces the bankruptcy stay. Congress has done this before. The northeastern railroads had been in bankruptcy for years when Congress took their assets out of the bankruptcy courts and conveyed them to the publicly created Conrail, and the Supreme Court upheld it.

Why does the FDIC handle failed technology companies?

Because a new agency cannot be built in the middle of a panic, and the FDIC already has the staff, the procedures, and the legal machinery for taking over a failed institution in a weekend. It acts as receiver from the day the Act is signed, with the Department of Energy's national laboratories supplying technical operators. A separate agency, the FDIC-Tech, is established within 18 months and takes over.

How do you keep a company in receivership running when nothing was prepared?

With blunter tools than the plan-ahead version needs. Destroying data, deleting models, withholding passwords, or sabotaging systems at a firm in receivership is a federal crime, and officers must hand over system access within 24 hours. Retention pay is offered widely, because the people who hold the knowledge are the only continuity plan there is. The receiver keeps everything running for 30 days while it learns which services are critical, and then stops speculative spending while keeping paying customers

and critical users online. Sites that have already gone dark get a decision within 90 days on whether restarting them is worth it.

Why does waiting cost more?

Some data centers will already have gone dark and their staff will have left. Some of the best hardware will already have been sold. Lenders who never expected losses will have lent more, against worse collateral. Their lawsuits will be stronger. The public has to advance all of the money before the industry has paid in anything. And the damage to the rest of the economy grows with every week of improvisation.

Why is this a full-length plan and not a short emergency bill?

Because of what happened in 2008. The first draft of the bank rescue was three pages long. It asked for $700 billion and said the Treasury Secretary's decisions could not be reviewed by any court. Congress read it as a blank check and voted it down, and the bill that passed days later was hundreds of pages written in a hurry, much of it by the industry. In a crisis, whoever has a finished text ready has the advantage. The parts of the emergency bill that must work immediately take effect at once, and the parts that take time to build, such as the chip registry, phase in over the following 18 months.

Politics

What would it take to pass this?

A crisis is the moment when it becomes possible. The largest banks, technology companies, and lenders have strong reasons to oppose it, and in ordinary times that is usually enough. When institutions are failing and the public is being asked to pay for another rescue, the question changes from whether to act to which bill to pass. In 2008 the only plan on the table was the bailout, so the bailout is what passed. This version exists so that there is another finished bill on the table.

Who is this for?

Anyone who wants the next crisis handled differently from the last two. That includes community bankers squeezed by consolidation, smaller technology companies that cannot match the giants' computing budgets, researchers priced out of their own field, conservatives who believe that failure should have consequences, and working people who have already paid for two rescues.

What happens after the emergency is over?

The temporary measures give way to the permanent ones in the plan-ahead version. The formula for barred buyers is replaced by the chip registry and the 5% rule. The free deposit guarantee is replaced by risk-based premiums. Designation by statute is replaced by a council with a written record and a right of reply. The FDIC hands over to the FDIC-Tech. Firms that remain private file conversion plans and place system access in escrow. The industry goes on repaying the Treasury until the debt is cleared, and

from then on it is paying for insurance against the next crash, in advance, the way it should have been all along.