Vero
Technologies

MODERNIZATION PLANNING · FRAMEWORK

Start in 2030. Work backward.

Most modernization plans start from what is possible next quarter and stop short. This one starts from the lender you need to be in 2030 and works back to what has to be true, in what order, with each phase paying for the next.

METHOD	HORIZON	COMPANION
Backcasting	**Five years**	**Business Case for Change**

Why forward planning stalls

Ask a lending team to plan modernization and the conversation almost always starts in the present: here is our budget, here is what IT can absorb this year, what can we fit?

That question produces a roadmap of adjacent improvements. Each one is defensible. Together they rarely add up to a different operating model, because nothing in the process ever asked what the destination was. Five years of sensible next steps can leave a lender exactly where it started, with newer versions of the same constraints.

Backcasting inverts the question. You define the state you need to be in, then reason backward to the conditions that make it possible, and only then ask what is affordable this year. The near-term decisions look similar. What changes is which ones you make, and the order.

The difference in practice

FORWARD PLANNING

"We can replace the audit module this year." Chosen because it fits the budget. May or may not be on the path to anything.

BACKCASTING

"In 2030 risk has to be continuous. That needs event data out of servicing. That starts with the audit module." Same first move, chosen for a reason, sequenced deliberately.

The method has one honest cost. It forces a leadership team to commit to a description of the future in enough detail to be wrong about it. Most teams find that harder than the budgeting.

A roadmap is not a list of improvements. It is an argument about sequence.

Step 1 · Describe 2030 in operating terms

Not aspirations. Observable facts about how the business runs on an ordinary Tuesday in 2030.

The discipline is to write statements someone could later verify. “Best in class dealer experience” is not a statement. “A dealer can floor a unit, request a payoff and settle it without contacting a person” is.

Six prompts that produce usable statements

- **Scale.** What size is the portfolio, across how many dealers and which asset classes?
- **Headcount.** How many people run it, and what are they doing with their day?
- **Dealer experience.** What can a dealer do without contacting anyone?
- **Risk posture.** How current is your view of a dealer, and how is an emerging problem surfaced?
- **Product range.** What can you launch, and how long does launching it take?
- **Evidence.** When an examiner or a credit committee asks how a decision was made, what do you hand them?

> **THE TEST FOR A GOOD 2030 STATEMENT**
>
> Could a colleague who disagrees with you tell whether it had come true? If not, it is a value, not a target. Rewrite it until the disagreement would be about facts.

Most teams produce eight to twelve statements. That is the right order of magnitude. Fewer and the picture is too vague to reason backward from; many more and the exercise becomes a wish list.

Step 2 · Name what has to be true

Take each 2030 statement and ask what conditions it depends on. Four keep appearing, in almost every lender, in roughly this order of dependency.

GATE 01

An open architecture

Every downstream capability depends on whether data can move. A closed system forecloses options you have not thought of yet, which is precisely the problem when the target is five years out. Openness is not a feature you add later; it is the property that determines whether later is possible.

The question to ask a vendor: what can I get out, on what schedule, without asking you?

GATE 02

APIs at the events that matter

Not an API programme in the abstract. Specific endpoints on the specific events other systems need to react to: funding, curtailment, payoff, audit result, title status, dealer onboarding, limit change. If an event cannot be subscribed to, every consumer of it becomes a batch file and a reconciliation.

The question to ask: which of my seven core events can another system subscribe to today?

GATE 03

An AI-ready system of record

Useful automation needs structured, current, permissioned data with a clear owner. Most lenders discover their real constraint is not model access but that the operational truth lives in attachments, inboxes and one person's spreadsheet. A system of record that holds events rather than documents is the precondition for anything agentic.

The question to ask: if an agent had to answer "what changed on this dealer this week", where would it read from?

GATE 04

A compliance framework for AI

This gate is usually sequenced last and should be sequenced first. What you need is not policy language but a working set of answers: what may a model touch, what must a human approve, what is logged, what is redacted, how is a decision reconstructed months later. Without it, every promising use case stalls in review.

The question to ask: who signs off on the first production use, and what do they need to see?

Step 3 · Sequence it so each phase pays for the next

The gates are dependencies, not a schedule. A schedule has to answer a different question: what can we afford, and what does each step return while the next one is still being built?

That constraint is what makes a five-year plan survivable inside a real institution. A phase that returns nothing for eighteen months will be cancelled in month nine, regardless of how sound the destination is. Sequence so that every phase produces a result visible to whoever controls the budget.

PHASE 1 · MONTHS 0 TO 6

Instrument the highest-drag workflow

Pick the single workflow consuming the most manual hours. Move it, measure it before and after, and publish the number. Nothing structural changes yet, which is the point: this phase exists to establish that the programme produces measurable results.

Returns: recovered hours, and the credibility to fund Phase 2

PHASE 2 · MONTHS 6 TO 18

Move the system of record, open the architecture

The heavy phase, and the one that clears Gates 01 and 02. Servicing events become structured and subscribable. This is where the temptation to do everything at once has to be resisted: migrate workflow by workflow, keeping the tools that already work through integration rather than replacement.

Returns: headcount avoided as volume grows, plus retired maintenance

Step 3 · Sequencing, continued

PHASE 3 · MONTHS 18 TO 36

Make risk continuous, and put the AI framework in place

With events flowing, monitoring stops being a monthly report and becomes a signal. Clear Gate 04 here, on low-stakes ground, before anything consequential depends on it. Document extraction and reconciliation are the right first uses: high volume, low judgement, easy to audit.

Returns: lower loss severity from earlier detection, reduced audit cost, analyst hours

PHASE 4 · MONTHS 36 TO 60

Compound it

The phase the first three exist to make possible. New asset classes and new programmes launch against infrastructure that already exists, at a marginal cost that no longer scales with headcount. This is where the 2030 statements either come true or reveal which gate was skipped.

Returns: revenue from products that were previously uneconomic to launch

NO NUMBERS HERE ON PURPOSE

Every figure in this phasing depends on your volume, your cost base and your growth plan. Anyone who hands you a five-year ROI percentage without those inputs is selling. Use the Business Case for Change worksheet to put your own numbers against each phase.

Running the session

This works as a half-day with the right eight people in the room and badly as a document circulated for comment.

- **Who.** Credit, servicing, operations, risk, technology, and someone who talks to dealers weekly. The last one is the most often missed and the most useful.
- **Sequence.** Ninety minutes on 2030 statements. Ninety on the four gates. The rest on sequencing and on what Phase 1 is.
- **Rule.** No vendor names and no budget numbers until the gates are agreed. Both collapse the conversation into procurement before the destination exists.
- **Output.** Eight to twelve statements, four gate assessments with an honest current state, and one Phase 1 with a named owner and a date.

> **WHO TO LEAVE OUT**
>
> Not procurement, and not anyone who will spend the session defending a system they chose. Both are needed later. Neither helps while the group is still trying to describe a destination honestly.

The three ways this goes wrong

01 The 2030 picture is a wish list

If nothing in it would require a hard trade-off, the group has described a better version of today rather than a different operating model. Push until something in the list is genuinely uncomfortable.

02 Phase 1 is chosen for ease rather than drag

The easiest workflow to move is rarely the one costing the most. Easy produces a successful project and an unconvincing number.

03 The compliance gate is deferred

Deferring Gate 04 feels efficient because nothing depends on it yet. It becomes the thing that stalls Phase 3 entirely, eighteen months after it could have been settled cheaply.

Run it with us in the room

We will facilitate the session against your portfolio, bring what we have seen at other programmes, and give you an honest read on which gates you have already cleared. You keep the output whether or not we ever work together.

Book the session

vero-technologies.com