

\$120M+ ARR B2B SAAS · 2025–2026

Built what customers were leaving over.

Retention up five points in a quarter.

A product and engineering organization rebuilt around AI in six months. Every customer conversation analyzed, priorities ranked by revenue at risk, and code drafted by AI and reviewed by engineers before release. Same team, same budget: four times the output, and net revenue retention up five points in one quarter.

+5 pts net revenue retention in one quarter, based on a cohort analysis	4× features shipped 20 → 81 per quarter; 2× by month 3, 4× by month 6	2–4 wks idea to production down from 4–6 months	0 added headcount product, design or engineering
---	---	---	--

THE STARTING POINT

A dull edge: the team was guessing, slowly.

Customers heard through a sample More than 16,000 sales and customer success calls a year sat unused in the call recorder. Roadmaps were built on assumptions and a handful of interviews.	The loudest voice won Prioritization depended on who was in the room, not on the bookings at stake or the revenue at risk of churning.	Quarters to ship Idea to production took four to six months. By the time a feature shipped, the customer who asked for it had often moved on.
--	--	---

THE APPROACH

Sharpen the four Edge pillars, end to end.

The work followed the four Edge pillars of the NXProduct 7-Pillar AI-Native framework: hear every customer, decide on revenue evidence, deliver with AI behind human gates, and measure what each release changed. AI does the high-volume work between the gates; people make every call that matters.

The next page shows what AI automated and what the team owned at each pillar.

THE EDGE · PILLARS 01–04

From customer conversation to measured release.

01 Customer Signal

Call recordings (16,000+ a year) → AI extraction → structured evidence file → an internal ranking app, refreshed weekly

AI AUTOMATES

Reads every sales and success call, extracts each product request with the exact customer quote and a link to the transcript, and groups them into recurring signals. A weekly batch keeps it current.

PEOPLE OWN

Product managers validate the signals and correct misreads. Leadership sets the sources and privacy rules.

RESULT

~700 recurring signals, each traceable to the customers who voiced them.

02 Decision System

AI AUTOMATES

Ranks every signal on two axes: bookings potential (requests prospects cite as purchase blockers) and churn correlation (requests customers cite when they cancel or downgrade), weighted by the ARR attached.

PEOPLE OWN

PMs and leadership choose what to build. Every decision shows the quotes, the revenue at stake and the churn risk, so prioritization debates end.

RESULT

~700 signals narrowed to 150 high-impact features.

03 AI-Native Delivery

Requirements → design → code and pull request → engineering review → release

AI AUTOMATES

Drafts the requirements from the customer evidence, generates wireframes and design specs from them, then writes production code and opens a pull request tagged for review.

PEOPLE OWN

The PM edits and approves the spec; designers own the experience; engineers review for correctness, security and edge cases, test, fix and release.

RESULT

Idea to production cut from 4–6 months to 2–4 weeks; 20 → 81 features a quarter.

04 Measured Impact

AI AUTOMATES

Links each release to the requests and accounts behind it, and tracks retention, bookings and adoption afterwards.

PEOPLE OWN

Leadership judges the outcome and reprioritizes. The team shipped the top churn-driving requests first.

RESULT

Net revenue retention up 5 points in one quarter, based on a cohort analysis.

THE SPINE · WHAT HELD IT UP

A sharper edge needs a team built to use it.

06 · Operating Model

No one was cut and no one was hired. Roles were redefined: engineers moved from writing code from scratch to validating and shipping; PMs from writing tickets to deciding. A PM, a designer and a couple of developers now cover a whole bet, and small squads cleared two years of technical debt.

07 · R&D Economics

Four times the output at the same cost: cost per shipped feature down about 75%. The freed capacity went to quality, technical debt and the next bets, not to cuts.

RESULTS

	BEFORE	AFTER
Net revenue retention	Baseline	+5 points in one quarter (cohort analysis)
Features released per quarter	20	81 (2× by month 3, 4× by month 6)
Idea to production	4–6 months	2–4 weeks
Customer conversations analyzed	A sample	16,000+ a year, refreshed weekly
How priorities were set	Loudest voice in the room	~700 signals ranked by revenue → 150 high-impact features
Product, design and engineering headcount	Baseline	No change

“We didn’t hire a single engineer. We stopped asking engineers to do the work AI does faster. The same team that shipped 20 features a quarter now ships 81, and because we built what customers were leaving over, net revenue retention rose five points in a quarter.”

Former SVP Product at a \$120M+ ARR B2B SaaS company

ABOUT THE AUTHOR

Guillaume “G” Vives has led product at enterprise SaaS companies for 25 years, including Zuora, Oracle, RMS, 360Learning and Lattice: two IPOs and \$4.5B in exit value. He designed and led the transformation described here, and now helps SaaS leadership teams do the same.

HOW MEASURED

Net revenue retention from a cohort analysis comparing the quarter before and after the churn-driving requests shipped. Throughput is features released to production per quarter, counted the same way before and after, with the same team and budget, over six months. The company is anonymized; details can be validated privately on request.

Ready to find *your edge*?

Start with a leadership diagnostic: one working session, a first-read maturity score and three priority moves.

nxproduct.com/contact