

AI Transformation Playbook

STAGE III: EMBEDDED AI PRODUCT FUNCTIONALITY

Julian Castelli Growth Scaling Partners
Presented by **Julian Castelli**, Board Chair and Operating Partner



The Final Sprint!


Stage 1	Stage 2	Stage 3 - You are Here!
Knowledge Worker AI Enablement	AI Process Transformation	Embedded AI Product Functionality
We started here	Builds on Stage 1	Builds on Stage 2 ►

AI-native competitors are NOT building better dashboards

- Designed outside-in
They start from the customer's job to be done — not your legacy architecture
- They don't improve — they replace
Existing architecture isn't iterated on. It's bypassed entirely.
- Competing for 100% of the job
Not 25–50% of it. The whole workflow. Every step.

They're not coming for your features.

They're coming for your customer's entire workflow.



Most SaaS Companies are Stuck in the Wrong Conversation

THE CONVERSATION YOU'RE HAVING

What's our AI strategy?

How do we add AI to our roadmap?

Which AI features should we add?


THE CONVERSATION YOU SHOULD BE HAVING

What is the complete job our customer needs done?

What would a *Level 6 Product* look like for them?

Are we building toward that — or away from it?

From “Descriptive” to Jobs Complete




"You don't get to decide when your customers expect outcomes. They're already deciding."

SECTION 2

The 6-Level AI Product Maturity Framework

The AI Product Maturity Framework



*Built on top of the Gartner Data Analytics Maturity Model

Level 1 — Descriptive: What happened?

Dashboards


Moves from past data to forward predictability

Passive readout

The product delivers numbers. That's the full value.

Human interprets and decides

All analysis and action lives outside the product.



Almost all legacy SaaS products were born here. Most have never left.

Level 2 — Diagnostic: Why did it happen?

Surfaces root cause


Connects outcomes to contributing factors.

Explains the variance

Goes beyond the number to the story behind it.

Still advisory

Diagnosis without prescription. The human acts.



Knowing why is valuable. But the human still has to decide what to do about it.

Level 3 — Predictive: What Will Happen Next?

Forward Looking Signals


Moves from past data to future probability.

Probability Modelling

Confidence-scored forecasts, not guesses.

Pre-emptive alerts

Flags risk before it becomes a crisis.



Prediction is powerful. Significant increase in value creation! (Influence Outcomes)

Level 4 — Prescriptive: What should I do?

Specific Recommended Action

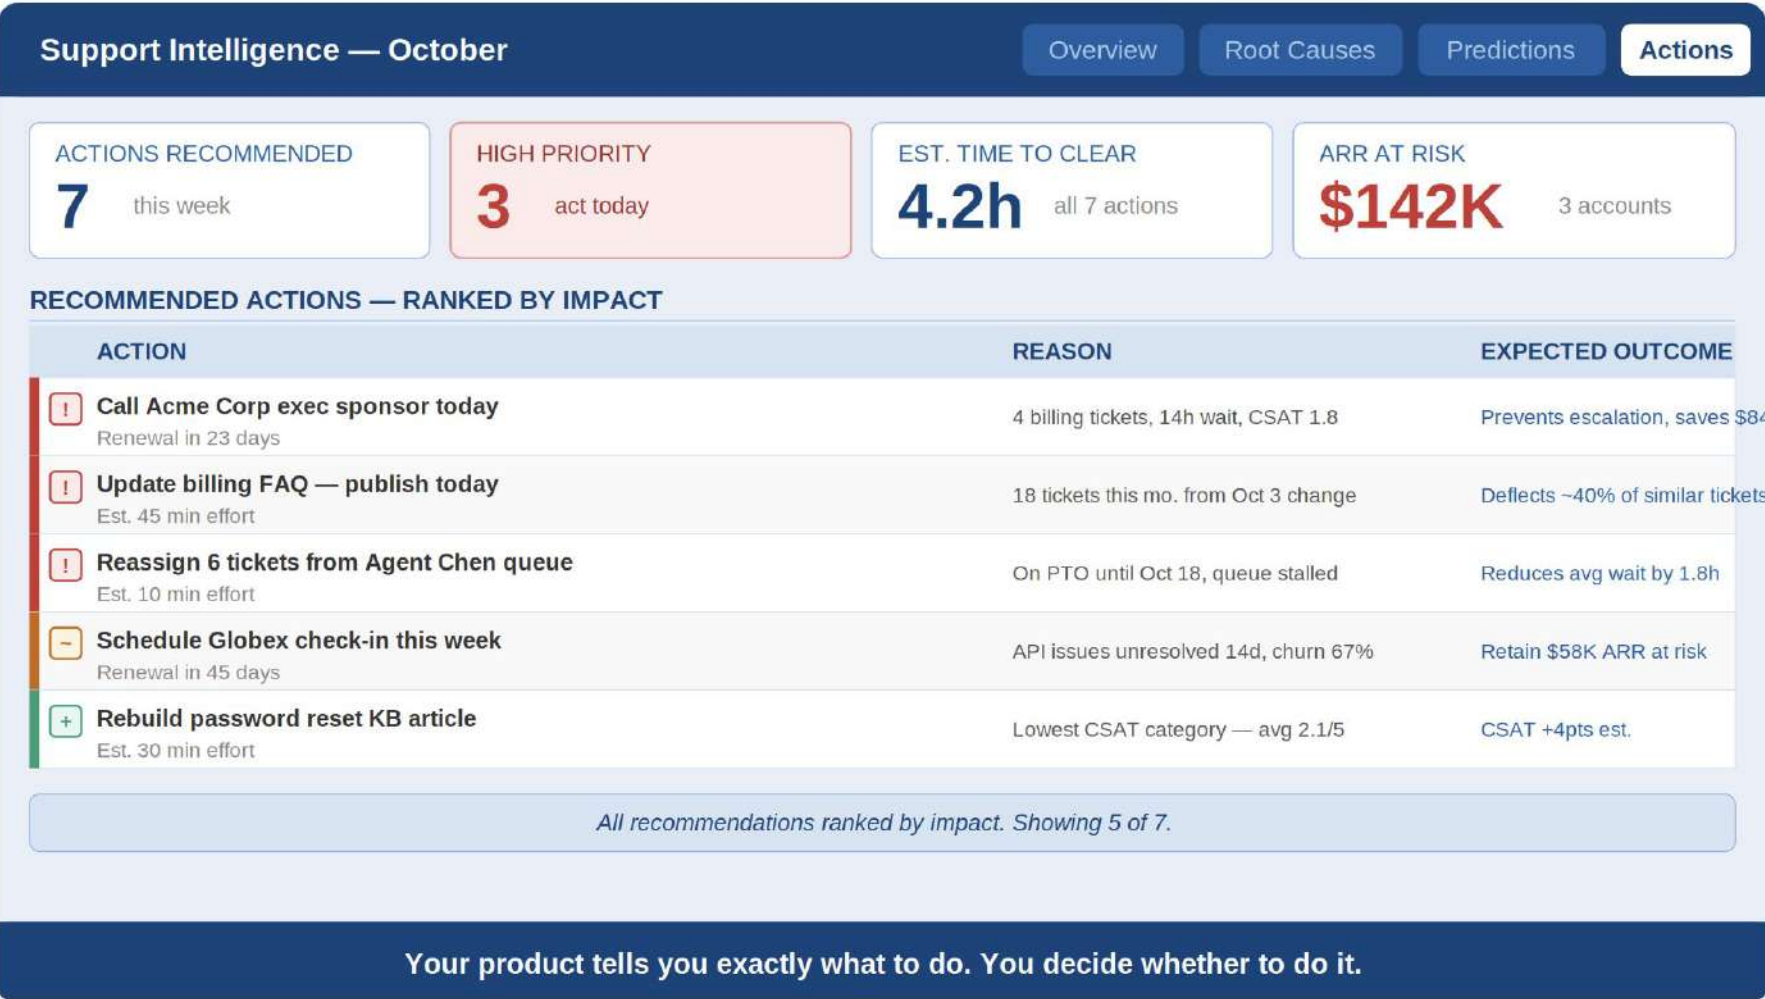
Not a range of options — a clear directive.

Ranked Options

Shows its work. Earns trust through transparency.
w/Reasoning

Explains Reasoning

The product can be questioned. And it can answer.



The product is now co-piloting. You're still flying — but it knows the route.

Level 5 — Human in the Loop - Approves AI Actions

Telegraphs The Action

Shows you what it plans to do – before doing it.

Gets Confirmation

One-click approval. You stay in control of the threshold.

Executes & Reports Back

The work is done. You receive the outcome summary.

Support Intelligence — October

OverviewRoot CausesPredictionsActionsAutopilot

Autopilot active — 3 actions awaiting your approval · 4 completed this week

All systems on

AWAITING YOUR APPROVAL

Approved

Renewal outreach email — Acme Corp exec sponsor

Draft ready: exec tone, references billing issue, proposes Oct 14 call. Sending in 10 min.

PreviewUndoEdit

Editing...

Billing FAQ update — publish to help center

You are editing: adjusting tone in paragraph 2. AI draft 90% ready.

PreviewApprove & Send

Approved

Ticket reassignment — 6 tickets redistributed from Agent Chen

Assigned to Sarah Kim (3) and Marcus L. (3). Avg wait projected to drop 1.8h.

PreviewUndoEdit

COMPLETED THIS WEEK — AFTER YOUR APPROVAL

Globex check-in email sent — Oct 6, 9:14am

Response received. Meeting booked Oct 12.

Password reset KB article rebuilt and published — Oct 5

CSAT for category: 2.1 → 3.8 in 48h.

3 tickets auto-resolved via updated FAQ — Oct 7

No agent time required.

Your product acts. You supervise.

Now you are on Auto-Pilot, but you retain control through Approvals

12

Level 6 — Autonomy - AI Executes 80% of Activities.

No HIL for 80% of Activities


The system executes end-to-end. No approval gate.

Outcomes Delivered & Reported

Results surface automatically. You review — you don't approve.

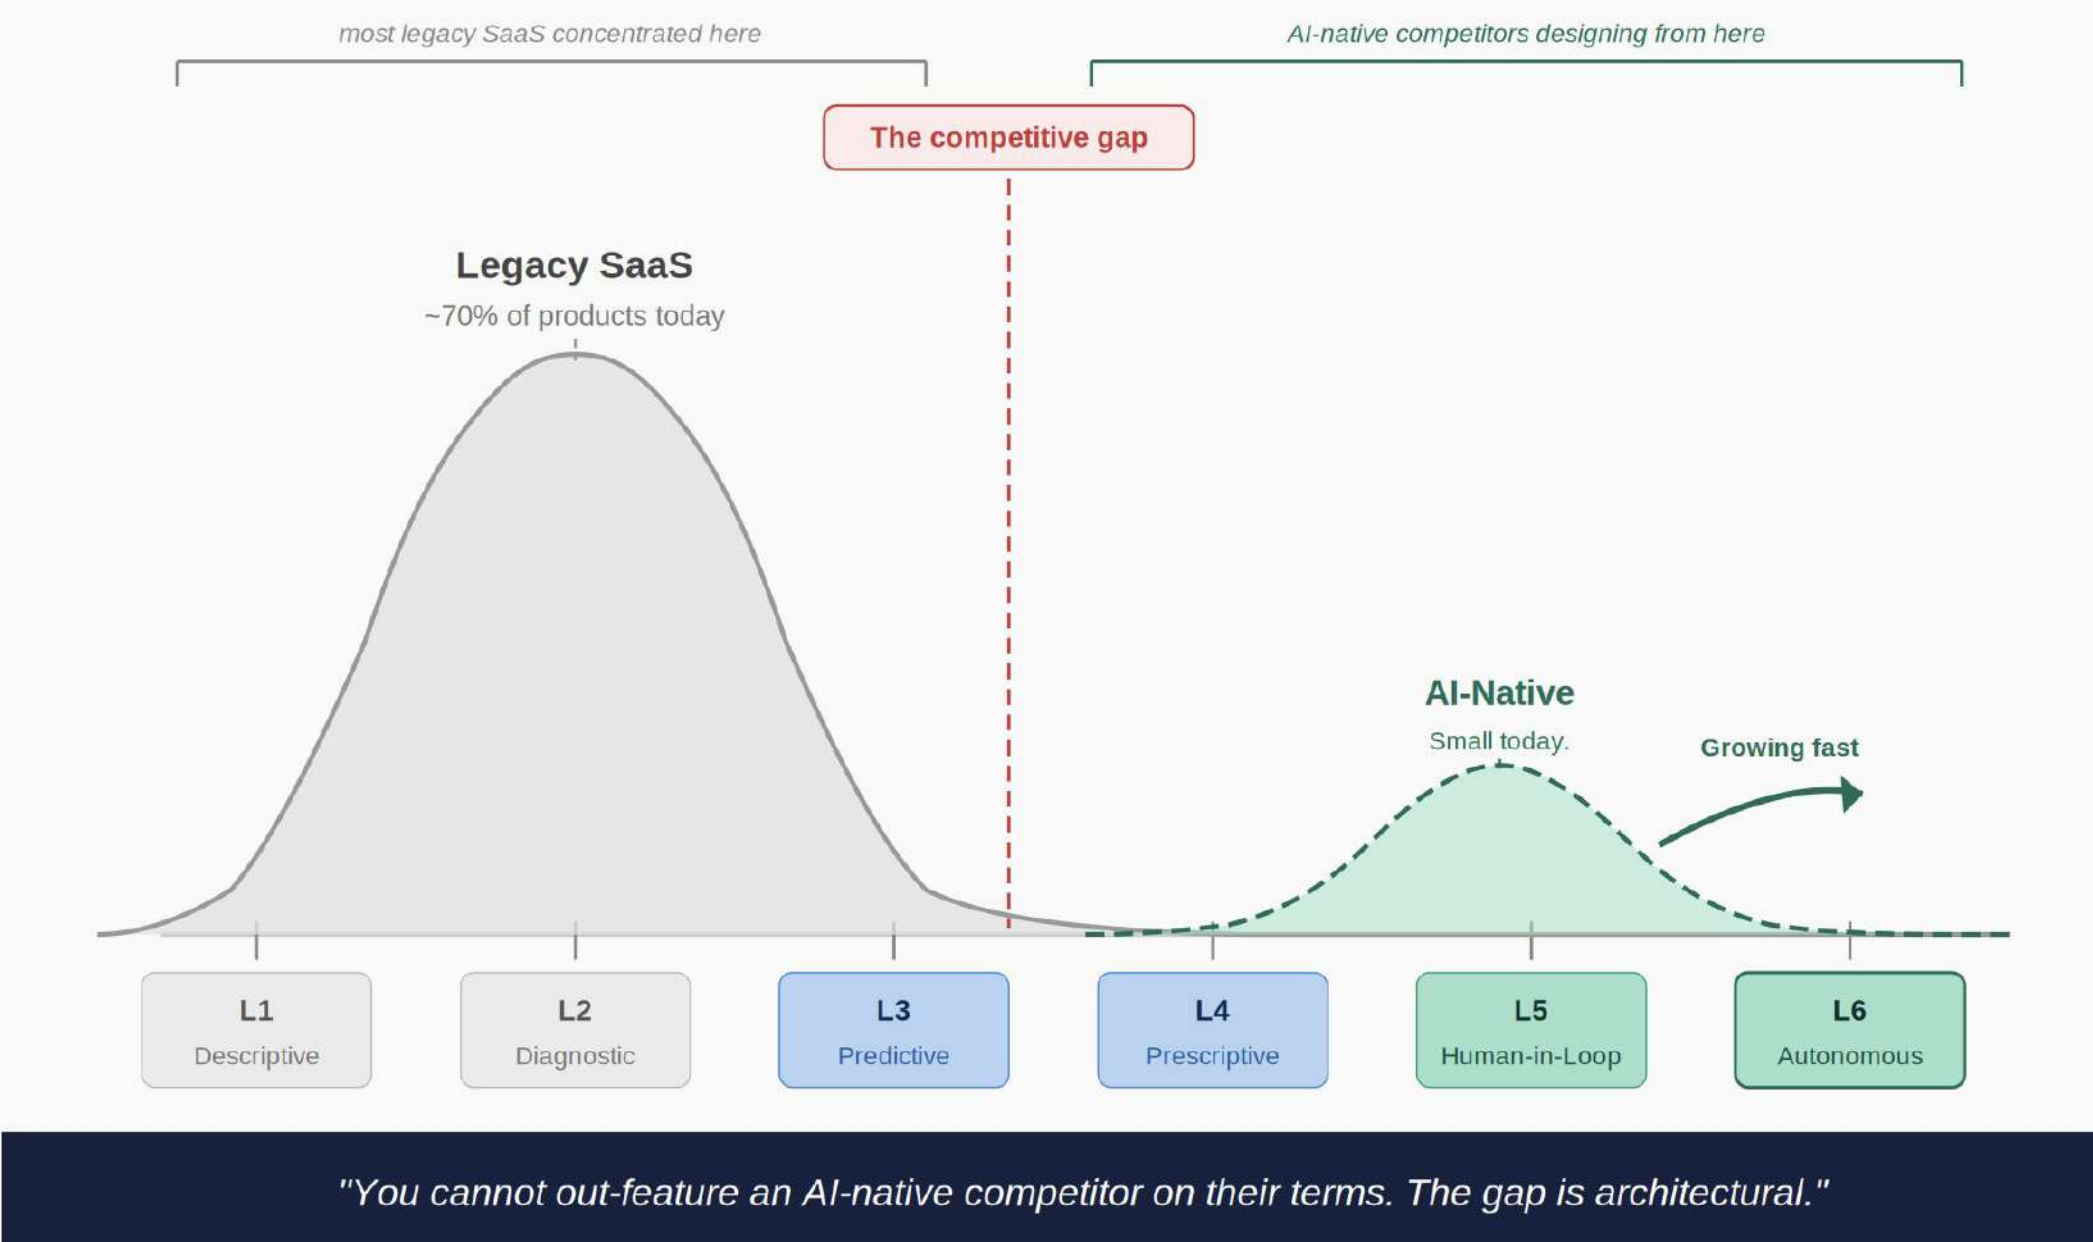
Human Manages Outcomes

Exceptions escalate. Everything else runs itself.




Human is Now Managing Outcomes and Autonomous Optimization

Where Are You Today?



AI Native Companies are Launching with a Head Start

Your pricing model is a product declaration




Levels 4-6 Unlock Huge Pricing and TAM Opportunities

SECTION 3

It's Already Happening

Real companies. Real transformations. Mapped to the 6-Level Framework.

Per Resolution Pricing



Three Things Tesla Teaches Every SaaS CEO

1

The journey is longer
than anyone thinks
Tesla - 12 years in

2

L4 → **L5** was expected to
be hard - it wasn't.

L5 → **L6** is the hardest
transition


3

Start with the **End in
Mind!**

Working on FSD for over 10 Years!!!



Tesla's AI Transformation Journey —



"Hardware shipped in 2014. Architecture for Level 6 locked in by 2016.
Twelve years later — still at Level 5 in 2026.
You cannot skip the staircase."

Three Things Tesla Teaches Every SaaS CEO

1

The journey is longer than anyone thinks

2

L4 → **L5** was expected to be hard - it wasn't.

L5 → **L6** is the hardest transition

3

Start with the **End in Mind!**



Autopilot hack!!!!

FOREVER

MODE!



Three Things Tesla Teaches Every SaaS CEO

1

The journey is longer than anyone thinks

2

L4 → **L5** was expected to be hard - it wasn't.

L5 → **L6** is the hardest transition

3

Start with the **End in Mind!**

Elon Shipped the Hardware Before the Software Was Ready



SECTION 4

How To Get There

Starting With the End in Mind

The destination determines the architecture. Not the other way around.

Avoid These 2 Failure Modes

Failure Mode 1

Bolting AI onto existing architecture. The architecture fights you. You're stuck at 50% of the job.

Failure Mode 2

Building features without a clear picture of the destination. You reach 80% product complete and discover you can't get to 100%. The car with no wheels.

Start with the end in mind isn't a vision statement. It's an architectural decision.

What Does Level 6 Look Like for Your Customer?

1

The Complete Job

Describe the complete job your customer is trying to get done – not the features they use today.

2

The Autonomous Morning

If your product were fully autonomous – what would your customer wake up to every morning?

3

The Architecture Test

What data, trust, and integrations would you need to build that? Are you building toward it – or away from it?

You can't architect for a destination you haven't defined.

Level 6 Products Don't Manage Workflows — They Complete Jobs

Revenue Intelligence

Deal closed. Here's what I did.

Customer Success

Renewal secured. Churn risk resolved.

Summary in your inbox.

HR Recruiting

3 candidates advanced. 12 screened out. Interviews scheduled.

The product doesn't ask for instructions. It delivers outcomes and reports back.

SECTION 5


The Diagnostic — Do We Have What it Takes to Get to AI Automation?

6 Readiness Criteria

Data	Is our customer data ready to act on?
Access	Can our product reach the systems it needs?
Capability	Can our AI produce actions — or just insights?
Integration	Are we integrated to execute — or just to display?
Trust	Has our customer given us permission to act?
Pricing	Does our pricing reflect the value we deliver?

DAT

Is Our Customer Data Ready to Act on?



Real-time visibility

Do we know what our customers do in our product in real time?

Outcome linkage

Can we tie product behavior to business outcomes?

Unified or siloed?

Is our data siloed by function or unified by customer?

Clean, structured, accessible data is the price of entry – not a nice-to-have.

What Systems Do We Need to Reach — and Can We?

Read and write

Can our product read and write to the systems our customers use?

Act, not just advise

Do we have the integrations required to act — not just advise?

Trust before access

Are we asking for data access before we've earned the trust to use it?



Access without trust is a dead end. Earn the right before you ask for the keys.

Can We Act on What We Know?

Actions or insights?

Do our AI models produce actions – or just insights?

Close the loop

Can the product close the loop without human intervention?

Action vs. dashboards

Are we building toward action – or investing in better dashboards?

Reporting is not action. If your AI only tells you what happened, you're still at L2.

Are We Integrated to Execute — or Just to Display?



Execute or display?

Are we integrated deeply enough to execute
— or just to display?

Bidirectional access

Do we have bidirectional connections or
read-only access?

L3 → L5 gap

What would we need to build to move from
L3 to L5?

Read-only is not enough. The gap between insight and outcome is integration depth.

Has Our Customer Given us Permission to Act?

Earned permission

Would our customers trust us to send an email on their behalf? Make a decision?

Accuracy track record

Have we demonstrated enough accuracy and reliability to earn autonomous action?

Have we even asked?

Is trust a blocker – or have we never asked the question?



Trust is not assumed. It is earned through accuracy, transparency, and restraint.

PRICING

Does Our Pricing Reflect the Value we Deliver?

Seats vs. outcomes

Are we charging for seats when we're delivering outcomes?


Shift readiness

Could we shift to usage or outcome-based pricing today?

Signal mismatch

Does our pricing signal Level 2 while we claim to be Level 4?

Your pricing model is your product thesis made financial. Misalignment is a credibility problem.



Score Ourselves - To Track AI Automation Readiness

	L1-L2	L3	L4	L5	L6	Your score
Data	Siloed, unstructured	Structured, accessible	Unified by customer	Real-time + outcome-linked	Self-improving loops	
Access	No integrations	Read-only APIs	Read/write key systems	Bidirectional, multi-system	Full system authority	
Capability	Reports only	Predictions generated	Recommends surfaced	Supervised execution	Autonomous action	
Integration	Display only	Single system read	Multi-system read/write	Closed-loop execution	Full workflow ownership	
Trust	Not asked	Advisory only	Conditional approval	Earned, supervised	Full delegation	
Pricing	Per-seat, feature-based	Usage tiers	Value-linked	Outcome-based	Shared upside	

SECTION 7

What To Do Next



Three conversations to have this week

1

With Your Product Team

Where does our product sit on the 6-level framework today – honestly?
What would it take to move one level?

2

With Your Customers

What is the complete job you're trying to get done? What would it look like if our product did that for you – without you having to manage it?

3

With Your Investors

Here is our current level. Here is our Level 6 destination. Here is the architectural sequence to get there.

The hardest part isn't the technology. It's the clarity about where you're going.

What does Level 6 look like for your customer — and are you building toward it?

Not the features. Not the roadmap. **The destination.**

Every company in this room has a window. The question is whether you use it.

Who I Help

Experienced Venture/Operating Partner:



Board Chair, Lead Independent Board Director:



Founder/ CEO / Executive / Operator:



Professional Business Foundation:



How I Help



Julian Castelli

OPERATING PARTNER

EXPERTISE:

Board Leadership, Value Creation, AI Strategy, GTM Execution, Team Building, Strategic Finance, Funding & Exits.



[Contact](#)



[Profile](#)



[Blog & Podcast](#)



[Community](#)

Appendix

SECTION 06

The Staircase — Your Level-by-Level Work Plan

A structured transition path from passive data to autonomous action — one level at a time.

Moving from Descriptive to Diagnostic

What Has to Be True

- Structured, queryable data in place
- Event tracking instrumented across key workflows
- Consistent data definitions and taxonomy

What You Build

- Root cause models
- Variance analysis engines
- Anomaly detection pipelines

What Changes

Your product stops *showing* data – it starts *explaining* it.

Moving from Diagnostic to Predictive

What Has to Be True

- Historical data depth – minimum 12–18 months
- Labeled outcomes for model training
- Model training and versioning infrastructure

What You Build

- Churn prediction models
- Expansion and upsell signals
- Risk scoring by account or workflow

What Changes

Your product stops explaining the past – it starts *seeing the future*.

Moving from Predictive to Prescriptive

What Has to Be True

- Validated model accuracy – confidence thresholds defined
- Action catalog built and mapped
- Recommendation logic and ranking logic established

What You Build

- Next best action engine
- Ranked recommendation surfaces
- Reasoning transparency layer (show your work)

What Changes

Your product stops *observing* – it starts *advising*.

Moving from Prescriptive to Human-in-the-Loop

This is the hardest transition on the staircase. You are no longer building a feature – you are establishing a trust relationship.

What Has to Be True

- Customer trust earned through L4 accuracy track record
- Accountability framework and governance model established
- Full audit trail and explainability built in

What You Build

- Action execution layer (system acts on approval)
- Confirmation and approval workflows
- Closed-loop reporting and audit log

What Changes

Accountability shifts from human to system. This is not a feature upgrade. It's a trust relationship.

Most companies underestimate this transition. Tesla spent 6 years on it.

Moving from Human-in-the-Loop to Fully Autonomous

What Has to Be True

- Proven L5 accuracy track record across customer base
- Explicit customer opt-in to full autonomy
- Regulatory clearance secured where applicable

What You Build

- Exception handling and edge-case escalation logic
- Autonomous action reporting (post-hoc transparency)
- Override and rollback controls

What Changes

The product doesn't ask for approval. It **acts** and reports. You can sleep.

*Level 6 isn't a feature. It's a new category of product.
Most companies underestimate this transition. Tesla spent 6 years on it.*