

SERVICENOW AI · READINESS & RUNSTATE

How ready are you to **let AI act?**

From AI ambition to readiness, action and continuous value.



THE PROBLEM

We've invested in AI. Where are the results?



Five leaders. One question. What can you safely hand over?

81%

of executives believe more AI is running than they can account for.

ServiceNow with ThoughtLab, October 2026 · 550 executives

69%

of organisations suspect or have evidence of prohibited public GenAI use.


Gartner, November 2025 · 302 cybersecurity leaders

YOUR SERVICENOW AI READINESS GAP

What's stopping us from letting AI do the work?

Can we trust the data, rely on the workflow, control the action and justify the consumption?

EXHIBIT 01 • THE EXECUTION GAP



ServiceNow AI is ready. The question is whether you are.


Crossfuze assesses readiness across the foundations of data, AI and agents, workflows and governance to identify what must change before autonomy can safely increase.

[Open the tools](#) →

THE ROUTE

What does the path to value look like?

Three steps. The first two cost nothing. Explore with the tools, find your first outcome in 90 minutes, then deliver it with AI RunState.



Clear next steps in 90 minutes. Outcomes every 90 days.

PREPARING FOR THE AUTONOMOUS ENTERPRISE

Autonomy is earned.

“

AI makes execution easier. Trust becomes the harder problem.

You prove the data can be trusted. You prove the workflow behaves as expected. You prove the controls work. Then you expand the boundary. **Permission to automate is earned.** ”



Oliver Nowak


AI Practice Lead

CROSSFUZE

Crossfuze helps you decide what AI is ready to do.

THE FOUR FOUNDATIONS

What does AI need to stand on?



AI can only act on what sits beneath it.

Your weakest foundation sets your autonomy ceiling.

Named people remain accountable.

Weights are each foundation's share of the seven-dimension readiness score on page 08; dimension 06, platform and entitlement (10%), is assessed per Zero.

YOUR MATURITY


How ready are we?

ServiceNow's Enterprise AI Maturity Index gives the direction. **Our evidence determines where our AI readiness actually sits on the five-stage path to greater autonomy.**

Autonomy becomes practical at stage 03, and each stage is earned by lifting your weakest foundation.

THE READINESS JOURNEY

AUTONOMY BECOMES PRACTICAL AT STAGE 03



Crossfuze readiness scale. Stages 02-05 align to the Experimenter, Scaler, Advancer and Transformer stages of the ServiceNow Enterprise AI Maturity Index 2026. The Five Zero outcomes and the Autonomous IT ambition are ServiceNow's. The readiness scale is Crossfuze's.

READY TO PLACE YOURSELF?

Take the ServiceNow AI Readiness tool →

THE SCORE

Can we trust the score?

Every AI decision starts with an AI Readiness score. Each dimension is scored against evidence, weighted, and shown with its status, so nobody has to take the number on trust.

SEVEN WEIGHTED DIMENSIONS

01	Data quality DATA Reliable fields, records and classifications		20%
02	Process determinism WORKFLOWS Decisions, exceptions and rollback understood		15%
03	Knowledge and content AI AND AGENTS Approved, current content to ground answers		15%
04	Workflow and integration WORKFLOWS The technical estate can run the work end to end		15%
05	Governance and security GOVERNANCE Access, approval, monitoring and stop controls		15%
06	Platform and entitlement ESTIMATOR Package and Assist allocation fit the plan		10%
07	Ownership and adoption GOVERNANCE Named owners, real users, capacity to measure		10%

FIVE ANCHORS, NOTHING IN BETWEEN

0	3	5	8	10
None	Weak	Mixed	Strong	Direct

Every score has to be argued from evidence.

EVERY FINDING CARRIES A STATUS

Observed	Calculated	Inferred
Assumed	Validated	Corrected

The wording of each finding matches its status.

GAPS STAY VISIBLE

Full	100% of the weight evidenced
Provisional	At least 80%, including every 20% dimension
Unavailable	Otherwise

Missing evidence is never scored as zero.

The score decides
what happens next.

Not ready

Fix the limiting foundation first, then reassess.

Ready

Design the roadmap for the chosen use case. Confirm the required controls before autonomous action.

THE BILL

What will AI cost us as it scales?

The agentic cost paradox. The more AI does, the more it consumes. Readiness decides whether that consumption turns into value or into an unplanned bill.

THREE PACKAGES

AI IS INCLUDED AT EVERY LEVEL

Foundation

ASSISTED

Summarisation, categorisation and task-based help.

Advanced

AGENTIC

Agentic workflows that understand context and automate a whole process.

Prime

AUTONOMOUS

Custom AI specialists and agents that make decisions with role-based expertise.

[Open the Assist Estimator →](#)

HOW ASSISTS SCALE

CONSUMPTION FOLLOWS THE WORK, NOT THE SEAT

25

Assists for a narrow run, 0 to 4 tools

50

Assists for a run of 5 to 8 tools

150

Assists for a wide run, 9 to 20 tools:
6× a narrow run



CROSSFUZE VIEW

We model what your chosen outcome may consume before you commit.


[Book your session →](#)

Sources: ServiceNow New Commercial Model Overview, April 2026; ServiceNow assist-consumption guidance, 2026. Entitlements vary by product and package. Assists and Data Fabric credits are separate meters. Total consumption depends on run volume and mix.

THE DELIVERY PROBLEM

Traditional project plans don't work for AI.

Readiness tells you what you could do safely. It does not solve the delivery problem. Traditional projects lock decisions in early, make changes expensive and delay value until the end, exactly when the market may already have moved on.

**Decided too early**

Scope is locked before the platform and the market settle.

Expensive to change

Every pivot needs a change order, with more cost and more time.

Value arrives last


The payoff lands at go-live, if it is still relevant.

Illustrative. Not every long programme fails, but AI moves faster than a fixed plan can.

THE 18-MONTH TRAP

Why do most AI visions stay in the slide deck?

The problem is not a lack of vision. An 18-month fixed-scope project can lock priorities before the market, the platform and the business have finished changing. Each pivot becomes harder, while measurable value arrives late.



Frozen priorities

Month 1 decisions drive the outcomes in month 18.

Change orders

Anything new costs extra time and money.

Late ROI

Most of the payoff lands at the end, **if it is still relevant.**

OUTDATED ON ARRIVAL

By go-live, the platform and the market have moved on.

So the vision waits. And the market keeps moving.

Illustrative timeline of a traditional fixed-scope AI programme.

INTRODUCING CARS

THE DELIVERY METHODOLOGY BUILT FOR AI

Crossfuze AI RunState (CARS)

Built for how AI moves. CARS is a delivery service built for AI. It manages your ServiceNow AI Action Plan, prioritises what gets built and delivers it to production every 90 days.

BUILT FOR TODAY'S CHALLENGES

Where should we use AI?

AI Action Plan

UPDATED EVERY 6 MONTHS

What do we build next?

Prioritization & Design Sprints

QUARTERLY

How do we get it live?

Outcome Delivery Cycles

CONTINUOUS



It starts from your free 90-minute assessment. Your readiness score, expected Assist usage and prioritised 12-month roadmap become the brief.

One team. One service. One flat monthly rate.

Board-ready AI Action Plan

A strategy leadership can stand behind, in 4 weeks.

Pivot without change orders

Priorities reset every 90 days, at no extra cost.

Innovation every 90 days

Senior advisors and hybrid delivery, working as one team.

The free 90-minute assessment and the paid AI RunState service are separate. You decide whether to continue.

WHAT YOU GET

What do we get at every step?

Every step ends in something you can sign off. You make the decisions. We carry the work.

01 STEP 1 • 4 WEEKS Plan. Leadership decides the role of AI before anyone builds anything.	WHAT YOU GET 12-month incident and request clustering, turned into an AI opportunity and value map - End-to-end agent journey maps with a friction log - Executive AI Action Plan with future-state process and ServiceNow capability map
02 STEP 2 • EVERY 90 DAYS Prioritise. Priorities updated every quarter, without new paperwork.	WHAT YOU GET - Approved three-month outcome set, capacity plan and signed-off backlog, handed to delivery - Outcomes scored on value, confidence and effort, with a parking-lot register and return points - Build-ready stories with technical feasibility, estimates and a sequenced dependency map
03 STEP 3 • 90-DAY CYCLES Deliver. Innovation into production every 90 days.	WHAT YOU GET 80 hours a month of hybrid delivery capacity, led by a dedicated onshore Technical Lead - Sprints, showbacks, peer review and testing to ServiceNow best practice - Release management, production verification and monthly status reporting
04 STEP 4 • MID-YEAR Recalibrate.	Your strategy never goes stale. It is re-tested against what has shipped and what is new, and feeds straight into the next sprint.

Seven chances to pivot in 18 months. No change orders.


A sprint every 90 days, a mid-year recalibration and a fresh AI Action Plan at renewal, while delivery keeps running.

Indicative structure. Deliverables, hours and timings are confirmed in your statement of work.


THE FIRST 12 MONTHS

What does our first year look like?

Months 1 to 3 establish the AI Action Plan and approve the first build-ready priorities. Delivery cycles then bring approved changes into production every 90 days, with a mid-year strategy recalibration.



EXAMPLE · ONE ILLUSTRATIVE ITSM YEAR



One engagement. The effort moves to where the value is.

Illustrative ITSM year, not a commitment or scope of work. Nothing is built until you approve it. *Requires the right Now Assist licensing.

WHO DOES THE WORK, AND WHEN

How do we get started, and keep going?

One 12-month engagement. Plan in month 1, then prioritise and deliver every 90 days.



Oliver Nowak
with Crossfuze Advisory

01 **PLAN • MONTH 1** | 4 weeks
AI Action Plan
Workshops, use cases, roadmap

02 **PRIORITISE • EVERY 90 DAYS** | Design sprint
Prioritization & Design
Your team sets the priorities. Crossfuze turns them into stories and a three-month capacity plan, ready to build.

03 **DELIVER • 90-DAY CYCLES** | 80 hours a month
Outcome Delivery **TECHNICAL LEAD**
✓ **Build** Enhancements and new capability
✓ **Run** Administration, defects and support
✓ **Prove** Hours tracked, reviewed before acceptance



Your team

Owns the priorities and signs off every cycle

Crossfuze Advisory

Oliver Nowak leads the AI Action Plan and each quarterly reset

Hybrid delivery pod

A dedicated onshore Technical Lead leads 80 hours of delivery a month

You own the priorities. We ship the outcomes.

Indicative structure. Scope, workshop numbers and hours are confirmed in your statement of work.

Autonomy is earned.

Start by knowing where you stand.

Book a free 90-minute AI Readiness Assessment. Leave with your readiness score, expected Assist usage and the first outcome worth prioritising.

You decide what comes next.

Book your 90-minute assessment →

Open the tools →

90 minutes · No cost · Clear next steps

Companion guide: **The Road to Zero**, five outcomes for Autonomous IT on ServiceNow.



crossfuze.com/campaigns/servicenow-ai-readiness-assessment

© 2026 Crossfuze · October 2026

ServiceNow, Now Assist and related product names are trademarks of ServiceNow, Inc. Crossfuze is a ServiceNow Elite Partner.