

cognizant®

From execution to orchestration

Designing work for the age of adaptive agents

Abhijeet Singh
Associate Director, IPA Consulting
Cognizant | 2026

A new kind of resource

Automation was built for structured work: stable inputs, explicit rules and repeatable actions. Agentic AI reaches into a different class of work, where inputs vary, rules are incomplete and judgment often determines what happens next.

Peter Drucker recognized the underlying problem decades ago. Knowledge work does not decompose as neatly as manual work. It depends on interpretation, context and judgment. Agents can now operate inside that space in ways earlier automation could not.

That changes the resource available to the operating model. A bot follows a rule. A person interprets a situation. An agent can interpret context, choose a path, use tools and adjust when conditions change. The result is greater flexibility but also less predictability. Similar inputs can produce different outputs because the context or reasoning differs.

The economics change as well. Bot cost is tied largely to runs, licenses and maintenance.

Human cost is tied to time. Agent cost depends more directly on the retrieval, reasoning and tool use required by the work. Complexity starts to matter in a new way.

A new resource therefore requires more than a new technology layer. It changes how work should be designed.

The execution model and its ceiling

The operating logic behind most automation can be described simply: decompose the work, assign each task to an executable resource and govern by exception. We call this the execution model.

It works well when work can be structured in advance. Its ceiling appears when interpretation is required.

Consider an email-triggered workflow. The downstream process may be straightforward: extract data, match records, update a system and log the result. The difficulty often sits at the front. The email is unstructured. Requests are phrased differently. Attachments vary. A person has to read, interpret, structure and decide what happens next before the automated process can begin.

The structured tail has been automated. The judgment-heavy front has not.

That boundary matters because much of knowledge work sits above it: decisions, coordination, interpretation and context-dependent routing. The constraint shifts from task throughput to the speed and quality of judgment.

Agents can operate in this space, but inserting them into an operating model designed for execution creates a new problem.



The design gap

Many AI programs begin at the technology layer. A capability is deployed, a pilot performs well and the expectation is that business outcomes will follow.

Evidence cited in the original draft points to a persistent gap between pilots and production. One 2025 report attributed to MIT research put measurable production impact at roughly five percent. A March 2026 survey reported 78 percent of enterprises with agentic AI pilots but 14 percent at production scale. Gartner has projected that more than 40 percent of agentic AI projects will be canceled by the end of 2027 because of cost, unclear value and inadequate controls. [1][2][3] The operating model issue becomes visible when agents move beyond the demo. Accountability becomes harder because an interpretation layer sits between input and action. Governance by exception becomes less reliable because agent behavior is not simply compliant or non-compliant with a fixed rule. Cost also becomes less deterministic because reasoning effort varies with the work.

The technology may perform exactly as designed, while the surrounding system remains unchanged. That is the design gap.

From execution to orchestration

Orchestration is the operating logic required when work is performed by a changing mix of people, agents, bots, data and tools.

Execution asks: What task should this resource perform? Orchestration asks a broader question: What combination of resources, decisions, controls and handoffs will produce the outcome well?

An orchestrated process does not assume that every activity can be decomposed into a fixed sequence before work begins. Some paths can remain deterministic. Others require interpretation. An agent may gather context, choose among approved tools, request human judgment or route the work differently as conditions change.

The orchestration layer coordinates those choices. It defines which resources can act, what authority they have, when work must escalate and how evidence is preserved.

Humans remain part of the design, as well as bots and conventional automation. Orchestration does not replace execution. It determines where execution is enough and where adaptive resources are needed.

From execution to orchestration

Execution model

1. Decompose work
2. Assign each task to an executable resource
3. Govern by exception

→

Orchestration model

1. Start with the outcome and process
2. Compose people, agents and bots around the work
3. Set decision rights, controls and handoffs
4. Adapt the path as context changes

Orchestration does not replace execution. It decides where fixed execution is enough and where adaptive resources are needed

Figure 1. The operating logic shifts from fixed execution to outcome-led orchestration.

The outcomes stack

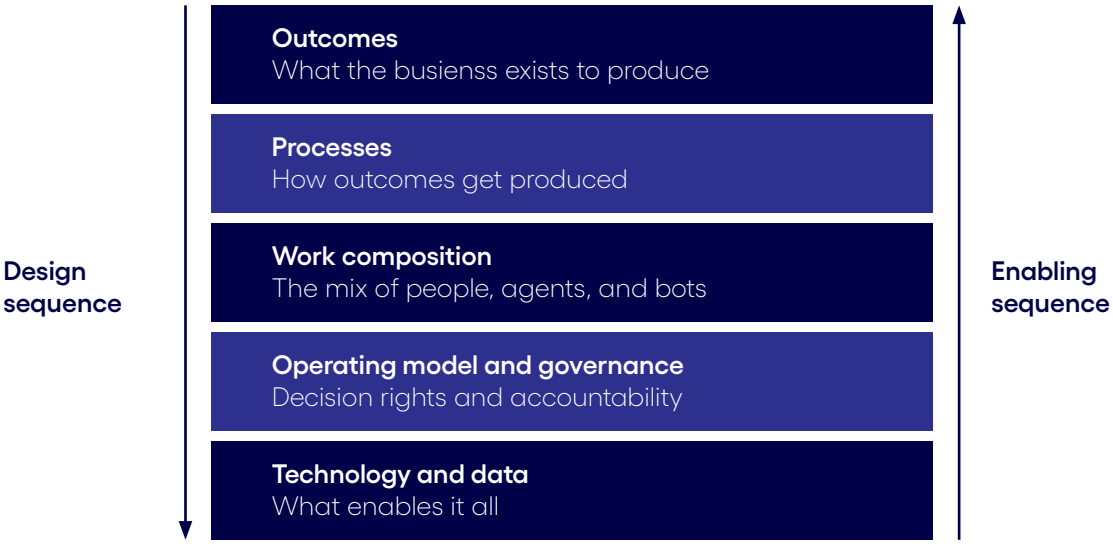
The outcomes stack provides a way to design that operating model from the top down.

Start with the outcome the business needs to produce, and then examine the process that creates it. From there, design the work composition: which parts belong with people, agents, bots or some combination. Governance defines decision rights, accountability, escalation and control. Technology and data enable the design.

The sequence matters. Many programs begin with technology and search for places to apply it. The stack starts with the outcome and works downward until the required technology becomes clear. Enablement then works back upward from technology and data through the operating model to the outcome.

This is where orchestration becomes practical. It is the discipline of designing the process, resource mix and governance together.

The outcomes stack



The outcomes stack: design top-down, enable bottom-up
Every layer necessary. no layer sufficient

Figure 2. The outcomes stack: design top-down, enable bottom-up.

Designing for orchestration

The shift from execution to orchestration changes the question leaders need to ask.

"What can this technology do?" is too narrow. The better starting point is: What outcome matters, how should the work be structured and what is the right mix of resources to produce it?

That question forces process design, work composition, governance and technology into the same conversation. It also makes accountability visible before agents are deployed at scale.

Agentic AI expands the range of work that can be handled by digital resources. The value will come from redesigning the system around that capability.

Execution remains part of the model, but orchestration is what makes the model adaptive.

Sources cited in the original draft

[1] Virtualization Review, "MIT Report Finds Most AI Business Investments Fail, Reveals GenAI Divide," 2025. https://virtualizationreview.com/articles/2025/08/19/mit-report-finds-most-ai-business-investments-fail-reveals-genai-divide.aspx

[2] Digital Applied, "AI Agent Scaling Gap: March 2026 Pilot to Production," 2026. https://www.digitalapplied.com/blog/ai-agent-scaling-gap-march-2026-pilot-to-production

[3] Gartner, "Gartner Predicts Over 40% of Agentic AI Projects Will Be Canceled by End of 2027," June 25, 2025. https://www.gartner.com/en/newsroom/press-releases/2025-06-25-gartner-predicts-over-40-percent-of-agentic-ai-projects-will-be-canceled-by-end-of-2027



Cognizant (Nasdaq: CTSH) is an AI Builder and technology services provider, bridging the gap between AI investment and enterprise value by building full-stack AI solutions for our clients. Our deep industry, process and engineering expertise enables us to build an organization's unique context into technology systems that amplify human potential, drive tangible outcomes and keep global enterprises ahead in a fast-changing world. See how at cognizant.ai or @cognizant.

World Headquarters
300 Frank W. Burr Blvd.
Suite 36, 6th Floor
Teaneck, NJ 07666 USA
Phone: +1 201 801 0233

European Headquarters
280 Bishopsgate
London
EC2M 4AG
England
Tel: +44 (01) 020 7297 7600

India Corporate Office
Siruseri-Software Technology Park of India (STPI)
SDB Block—Ground Floor North Wing
Plot No H4, SIPCOT IT Park
Chengalpattu District
Chennai 603103, Tamil Nadu
Tel: 1800 208 6999

APAC Headquarters
1 Fusionopolis Link,
Level 5 NEXUS@One-North,
North Tower Singapore 138542
Phone: + 65 6812 4000

© Copyright 2025–2027, Cognizant. All rights reserved. No part of this document may be reproduced, stored in a retrieval system, transmitted in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, without the express written permission of Cognizant. The information contained herein is subject to change without notice. All other trademarks mentioned herein are the property of their respective owners.

WF 5437001