cognizant® | aws

Accelerating enterprise AI at scale

From cloud-native foundations to agent-driven enterprise execution—designing, deploying and governing production-grade AI on AWS.



Contents

3 **A partnership for the AI era**
- Our joint value

4 **Why enterprise AI stalls**
- Understanding the challenges
- Why these challenges matter

5 **Cognizant–AWS integrated value**
- Cognizant's AI builder approach and accelerators
- Enterprise AI industrialization

7 **Enterprise agentification framework**
- Cognizant® Agent Foundry
- Building blocks of the enterprise agentification framework
- AWS-hosted accelerators portfolio

10 **Industry-specific solutions**
- Banking & insurance
- Travel, hospitality and media
- Case studies

16 **Engagement model**
- Enterprise AI journey
- Getting started

17 **Conclusion**

A partnership for the AI era

Engineering enterprise AI from strategy to industrialization

Artificial Intelligence is rapidly becoming an enterprise capability that extends beyond automation to enable intelligent decision-making, autonomous operations and business transformation. Realizing this potential requires more than powerful models—it requires production-grade engineering, responsible governance and a scalable cloud foundation.

Built on more than a decade of collaboration, Cognizant and AWS combine enterprise engineering expertise with AWS-native AI innovation to help organizations operationalize generative AI and agentic AI securely, responsibly and at scale—from strategy and design to deployment, governance and ongoing operations.



Proven together, ready to scale—your path to agentic AI

15+ years of strategic collaboration

560+ AI and analytics engagements

8,700+ trained AI professionals

100+ AWS customer launches

19,000+ AWS skills completed

13 AWS competencies

Why enterprise AI stalls

Understanding the recurring organizational and technology barriers preventing AI from scaling beyond successful pilots to enterprise-wide adoption

Pilot purgatory

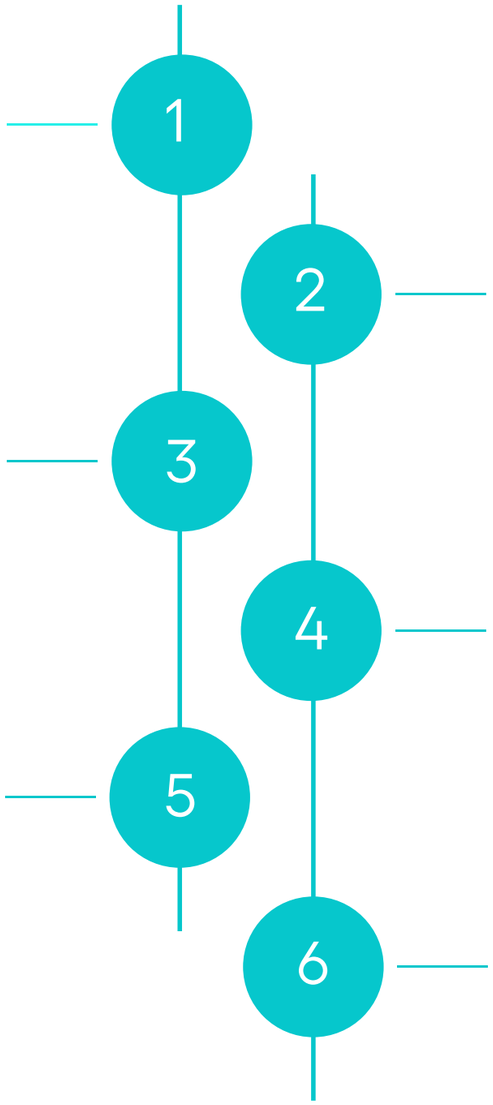
Many AI initiatives remain confined to experimentation/ proof-of-concept environments

Talent shortage

Demand for AI and machine learning expertise continues to outpace available skills, making it difficult to scale initiatives across the organization

Generic AI

General-purpose AI models lack enterprise context and domain understanding to address industry-specific business requirements effectively



Governance gap

Limited governance, risk management and intellectual property guardrails to support responsible AI adoption at scale

Fragmented stack

Siloed data and disconnected tools and platforms limit the AI systems from accessing complete organizational context to realize its full potential

Unclear RoI

Difficulty in converting proofs-of-concept into measurable business outcomes that demonstrate quantifiable business value

Why these challenges matter

Delayed enterprise-wide AI adoption

Increased implementation complexity and cost

Reduced confidence in AI initiatives

Slow transition from pilot to production

The Cognizant–AWS integrated value

A joint approach to building, deploying and scaling enterprise AI

Enterprise AI requires more than individual products or services. Cognizant and AWS address these requirements through a joint approach that supports every stage of the enterprise AI lifecycle. Core components include

Cognizant's AI builder approach & accelerators

Full-stack IP & AI Factory - proprietary platform stack plus enterprise-grade AI engine

- **BASIS™ framework** - consulting methodology to move pilots to scale
- **Cognizant® Agent Foundry** - agent design, build, deploy
- **Neuro® AI multi-agent accelerator & Neuro® AI engineering** - gen AI platform with multi-agent orchestration
- **Cognizant® Agent Foundry Composer & Ops, Cognizant Ignition™ Ontology** - agent grounding, QA, observability, FinOps, data readiness
- **Bluebolt** - rapid prototyping
- **Context engineering core** - captures enterprise context to ground gen AI
- **Industry solutions** - TriZetto, Duck Creek/Guidewire, Belcan
- **ISO/IEC 42001 governance** - accredited responsible AI and Trust™ framework

The AWS Cloud-native AI services

Comprise **foundation models, enterprise infrastructure, security capabilities** and **dedicated engineering support** required to deploy AI at scale

- **Amazon Bedrock** - agents, guardrails, knowledge bases, RAG
- **Amazon Bedrock AgentCore** - production AI agents at scale
- **SageMaker + Amazon Q** - business, developer, GenBI
- **FM catalog** - Anthropic, Amazon, Meta, Mistral models
- **Built-in security** - observability, guardrails, evaluation
- **Funded PoCs** - marketplace, ISV Accelerate, co-sell motion
- **Dedicated solution architects** - co-engineering support across the build
- **Global infrastructure** - scalable, secure, region-resilient cloud

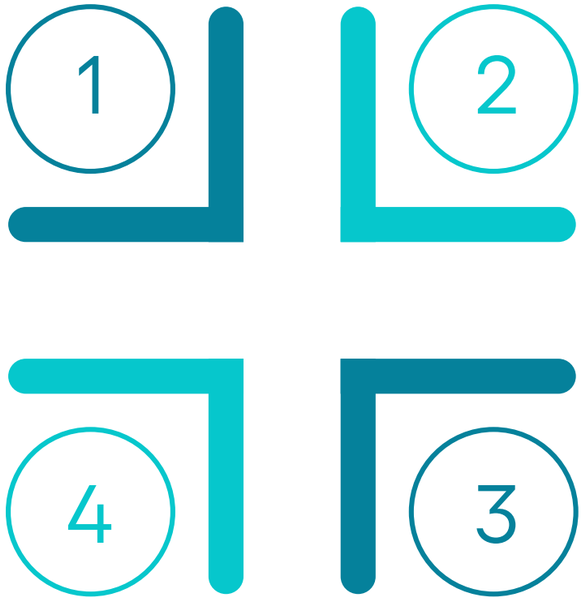
Together, these complementary strengths create a delivery model that combines enterprise engineering with AWS-native innovation, enabling organizations to accelerate AI adoption while maintaining governance, scalability and measurable business outcomes.

Enterprise AI industrialization

The guiding principles that shape every stage of the Cognizant–AWS AI engagement

Joint co-engineering

Cognizant and AWS teams work together throughout the AI lifecycle to design, build and deploy enterprise AI solutions, combining engineering expertise with AWS-native innovation.



Industrialized delivery

Reusable/repeatable implementation patterns, factory-grade engineering practices and enterprise accelerators help organizations move from pilots to production at speed, more efficiently and consistently.

Responsible AI by design

Governance, compliance and security guardrails are embedded throughout the implementation lifecycle, enabling organizations to deploy AI solutions responsibly from the outset.

Outcome-based delivery

Funded proof-of-concept initiatives and value-driven engagement models help organizations accelerate implementation while maintaining a focus on measurable business outcomes.

These guiding principles are further reinforced through a portfolio of proprietary frameworks, accelerators and AWS-native integrations that help organizations move seamlessly from AI ambition to production-grade execution.

Cognizant® Agent Foundry - enterprise agentification framework

An integrated approach for building, deploying and scaling enterprise AI

As organizations embrace Agentic AI, the challenge extends beyond building intelligent agents. Unlike individual AI applications, Cognizant® Agent Foundry provides a reusable framework that supports enterprise-wide agentification.

- Built on years of joint engineering with AWS
- Takes governed, autonomous agents from prototype to production
- Native integration with Amazon Bedrock and AWS infrastructure
- Combined with built-in observability, guardrails and human-in-the-loop controls



The building blocks of the enterprise agentification framework

Cognizant® Agent Foundry brings together a unified framework of six integrated capabilities for designing, developing and operating production-grade AI agents.

Capability	Role in the framework
AI-first reimagined process design	Contextual, iterative flow modeling with right tool/API integrations. Built for autonomous process design at scale—replacing static workflows with dynamic, context-aware orchestration
Reusable templates, libraries and modules	Built on the Cognizant Neuro® AI multi-agent accelerator and Cognizant® Agent Foundry composer and ops accelerators—deeply integrated with the AWS stack: Bedrock, SageMaker, Lambda, OpenSearch & Knowledge Bases
Industry and horizontal solutions	Verticalized agents (wealth advisor, underwriter, claims) and horizontal agents (customer service, coding, testing)—covering the breadth of enterprise domains
Pre-built agent catalog and marketplace	A reusable agent library blending deep domain knowledge with agentic technology—democratizing access to production-ready agents across the enterprise
Persona-driven user interface	Multi-persona UX empowering users to interact naturally with data — voice, conversation and goal-solver interfaces tailored to each role and workflow
AWS foundation layer	Integration and interoperability, governance and compliance, AgentOps observability, security guardrails — all powered by AWS infrastructure and Bedrock-native services

Together, these six capabilities create a structured framework that helps organizations:

- **Reimagine business processes for AI**
- **Accelerate development through reusable assets**
- **Leverage industry-ready AI capabilities**
- **Improve user adoption with persona-driven experiences**
- **Operate AI within an AWS-native governance framework**

Supporting these framework capabilities is a portfolio of AWS-hosted accelerators that help enterprises simplify data readiness, industrialize AI solutions and govern AI operations at scale.

AWS-hosted accelerators

Accelerating enterprise AI through reusable assets built for data, deployment and operations

Complementing the Cognizant® Agent Foundry is a portfolio of AWS-hosted accelerators designed to simplify key stages of the enterprise AI lifecycle—prepare enterprise data, industrialize AI solutions and operate AI environments with greater efficiency and governance.

Acclerator portfolio

Cognizant® Agent Marketplace

A persona-driven catalog of pre-built agents

- Curated **industry** and **horizontal agents**
- **Persona-driven UX** for rapid adoption
- Democratizes **agentic AI assets**

- **200+ agents ready to deploy, enabling faster deployment and wider adoption**

Cognizant Ignition™ Ontology

Creates an AI-first data foundation by

- Connecting enterprise data **in place** – zero migration
- Building **metadata and ontology**
- Powering **graph-RAG and NLQ-to-SQL** capabilities

- **> 50% reduction in onboarding effort**
- **> 85% insight accuracy**

Cognizant® Agent Foundry Composer

Vibe-to-verified industrialization

- Converts proofs-of-concept to enterprise-grade agents
- Built-in grounding & HITL approvals
- Reusable SKILL.md playbooks

- **Accelerated PoC to production journey**

Cognizant® Agent Foundry Ops

Operate and govern agent fleets

- Single command center for agents
- Self-healing & auto-remediation
- Built-in FinOps · MLOps · LLMOps

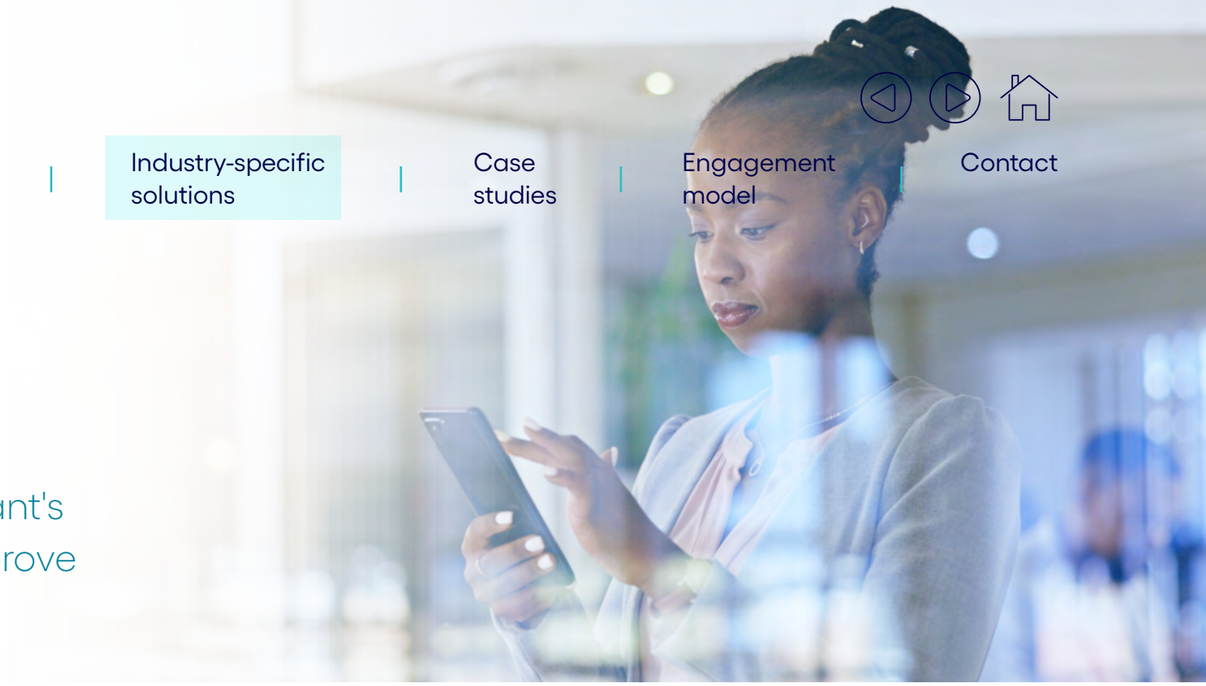
- **Real-time anomaly detection through centralized command centre for enterprise AI operations**

Together, these accelerators help reduce implementation effort while improving data readiness, operational governance and reuse across enterprise AI initiatives.

Industry-specific solutions built for enterprise AI

Pre-built, AWS-hosted models designed to accelerate deployment across industries





Banking & insurance

Banking and insurance organizations manage complex, document-intensive and highly regulated business processes. Cognizant's industry solutions combine agentic AI with AWS-native services to improve decision-making, automate workflows and accelerate operational efficiency across these environments.

Banking

Fraud and dispute management

An agentic operations workflow that accelerates fraud triage, dispute resolution and case investigation by orchestrating enterprise data, decisions and analyst actions across the dispute lifecycle.

RoI levers

- **Less investigation effort and faster case closure**
- **Better detection accuracy and lower operational losses**

Trade exception management

An agentic assistant orchestrating trade exception triage and resolution across complex post-trade workflows—reducing backlogs and accelerating resolution while maintaining strict compliance adherence.

- **Fewer backlogs and faster resolution cycles**
- **Compliance adherence with full audit traceability**

Insurance

L and A underwriting assistant

An agentic workflow on AWS Bedrock that assists underwriters with risk assessment, policy decisioning and document-driven reasoning —accelerating cycle times and improving consistency across complex life and annuity cases.

- **Faster cycle time and better decisioning consistency**
- **Less rework and higher straight-through processing rates**

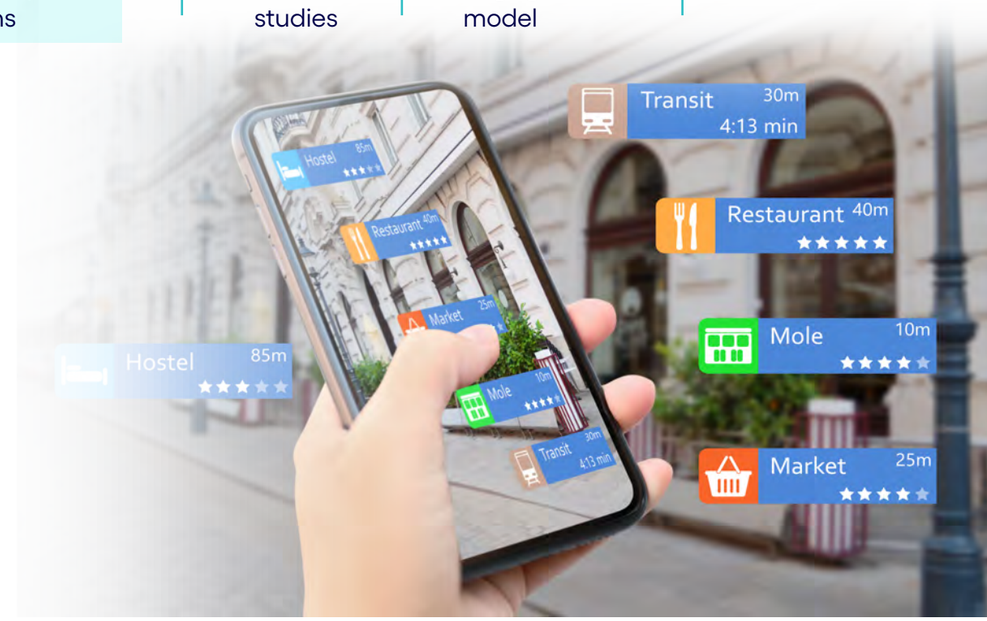
Intelligent document processing

Gen AI-powered IDP for the document-heavy insurance industry—automated extraction across PDFs, forms, tables and images using AWS Textract and Bedrock to reduce manual review and accelerate downstream processes.

- **Drive efficiency and faster access to insights**
- **Increase accuracy and improve downstream**

Travel, hospitality and media

Customer engagement and content creation require fast, contextual and highly personalized experiences. Cognizant's industry accelerators combine agentic AI with AWS services to automate customer interactions and improve CX while simplifying AI-powered content generation.



Travel and hospitality

Concierge AI companion

A Bedrock agents-based accelerator providing seamless cross-ecosystem concierge assistance — autonomous travel bookings, dining reservations and car rental guidance across the client's ecosystem with natural-language interactions.

RoI levers

- **Reduced operational costs via automated travel assistance**
- **Enhanced user experience with higher 'First Time'**

Media and communications

AI TechCast platform

An AI-driven podcast generator turning topics into scripts with synchronized slides, human-quality voice-overs and speaker visualization — built on LangChain and AWS for rapid knowledge dissemination.

- **Simplified content creation via AI workflows**
- **Improved engagement and multi-format content**

Customer success - enterprise AI in action

Improving operational efficiency

Using AI to simplify knowledge-intensive and operationally complex business processes while reducing manual effort



Gen AI customer support agent

Industry: leading healthcare and agriculture company

A **self-service support experience** resolving user-guidance issues from internal knowledge sources—deployed on a scalable AWS architecture with semantic retrieval.

20% reduction in incidents

15% cost savings

53% faster than human agents



Agentic AI for mortgage guideline agent

Industry: leading financial services provider

LangGraph-orchestrated multi-agent system delivering instant access to 10,000+ mortgage guideline sections across 5 major sources—with vector-powered semantic search.

85% less manual research

60% faster query performance

10,000+ searchable guideline sections



Gen AI intelligent complaint processing

Industry: leading healthcare manufacturer

Real-time gen AI triage of tele-call complaints—automating classification, duplicate detection and damage-code tagging at scale on AWS.

45% reduction in triage effort

€10K savings per day

~97% classification accuracy

Improving analysis and policy development

Enabling better decisions

Beyond automation, enterprise AI is enabling organizations to process information more effectively, improve analytical consistency and accelerate complex decision-making.



Gen AI-led reimbursement policy and payment integrity

Industry: leading healthcare payer

Gen AI-powered research and drafting automating market scan and evidence extraction — improving policy creation across multiple payers in parallel.

$1.3M annual value

20x parallel payer processing

20% policy improvements



Agentic AI-led earnings analysis and reporting agent

Industry: leading P and C insurer

A ReAct-based autonomous agent that automates financial statement analysis and analyst-call preparation by leveraging AWS Lambda and Amazon Bedrock—continuously improves output quality through reflective reasoning, enabling faster, consistent earnings insights and reporting.

95–98% positive feedback on generated outputs

24% reduction in drafting time

Direct alignment with Q4 earnings questions



Agentic AI-based contract intelligence platform

Industry: leading biopharma enterprise

A **production-grade agentic platform** automating multilingual contract data extraction with a dual-layer Critic Agent architecture and Corrective RAG for accuracy.

30–40% less review time

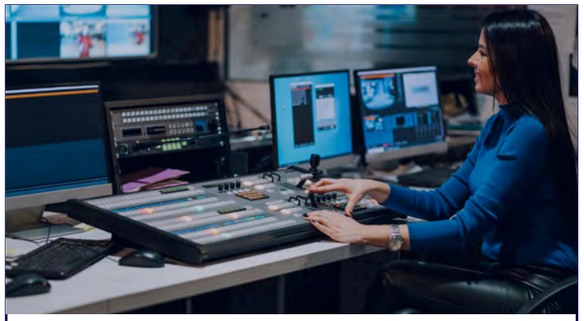
>88% accuracy on 40+ fields

Critic agent - dual-layer validation

Accelerating autonomous operations

Intelligent assistance to autonomous execution

Extending enterprise to real-world production deployments demonstrates the growing role of agentic AI in automating workflows, generating business artefacts and supporting autonomous operational processes.



Requirements data sheet generator

Industry: leading US digital TV and media company

Multi-agent AI assistant uses Amazon Bedrock with Claude Sonnet and OpenSearch to generate use-case requirement sheets across the 'Offers' value stream—collapsing time-to-market from weeks to minutes.

Reduced time-to-market from 3–4 weeks to

15 minutes

~90% faster than manual processes



Agentic AI-led Project Falcon—self-healing SFTP

Industry: leading EU electric utility

Agentic RCA & auto-remediation —agent analyzes MaaPS and Elasticsearch logs to detect failures, infer root cause and apply code-fixes, closing the loop.

Multi-level root cause analysis

Automated remediation (Heal + code-fix)

Reduced mean time to resolution (MTTR) - prevents repeat incidents

From opportunity to enterprise AI at scale

A structured engagement model for moving from AI discovery to full scale adoption

The Cognizant–AWS engagement model provides **a four-stage approach that helps organizations progressively build, scale and operationalize generative and agentic AI** capabilities while maintaining governance and operational readiness.



The enterprise AI journey

	Discover and define	Pilot and prove value	Industrialize at scale	Operate
Focus	**Identify opportunities**	**Validate business value**	**Production deployment**	**Continuous improvement**
Key activities	AI opportunity identification, use-case prioritization and AWS reference architectures	Rapid PoCs using Amazon Bedrock, SageMaker and Cognizant® Agent Foundry accelerators	MLOps, AgentOps, governance, security and operating model	Monitoring, optimization and evolution of AI and agent ecosystems

Each stage builds on the previous one, enabling organizations to validate business value before expanding AI initiatives across the enterprise.

Getting started

Turning enterprise AI priorities into production initiatives

Successful enterprise AI initiatives begin with clearly defined business priorities, a scalable technology foundation and an implementation approach aligned to organizational goals. Cognizant and AWS work together to help organizations establish this foundation through a structured engagement model:

Prioritize

Identify **two or three high-value gen AI or agentic AI use cases** aligned to business priorities.

Collaborate

Conduct a joint **Cognizant–AWS discovery workshop** to refine opportunities and align the solution approach.

Validate

Launch **a rapid pilot** using Cognizant® Agent Foundry together with AWS-native accelerators.

Scale

Industrialize at scale using **Cognizant® Agent Foundry Composer and Ops** and enterprise governance capabilities.

Conclusion

The future of enterprise AI is autonomous

Enterprise AI is no longer defined by experimentation—it is defined by execution. Together, Cognizant and AWS provide the capabilities, engineering discipline and governance required to operationalize gen AI and agentic AI at enterprise scale. By helping organizations move confidently from pilots to production, they lay the foundation for more intelligent, autonomous and adaptive enterprises.

Accelerating enterprise AI at scale

From cloud-native foundations to agent-driven enterprise execution—designing, deploying and governing production-grade AI on AWS

cognizant®

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
Phone: +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

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

© 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 here in are the property of their respective owners.