

AI Agents Enter the Workforce:

Takeaways from NiCE World 2026



 **opusresearch**

June 2026

Amy Stapleton, Senior Analyst, Opus Research

Ian Jacobs, VP & Lead Analyst, Opus Research

Derek Top, Research Director & Principal Analyst, Opus Research

Opus Research, Inc.
893 Hague Ave.
Saint Paul, MN 55104

www.opusresearch.net

Published June 2026 © Opus Research, Inc. All rights reserved.



AI Agents Enter the Workforce: Takeaways from **NiCE World 2026**

As NiCE World 2026 wrapped at Walt Disney World Resort last week, it became increasingly clear the company isn't just bolting AI agents onto CXone, it's trying to build something underneath the whole portfolio, a kind of operating layer for agentic CX. As CEO Scott Russell said in his opening keynote, it's all about orchestrating intelligence.

Imagine a system that watches customer interactions, spots friction, notices where automation could help, and then actually builds the fix. It can generate a new AI agent or tweak a workflow, test it, keep it governed, and hand the tricky cases back to a human when judgment or sign-off is needed. That's the ambition, and it's a lot bigger than a feature list.

Agentic Analytics and Proactive Service

For years, analytics in this space has mostly explained what already happened. NiCE is changing that. Now analytics runs quietly in the background looking for anomalies like friction points or sentiment dips that keep coming up. Instead of waiting for someone to dig through a dashboard, it surfaces recommendations on its own.

Here's where it gets interesting. Some of those recommendations can turn into automation. The system can reason through what an AI agent would need to handle a given intent and then spin up that agent so you can test and refine it. This automatic AI agent creation includes linking it to the knowledge, tools and workflow logic required for it to complete the full task. Analytics stops being a rearview mirror and starts becoming a discovery engine. It also provides the tools to move customer service from being reactive to proactive, addressing issues before the customer ever needs to reach out or help.

Cognigy isn't sitting off to the side as a recently acquired product. Cognigy AI is built right into Agentic Analytics, which shows NiCE is threading it through the intelligence, agent-building, and orchestration parts of the portfolio. If that's accurate, NiCE Cognigy is central to the shift from analytics and routing toward continuously finding and running automations that uplift CX.

NiCE's Wider Cast of Characters for Orchestrating Intelligence

There's more to the architecture than analytics: Agentic Experience Optimization, Guardian AI, Experience Memory, the Agentic Experience Plane, and an AI-first desktop all show up. A few are worth calling out:

- Guardian AI is the oversight layer for AI agents. It keeps an eye on things like brand alignment, agents drifting off script, containment slipping, and sentiment dropping, then suggests fixes a human can approve or wave off.
- Experience Memory builds a live record of each customer's journey by plugging into enterprise systems and tracking the events that actually matter to that customer.

- The Agentic Experience Plane is pitched as open and vendor-neutral, with room for third-party agents, dialers, and services through MCP (Model Context Protocol). Realistically, CXone customers will probably get the deepest native value, but the openness is the headline NiCE wants you to hear.

On the AI-first desktop, people aren't just working a queue of interactions. They're collaborating directly with AI agents, watching for where the agents need a hand, approving recommendations, stepping in when it counts, and leaning on journey context from Experience Memory to make sharper calls.

Broader Implications: The Hybrid Workforce Is Becoming the Operating Model

AI agents are moving from the edge of the contact center into the workforce model itself. They're still technology, of course. But NiCE, its customers, and partners increasingly described them as participants in the work, not just software sitting behind the scenes.

That framing requires a new way to think about AI agents. If they're truly becoming part of the workforce, the next CX stack can't treat them as isolated bots, copilots, or workflow add-ons. It needs a shared operating layer for all work performed on behalf of the customer, whether that work is done by a human agent, an AI assistant, an autonomous agent, or eventually a customer's own AI agent. Essentially, autonomy doesn't reduce the need for management. It changes what has to be managed.

It needs a shared operating layer for all work performed on behalf of the customer, whether that work is done by a human agent, an AI assistant, an autonomous agent, or eventually a customer's own AI agent.

That starts with a common view of work. The stack needs to know who or what is handling each step, what authority they have, what context they're using, what policy applies, and when a handoff is required. A copilot suggesting a refund, an autonomous agent approving a change, and a human agent handling an exception should all operate against the same customer state, knowledge sources, permissions, and audit trail.

It also means supervision has to expand. Managers will still coach people, but they'll also need to monitor AI behavior, evaluate outcomes, test new agent actions, approve policy exceptions, and see where automation is creating friction or risk. The stack will have to make AI performance visible in the same operational frame as human performance.

From Efficiency AI to Opportunity AI

Another clear theme from NiCE World was the shift from efficiency AI to opportunity AI. The first wave of CX AI was mostly about doing the same work faster or cheaper. That was perhaps necessary, but not sufficient for real transformation. A more interesting conversation at the event was about using AI to create outcomes that were difficult, expensive, or impossible to deliver before.

Arizona State University (ASU) was a great example of this. ASU wasn't just using AI to deflect calls or shorten interactions. It was using AI to predict where students might run into trouble, proactively reach out, and guide



them toward the next best action before they had to ask for help in the first place. In one example from the event, most contacted students acted and self-served based on the outbound contact; very few needed to call. Those are the first steps towards full-on opportunity AI, using agentic tools and CX infrastructure to expand what the organization can do for customers, students, patients, or citizens.

Those are the first steps towards full-on opportunity AI, using agentic tools and CX infrastructure to expand what the organization can do for customers, students, patients, or citizens.

Welcome to the Control Plane

NiCE didn't actually use Opus Research's AI Agent Control Plane framing, but much of what it introduced at the event points in that direction. The emerging story is not just about better bots or more capable copilots. It is about the operational layer needed to make AI agents useful, safe, measurable, and manageable at enterprise scale.

Quite a few of NiCE's announcements were in line with those ideas. For example, as mentioned above, the Agentic Experience Plane starts to address orchestration across human and AI work. Experience Memory speaks to the need for persistent journey and intent state, so every participant in the interaction has access to relevant context.

Guardian AI maps to the policy, guardrails, and continuous oversight layer. Agentic Analytics connects insight to action, moving analytics closer to operational control.

So, we see NiCE assembling many of the control-plane building blocks. This will create a clearer path for enterprises to govern, coordinate, and improve AI agents across channels, systems, vendors, and, eventually, customer-side AI agents.

Your Next Customer Might Not Be Human

Most of NiCE World was about putting agents to work for the enterprise, building them, governing them, folding them into the workforce. But what if your next customer isn't human? From the main stage, NICE Chief AI Officer Phil Heltweig showed a demo of a personal AI assistant rebooking a trip all while Phil was reclining in a comfy, La-Z-Boy chair. While the performance displayed the ease of AI agents working on our behalf, it also underscored an architectural vision of autonomous software systems and how contact centers need to be ready to serve it.

NiCE Cognigy's architecture does include MCP that runs in both directions. As a client, an AI agent consumes external tools (CRM, ERP, search, databases) through a single semantic interface. And as an MCP server, the architecture can expose its own agents and tools to other companies' AI agents, turning the contact center from a destination a human navigates into an endpoint another machine can call.

NiCE is framing the machine customer as just another channel, sitting alongside voice over WebRTC, traditional telephony, and digital. The platform is intended to help enterprises build agent-to-agent interactions additively, not by asking customers to re-architect.

Cognigy's dual position helps too. It's true that Cognigy is embedded in the core of CXone and is available CCaaS-agnostic on top of other platforms, which gives NiCE a path to both its own installed base and to customers who never standardized on it.

If a growing share of inbound volume arrives from agents rather than people, the platforms that can absorb it without linear cost growth win. Serving an agent is, in principle, cheaper and more deterministic than serving a person (structured requests, machine-readable responses) and none of the overhead a human conversation carries.

The Control Plane for Agent-to-Agent

With some predicting "your customers will all have shopping bots by next year" customer care organizations need to prepare for agent-to-agent interactions inside controlled environments: enterprise systems, partner ecosystems, and MCP-exposed tools where identity, authorization, and scope are known quantities. What's further out (and considerably harder) is the autonomous consumer agent transacting on a stranger's behalf across brands that have never authenticated it.

During the customer panel discussion, Nationwide made a point of saying they are "radically transparent" about disclosing when a customer is talking to an AI. Flip the interaction and the obligation flips with it: when the customer is an agent, how does the enterprise know who it is serving? Who authorized that agent, and who is accountable for the outcome? What does a CSAT score even mean when the respondent is a model?

None of this is governable without the AI agent control plane substrate (identity, policy, audit, and oversight) that has to sit underneath it. It will be important to treat the agent-facing endpoint as a first-class channel, get authentication and authorization right, and decide deliberately which tools you are willing to expose and to whom. The cost curve and the timeline make agent-to-agent interactions inside a contact center a question of not if but when. ■



About Opus Research

Opus Research is a diversified advisory and analysis firm providing critical insight on software and services that support digital transformation. Opus Research is focused on the merging of natural language understanding, machine learning, conversational AI, LLMs, conversational intelligence, intelligent authentication, and digital commerce.

For sales inquiries please e-mail info@opusresearch.net or call +1(415) 904-7666

This report shall be used solely for internal information purposes. Reproduction of this report without prior written permission is forbidden. Access to this report is limited to the license terms agreed to originally and any changes must be agreed upon in writing. The information contained herein has been obtained from sources believed to be reliable. However, Opus Research, Inc. accepts no responsibility whatsoever for the content or legality of the report. Opus Research, Inc. disclaims all warranties as to the accuracy, completeness or adequacy of such information. Further, Opus Research, Inc. shall have no liability for errors, omissions or inadequacies in the information contained herein or interpretations thereof. The opinions expressed herein may not necessarily coincide with the opinions and viewpoints of Opus Research, Inc. and are subject to change without notice.

Published June 2026 © Opus Research, Inc. All rights reserved.