How Salesforce Is Rethinking Customer Support and Success

A conversation with Sanjeev Balakrishnan on people, capacity, and proactive customer guidance.

Omid Razavi · September 28, 2026

A conversation with Sanjeev Balakrishnan at Dreamforce 2026

As AI takes on more support work, what should people and organizations do differently?



AI-generated depiction of Omid Razavi's conversation with Sanjeev Balakrishnan at Dreamforce 2026, based on photographs of them. Source: Omid Razavi.

That question ran through my latest conversation with Sanjeev Balakrishnan, Executive Vice President of Customer Success at Salesforce. Sanjeev also serves on the SuccessLab Advisory Board, and I'm grateful for the experience and thoughtful perspective he brings to our community.

Reconnecting at Dreamforce and other industry gatherings has become a regular opportunity to examine what is changing inside his organization.

In our 2024 conversation, we discussed scaling support through expertise, empathy, and preventive service. In May 2025, we explored how to prepare AI and people to operate together. At Dreamforce last year, we focused on Agentforce in production and the workforce around it.

© SuccessLab 2026 | Omid Razavi

This time, we discussed how AI is changing not just support work, but the organization around it. Sanjeev shared how he is involving employees in shaping that change, testing roles that bring support and proactive guidance together, and adapting to rising customer expectations. A central question ran through our discussion: how can leaders turn greater capacity into better customer outcomes?

The conversation has been edited, condensed, and organized by theme for clarity.

From Ideas to Working Agents

Omid: Since we last spoke, what has surprised you most about AI inside your organization?

Sanjeev: How quickly we can move from an idea to something that works.

During our second quarter, from May through July, we gave participating employees access to Claude and encouraged them to experiment. What information did they wish they had at their fingertips? What would help them serve customers better? Could they build a different dashboard or a better way to manage staffing?

We organized an innovation effort with a program manager and owners across our product teams. In three months, we received 216 ideas with working agent prototypes, not just suggestions.

"In three months, we received 216 ideas with working agent prototypes, not just suggestions."

People built workforce-management concepts, diagnostic tools, and new ways to visualize information.

We knew experimentation would cost, but I didn't want to constrain ideas before we understood what was possible. Now we are working with product teams to explore which of those ideas and lessons can become native capabilities.

Giving Employees a Role in the Change

Omid: Beyond efficiency, what did that change for your people?

Sanjeev: It helped people move beyond some of the fear around AI. They could try their own ideas and see what the technology allowed them to do.

As a leader, you can explain the opportunity repeatedly. Across a large organization, with different levels of experience, that message lands differently. Building something useful made the opportunity much more tangible.

What I am proudest of is the excitement and the beginner's mindset. Humans and these tools will be working together.

Turning Customer Context Into Better Guidance

Omid: Your responsibilities extend beyond support into adoption. How are you using AI to understand what help a customer actually needs?

© SuccessLab 2026 | Omid Razavi

Sanjeev: We built an adoption agent that looks at where a customer is in their Salesforce journey. It considers the features they have enabled, what they are using, their tenure, and their support history.

We organized features into foundational, intermediate, and advanced capabilities so the team can see where a customer might need guidance. The goal is to improve implementation health and help customers use what they have.

That context can change the intervention. If an established customer has repeated severe incidents or recurring performance problems, another support interaction may not be enough. Where their success plan includes it, we can bring in an architect to address the underlying issue.

Much of that information already existed. The opportunity is to bring it together so our people can make better recommendations.

Bringing Help Into the Customer Experience

Omid: Where else are you investing to make the experience more useful for customers?

Sanjeev: In-app assistance is a major priority. If a customer leaves the product to look for help elsewhere, we can lose the context needed to personalize support. Inside the application, we can better understand where they are stuck and recommend the right help, including onboarding where appropriate.

We are also investing in experience design. Finding the right content is one part of the problem. Having a helpful conversation across languages, cultures, voice, and text is another. That human side of the experience needs attention alongside the technical work.

How AI Is Changing the Support Engineer's Work

Omid: What is changing in the daily work of a human support engineer?

Sanjeev: I ask the team to look at three parts of the work.

First, what do we need to know about the case and the customer? That includes summarizing a prior digital conversation, identifying related issues, and classifying the case so it reaches the right expertise.

Second, diagnostics. AI can analyze attachments, logs, and images and surface relevant information. The engineer can inspect the source to validate it, rather than starting every investigation from scratch.

Third, resolution and documentation, including capturing the summary. We are seeing improvements across those stages, alongside capabilities that help bring in engineering or start a swarm when needed.

Meeting Customers Where They Work

Omid: As customers engage through more tools and channels, how do you decide where support should be available?

© SuccessLab 2026 | Omid Razavi

Sanjeev: Start with the customer and the industry. A developer using an AI coding tool may not want to leave it to visit a portal. A commerce customer may prefer a more visual interface or a messaging channel. We need to meet customers where they want to engage.

"We need to meet customers where they want to engage."

We have also seen an important operational shift. Chat now accounts for about 35 to 40 percent of our traffic. Digital agents do not sleep, and customers increasingly expect humans to be available too. We have had to adjust staffing around hourly arrival patterns.

Real-time translation across supported languages also helps us provide broader coverage and manage costs.

Understanding the True Cost of AI

Omid: Can you now trace the full cost of solving a customer's problem across AI interactions and human work?

Sanjeev: The initial experimentation period was about learning. We did not want to curb usage before seeing what people could create. Now we are working through the true cost.

Measuring Return Through Capacity and Redeployment

Omid: Where is the return clearest today, and what has it allowed your organization to do differently?

Sanjeev: Our biggest return has come from deploying Agentforce on the customer-facing Help experience. We went there before putting those capabilities into the hands of our human agents.

The important question is how much work would otherwise have reached a person. We also have to consider what traditional self-service was already handling.

I have barely hired in support over the past two years, even as we have introduced new products that need substantial support while they mature. Agentforce handles more of the work around mature products, which helps us redirect people toward those newer areas.

Some team members have also moved into forward-deployed engineering roles. That is part of how I think about the return: the capacity to support new products and redeploy expertise.

"That is part of how I think about the return: the capacity to support new products and redeploy expertise."

From Forward Deployment to Lasting Adoption

Omid: How do forward-deployed engineers fit into the customer journey, and who stays with the customer afterward?

Sanjeev: Our forward-deployed engineers work with selected customers to accelerate Agentforce adoption and time to value. They report into Sales, which determines where to bring them in. They are not assigned to every customer.

They then hand over to my organization, which provides the ongoing, scaled relationship through our success plans. The resources differ by plan and can include onboarding, success guides, architects, proactive monitoring, and named customer success managers.

Our responsibility continues beyond that initial engagement, through the customer's ongoing use and adoption.

Connecting Responsive Support With Proactive Success

Omid: You were also testing a Success Engineer role. What are you trying to change by combining support with proactive guidance?

Sanjeev: About half our support volume is how-to questions and half is break-fix. Over time, I want us to approach more of the how-to and adoption work proactively.

In Mexico City, we are piloting a Success Engineer role with about 200 people who previously worked in support, serving Latin America. It combines support expertise with the work of a success guide.

Why should it matter whether I call the customer or the customer calls me, if the issue and the expertise needed are the same?

"Why should it matter whether I call the customer or the customer calls me, if the issue and the expertise needed are the same?"

Traditional support is very transactional, partly because of the workload. At our volume, people move from case to case. The experiment is whether staying with the customer longer produces a better outcome. We are also looking at the economics of combining those responsibilities. It is a pilot, and those are the questions we are testing.

Aligning Measures With Customer Health

Omid: If the goal is a healthier customer journey, how far have performance measures moved beyond individual case targets?

Sanjeev: We are starting at the executive and senior leadership level by looking at overall customer health. We're also taking a team-based approach. We have not yet made that change to individual support incentives. We need more time and data to support it. Some roles, such as architects, are already measured on customer health rather than only transactional measures like CSAT.

Learning From the Conversations Themselves

Omid: When we meet again next year, what progress would you most like to show?

Sanjeev: Better use of the information in customer conversations. Chat and voice transcripts contain a wealth of data that can help us understand where customers are struggling and what they need next.

© SuccessLab 2026 | Omid Razavi

I am building a small command center in India to look across digital and human conversations. The vision is to recognize when an interaction is getting stuck and bring in the right help. That could mean bringing a human into a digital conversation that needs assistance.

By next Dreamforce, I want us to depend less on surveys to understand the experience.

That does not mean we will stop surveying customers. We can learn much more from the conversations themselves.

One of my biggest lessons from our experiments is that the customer information often already existed. It was buried or difficult to use.

"One of my biggest lessons from our experiments is that the customer information often already existed. It was buried or difficult to use."

Making it accessible can help us understand the customer and provide more personalized service.

Five leadership choices that shape AI's value

Here are a few takeaways from my conversation with Sanjeev.

1. Let employees shape the work.

The innovation effort brought employee expertise into working prototypes. It also gave people a role in shaping how AI changes their jobs. Participation becomes more meaningful when employees can help design the work, not simply adopt the tools.

2. Plan AI and human service together.

Always-on AI can raise expectations for human availability. Staffing, escalation, and language coverage need to evolve with the digital experience. The service promise must hold when a customer needs a person.

3. Turn customer context into better intervention.

Repeated incidents may call for architectural help rather than another isolated resolution. Connecting usage, adoption, and support history should change who intervenes, what they recommend, and whether the underlying cause gets addressed.

4. Match broader roles with broader measures.

The Success Engineer pilot tests whether one role can connect reactive support with proactive guidance. That broader purpose calls for measures and incentives that recognize both immediate resolution and longer-term customer progress.

5. Be precise about the return.

Greater capacity does not automatically mean lower spending. Supporting new products and redeploying experienced people can create value in other ways. Measure those benefits explicitly, alongside the technology and human effort required to achieve them.

Continuing the conversation

I'm grateful to Sanjeev for making time during Dreamforce, and for the candor and generosity he brings to our conversations. His willingness to share both progress and open questions gives our SuccessLab community a valuable opportunity to learn alongside him.

When we reconnect, I look forward to hearing what he and his team have learned. Are problems recurring less often? Is help arriving sooner? Can customers move from a technical fix to successful use without starting over with another team?

That is why I value this continuing dialogue with Sanjeev: we can revisit the ideas, understand what changed for customers, and learn from what followed.

Original LinkedIn publication