Support to Customer Leadership: The New Mandate for the AI Era

From Prevention to Verified Resolution

CCO PERSPECTIVES

SuccessLab Executive Roundtable | September 1, 2026
Hosted by Boston Consulting Group | Mountain View

Customer leaders convene at BCG Mountain View to examine the new mandate for AI-powered support.

We planned this one for BCG's boardroom.

By the time the invitations closed, the list had outgrown it, and our hosts opened a larger space. That keeps happening. On September 1, 50 senior Customer Service, Support, and Technical Success leaders spent several hours at BCG Mountain View working through what AI is doing to their profession. One keynote, one executive dialogue, one leadership panel, and three rapid-fire partner sessions. Dinner ran long, as it usually does. Some of the best conversations happened after the program ended, standing up and holding plates.

Somewhere in the middle we did a seventh-inning stretch and a group photo. Support executives threw their hands in the air on the count of three. I recommend it.

Leaders arrived with operating problems, success stories, lessons from what had gone wrong, and a willingness to share all three. That candor is the reason to convene this group.

Four executive shifts stood out for me. Together, they point to a larger opportunity for CCOs. Support leaders are increasingly positioned to take on broader customer leadership, especially when they connect service intelligence to customer outcomes, product decisions, retention, and growth.

THE CCO THESIS: CUSTOMER LEADERSHIP MOVES UPSTREAM

Omid Razavi opening the SuccessLab Executive Roundtable at BCG Mountain View.

I opened with five claims. They shape what I was listening for all night.

Customer leadership sits at an unusual intersection. We see customer needs, product reality, commercial commitments, and operating constraints at the same time. Almost no other role in the company sees all four at once. That perspective should not remain downstream of important decisions. Customer leaders should help determine what the company builds, where it invests, how it grows, and which outcomes it values.

For that to happen, customer evidence has to carry real authority. It earns that authority when it changes a priority, redirects an investment, alters a product decision, or changes how success is measured. If none of those happened last quarter, the voice of the customer in your company remains ceremonial.

Defining customer leadership for the agentic era takes a community. SuccessLab brings leaders from different roles, industries, and organizations together and builds continuity across meetings, so these questions can advance through shared experience.

The room earns its value through operating experience. What did you change? What worked? What did the economics show? What would you do differently?

Useful disagreement makes the work stronger.

The outcome I am after is influence. Greater weight for customer leaders in consequential decisions, broader opportunities to lead, and career paths that reflect the value of the work. That is how I think SuccessLab should be judged.

SHIFT 1: FROM AI PROJECTS TO AN OPERATING MODEL

Bryan Belmont, BCG advisor, formerly Microsoft CVP of Customer Service and Support

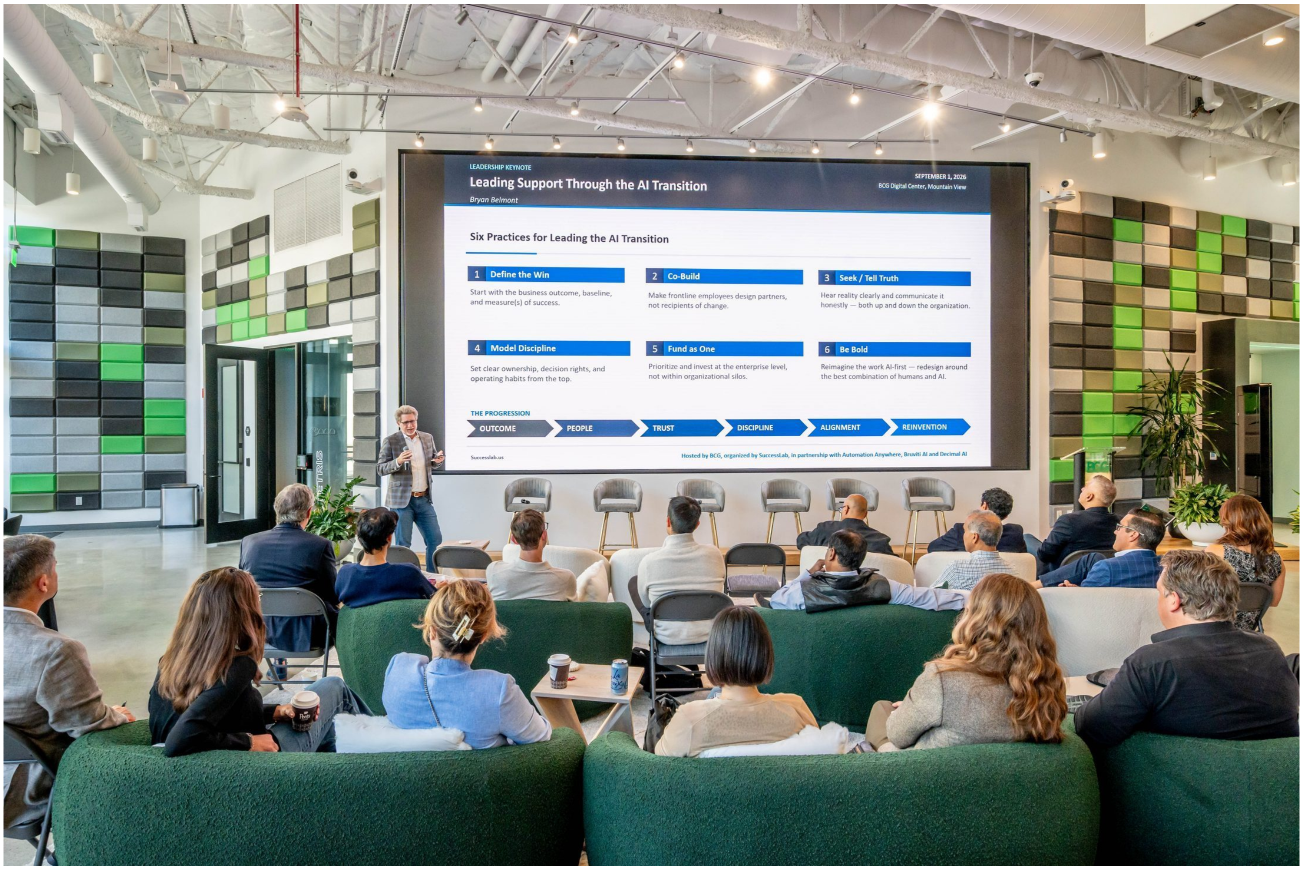

Bryan Belmont presents six practices for leading Support through the AI transition.

Bryan's team rolled out Copilot to 40,000 people across Microsoft's support organization before it was publicly available. Over the past two years, he has advised a wider range of companies through the same transition. He offered six practices that have become even more important in the AI era.

Define the win

He opened with The Hitchhiker's Guide to the Galaxy. A civilization spends seven and a half million years computing the meaning of life. The answer comes back as 42. Then they have to build a bigger machine to work out what the question was.

He described a client running 76 AI initiatives inside support. His first question was why 76. Why not 10, or 150. There was no answer. His verdict: plant a million seeds and hope a flower grows.

That is the pattern he keeps finding. Teams instrument the outcome and skip everything upstream of it. Do people understand what you asked them to do. Did they adopt it. Are they doing it consistently and correctly. Measure only the endpoint and the postmortem lands on the technology, which is rarely the problem.

Co-build with your people

The rollout began with a listening system. His team spent a week at each site, including outsourcer locations, sitting with engineers and asking what was broken. They changed the plan based on what they heard, then returned to explain what changed.

An Azure support engineer told him that every previous tool rollout made the job harder, and that this was the first time anyone had asked.

Simon Bamberger of BCG added the detail that makes this work. Bryan's team moved from listening to participation by seating frontline engineers on the transformation team, with protected time in the weekly meetings.

Bryan was careful to say this is additive. Frontline engineers know how to make today better. Leadership still owes the organization a view of where it needs to be in three years.

Seek truth and tell truth

This was the practice that drew the most conversation afterward, and it is the one worth dwelling on.

He described a client executive announcing that running all email through AI would cut half the headcount within a couple of months. The technical team present agreed on the spot. Everyone knew it would never happen. Nobody said so. What follows, in his description, is a kind of shared organizational delusion, where the plan of record has no relationship to what anyone believes.

He was equally direct about his own reviews at Microsoft. Every milestone met. Every status green. No movement in the outcome, quarter after quarter. He let that continue longer than he should have before stopping the program.

I have watched that same review, from the other side of the table, more than once in 25 years. The uncomfortable part is that nobody in it is lying. The milestones really were hit. The work really was done. Everyone present can see the outcome has not moved and nobody has the standing, or the appetite, to say the quiet part. It takes a senior person deciding that being wrong in public is cheaper than being wrong for another two quarters.

Bryan's framing is that both directions count. Telling hard truths upward, and telling them downward to your own organization. He described a leader who wanted to plan an outsourcing move without telling his team, and his advice was that the conversation gets three times harder later and costs all the trust you would have had.

Build the operating model

Bryan was once asked to assess a support organization with poor customer satisfaction. After three days with engineering, sales, offerings, and customer success, his finding was that support was fine. Everyone made decisions and no one made decisions. Budget and resourcing sat in silos and never lined up. There was no rhythm of business across the organizations that depended on each other.

The recommendation was a program office, shared metrics, and a real cadence. No process changes. No people changes. Results improved anyway.

He said the client was mildly annoyed that he had not recommended firing anyone in support. I believe it. When satisfaction is down, support is the easiest place to look and usually the wrong one.

Align the funding

Everyone agrees in the meeting. Then everyone goes back and budgets independently. Work that has to happen in one organization never gets funded because a different organization owns the outcome. At the end of the year there is a great deal of effort, a fair amount of good work, and no result.

Verbal alignment is free. Budget alignment is where you find out who actually agreed.

Be bold

The top few percent of software engineers see dramatic gains from AI. What separates them is that they reinvented how they work rather than adding AI to the job they already had. Bryan's argument is that the first five practices are what earn you the right to attempt that.

He put the question to his own team in the starkest terms available. What if there were no such thing as a case. How would you design support then. That question came back several times over the next two hours, which is the best thing you can say about a keynote.

A first-time attendee from ElevenLabs asked how companies escape the PowerPoint governance trap and start building. Bryan pointed to an active operating rhythm: accountable leaders, live data, and honest conversations about whether outcomes are moving.

SHIFT 2: FROM HEADCOUNT DEBATES TO CUSTOMER VALUE

Teresa Haun, Zendesk, and Atul Nanda, formerly Salesforce, Google Cloud, and Atlassian

Atul Nanda and Teresa Haun discuss how AI is reshaping the support workforce and customer demand.

The most thought-provoking hour of the evening was also the most useful because Teresa and Atul brought different evidence and reached different conclusions.

The demand picture

Teresa opened with a line I have been repeating since: no business has ever said it has enough customer support. There is a large stock of service debt in every company. Every moment a customer needed help, tried, and gave up.

The data Teresa presented, also shared by Adrian McDermott of Zendesk, complicated the standard narrative in several directions:

- Customer service job postings ran roughly 10 percentage points ahead of the overall market from about August 2025, narrowing in more recent data.
- Employment data tells a layered story. Call center employment softened in the US, Canada, and Germany while rising in the Philippines.
- Among customers automating around 60% of volume, total interaction volume doubled. Agent headcount rose about 3% while workload per agent rose roughly 90%.
- A Goldman Sachs comparison put an AI agent at about $93 per day against roughly $90 per day for an offshore human agent.

- In Zendesk's own customer base, 54% of customers contacted support meaningfully over a year. That group accounted for 94% of revenue and 95% of revenue growth.

That last figure is the one I would put on a slide for your CFO. Support talks disproportionately to the customers who generate the company's growth.

The operator's counterweight

Atul did something on stage that I rarely see. He said that a year ago he sat in a room and predicted 60% job loss in support, and that he was wrong.

He then went and checked his own claim. The Indeed hiring lab data shows support roles down about 10% since the start of the pandemic while other roles grew. Gartner research he cited found only about 20% of surveyed support leaders attributing reductions to AI. His read is that some of what gets announced as AI efficiency is ordinary cost reduction wearing a better story.

His broader point will be familiar to anyone who has run a support P&L. Support has always been asked to fund the rest of the company. That predates AI and will outlast it. His own projection is that meaningful reductions arrive later rather than never.

His advice for leaders was specific:

1. Get personally fluent. Build something yourself. You cannot set the pace for an organization on a capability you have not used.
2. Audit your org design. Move from region-based to product-based structure.
3. Hire data engineering and data science talent into support. He called it a priceless investment, and said the biggest wins his teams delivered came from having people who could build the story out of the data.
4. Shift outsourcing contracts toward outcomes rather than bodies and hours. He spent four days with four outsourcers in India last year having exactly that conversation.
5. Prove value in one important area with low risk. Do not wait for perfection. Do not start fifteen things at once.

I asked him a question that has been on my mind. He came from SAP on-premise to Salesforce, which was trailblazing SaaS support, then to Google Cloud, which was consumption-based. Each shift changed what support was for and how it was funded. We are at another one now, from consumption toward outcomes.

Atul's answer came back to organizational design, particularly the move from regional to product-based structures. I added that AI agents specializing by product can serve every time zone, weakening the logic behind the traditional follow-the-sun model. The org chart built around human working hours should now be examined.

The counter-evidence

Sonal Chablani reported the opposite pattern in her own organization. A 25% year-over-year decline in ticket volume, driven by more technically sophisticated customers who adopt AI and resolve issues before they escalate.

Both things are true. Which one you see depends on your product complexity and who your customers are. A leader with a consumer-facing product and a leader with a complex enterprise platform should expect different curves, and neither should assume the published research describes their situation.

Sanjeet Kaur Bali of Infor reframed the debate around talent density rather than headcount. The skill set is changing toward techno-functional, and support is becoming a talent factory for the rest of the company. That is closer to what I see in my advisory work than any of the hiring charts.

How the roles change

The most concrete artifact of the evening:

- Tier zero and tier one automate.
- Tier two and tier three become the VIP humans, for interactions where judgment, empathy, privacy, or compliance genuinely change the result.
- Team leads supervise AI and people together, deciding which work belongs where.
- Knowledge management, workforce management, quality, and operations converge into a role Teresa called the AI service architect.
- Support leaders are increasingly positioned to take on broader customer leadership, especially when they connect service intelligence to customer outcomes, product decisions, retention, and growth.

That expansion carries a larger mandate than interaction quality. It connects customer outcomes, advocacy, growth, and retention as those fields converge.

SHIFT 3: FROM DEFLECTION TO VERIFIED RESOLUTION

Moderated by Uche Monu, BCG. With Brian Lomeli, ServiceNow; Tanvir Kherada, Databricks; Sanjeet Kaur Bali, Infor; Avi Karnon, Cribl; Asad Bahou, Automation Anywhere

Executive leaders examine the path from prevention to verified resolution.

The panel moved from workforce to architecture, and this is where the practical detail lived.

Build it across functions from the start

Brian Lomeli gave the clearest answer of the evening on how to organize this work, building directly on the keynote. At ServiceNow, every AI effort brings product, IT, and support together from the beginning. This shared structure keeps the teams aligned around the same outcome.

He then named the part the keynote did not cover. Each of those organizations reports progress up a different management line, and each has an incentive to describe its own work favorably. Getting to the truth inside one organization is hard. Getting to it across three, when everyone's status deck is green for a different audience, is a different problem entirely. His answer is to agree the outcomes up front, in the same room, and then hold everyone to the same evidence.

He was equally direct about a risk he sees across the industry. Companies are harvesting the value before they create it. Headcount comes out first, on the strength of a projection, and the technology is expected to absorb the workload on a schedule nobody validated. If you have already let people go, he said, your AI had better take up that workforce, and he does not see it working as cleanly as the public reporting suggests.

Understand the demand before you route it

ServiceNow built what Brian called a funnel view of everything arriving at its front door, then works layer by layer to see where volume leaves the system. The goal is to understand why the demand exists, so the underlying issue can be solved outside the case entirely.

His concern is that most organizations are still fitting AI to a process designed for ticket queues. He returned to Bryan's question from the keynote about designing support as though the case did not exist, and made it operational. Where is demand actually coming from. Why. How do we route it to the right place, whether that is self-service, an expert, or a fix in the product. Then measure how much leaves the system at each layer.

He acknowledged that the answer remains imperfect and difficult, and that leaders across the industry are still working through it. That candor from a company building in this category carries more weight than a tidy framework.

What reliable resolution requires

Knowledge alone fails. Send the same answer to two customers reporting the same connectivity error and it will fail, because their environments differ. Reliable resolution requires customer context, problem context, and organizational memory of past resolutions, working together.

Inputs belong to the error class. Tanvir described defining which signals matter for each class of problem. Logs, configuration, prior tickets, customer code. Your product code is often the best-maintained documentation your company has.

Skill design is a frontline job. Tanvir's team builds resolution skills with the engineers who handle those cases daily, and he shared a warning worth repeating. A badly written skill description will pull the agent confidently down the wrong path. The agent picks the skill it was told about, not the skill that fits.

Support should start in the product. Avi argued that supportability is an investigation capability built into telemetry rather than documentation added afterward. His team is moving the tools they built into the product itself, and giving up some autonomy to do it.

Prove it internally first. Sanjeet described Infor's sequencing decision, which I found the most transferable idea of the panel. Make self-service work for your own engineers, with your own version and environment specificity, before you point it at customers. If it cannot satisfy the people who know the product, it will not satisfy the people who do not.

Escalation is a design decision

An orchestrator qualifies whether a case can be resolved autonomously. The response is scored against a configurable confidence threshold. Sentiment acts as a second trigger. Below threshold, a person takes it.

Automation Anywhere calls that handoff reassignment rather than escalation, and the case can pass back to the agent after a human turn or two. Every human action feeds the tuning set.

Infor made the opposite call for complex enterprise support and chose verified generation with a person in the loop over autonomy. Both leaders could explain exactly why, which made each position credible. Deliberate choices, grounded in customer and operating context, are what matter.

Who owns the outcome

Uche asked the panel who owns the outcome when a resolution fails. Brian said ownership spans Support, Product, and the cloud organization, with Support serving as the mouthpiece for what customers are experiencing and routing that signal to whoever can address the root cause. Sanjeet added that accountability remains with the people delivering the service, even when an AI agent is involved.

The cost question

The most useful question of the night came from Bryan Belmont, from the floor.

For twenty years, the deflection funnel worked because self-help had close to zero marginal cost. Write the article once, serve it a million times for free. That curve sits underneath every deflection target and headcount model since 2005.

Inference does not work that way. Cost accrues per interaction and grows with volume. On the Goldman comparison Teresa had shared an hour earlier, AI is already roughly the price of offshore labor, and today's token pricing is subsidized. Bryan's worry is that we are building something that falls over.

Sanjeet offered a practical hedge. Write every AI resolution back into your own knowledge base, every time. If licensing terms change, the asset stays where you control it.

Avi offered the counterargument, and it deserves equal weight. Cost optimization was never the first goal. Can we resolve issues three or four times faster. Can we serve customers better. Can we enable them. Inference costs will commoditize as vendors compete and models improve. Build the new support first and solve the cost layer as it arrives.

Both positions have merit, and the disagreement is useful. If you are quoting a three-year TCO to your CFO this quarter, you are quoting subsidized pricing. Move forward boldly while bringing that exposure into the renewal conversation early.

Evidence From The Field

Ten minutes each. Enough to show the work, not enough to pitch.

Automation Anywhere showed how specialized agents can work under an orchestrator, qualify cases for autonomous resolution, score confidence, and send uncertain cases to people. Each human intervention becomes part of the next tuning cycle.

Bruviti AI showed why complex service environments require more than strong answers. For equipment manufacturers, resolution often depends on field service, inventory, warranty, and installed-base telemetry. Their approach connects those operational systems through an ontology across parts, equipment, and components.

Decimal AI presented with Resilinc and brought measurable operating evidence. Their agents analyze code, logs, and data before handing engineers a diagnosis. Ravi reported resolution time falling from 6.5 days to 2 days. He also said Support was resolving roughly a third of issues, reducing the number sent directly to Engineering.

Across all three sessions, the same leadership lesson emerged. AI creates value when it connects customer context, technical evidence, organizational memory, and the systems required to complete the resolution.

SHIFT 4: FROM SUPPORT ECONOMICS TO ENTERPRISE INFLUENCE

Three conclusions crystallized, including one that changed my mind.

Lead with the value case

I have spent a good part of 25 years in rooms defending support budgets through cost per contact, deflection rates, and headcount ratios. Those measures have never won a strategic investment conversation. They position the support organization as an expense seeking approval.

The 94% figure is a different kind of argument. The customers who contact support are the customers who generate the growth. That reframes what deflection was ever for. One leader that evening said she is allergic to the word. She does not want to be deflected, and neither do her customers.

The economics under all of this deserve harder scrutiny than they are getting.

Model your inference costs at several times today's price and see what survives. If the answer is nothing, you have a strategy that depends on someone else's pricing decision. Then decide whether cost is even the right first question, as Avi argued it is not. Those two ideas are compatible. Build boldly, and know exactly what your exposure is.

Leadership determines the outcome

Bryan's six practices predate AI entirely, and every session that followed him arrived back at one of them from a different direction. Define the outcome. Measure the path. Bring frontline people inside the work. Tell the truth in both directions. Fund it as one thing.

That should be encouraging. Those are disciplines this profession already knows how to build.

Six questions to take back to your team

1. For each active initiative, what is the outcome, the baseline, and the measurement chain between them? And can you explain why you are running that number of initiatives rather than half or twice as many?
2. Which frontline engineers sit inside your transformation team with protected time, and when did you last go back and tell them what changed because of something they said?
3. Which program is green on every milestone and moving nothing? If product, IT, and support each report separately on this work, you may have three green decks and no shared outcome.
4. For your highest-volume error class, what inputs does the agent receive, what is the confidence threshold, and who decided both?
5. Where does every AI resolution get written back, and who owns that store if your licensing terms change?
6. Name one priority, investment, or product decision that changed last quarter because of customer evidence. If you cannot, that is the place to start.

Thank You

The SuccessLab community closes the evening with an all-hands-up group photo.

Thank you to BCG for hosting us, and to Bryan Belmont, Teresa Haun, Atul Nanda, Uche Monu, Simon Bamberger, and every panelist and partner who brought operating detail and candor to the conversation.

A particular welcome to the leaders joining us for the first time from Cresta, Decagon AI, ElevenLabs, and Sierra AI. Your questions sharpened the discussion and brought an important builder perspective into the room.

Most importantly, thank you to the customer leaders who contributed their experience and judgment. These gatherings become valuable when leaders share what they have tried, what they learned, and what they will do differently.

What Is Next

Explore past and upcoming SuccessLab events at successlab.us/events.

Read SuccessLab content, including the CCO Perspectives series, at successlab.us/content.

Customer leaders need a place to examine consequential changes with peers who understand the work. If you lead in one of those areas and want a seat, reach out.