

POV

Rethinking Application Support

Leveraging AI, data, and automation to transform the operating model

Author **Sinjan M Kashyap**



This paper envisions an AI-powered and data-centric operating model (referred to henceforth as APDC model) that includes a mix of human and virtual workforce for application support. The document discusses the overarching vision, benefits, building blocks, the model itself, as well as use cases.

A better way

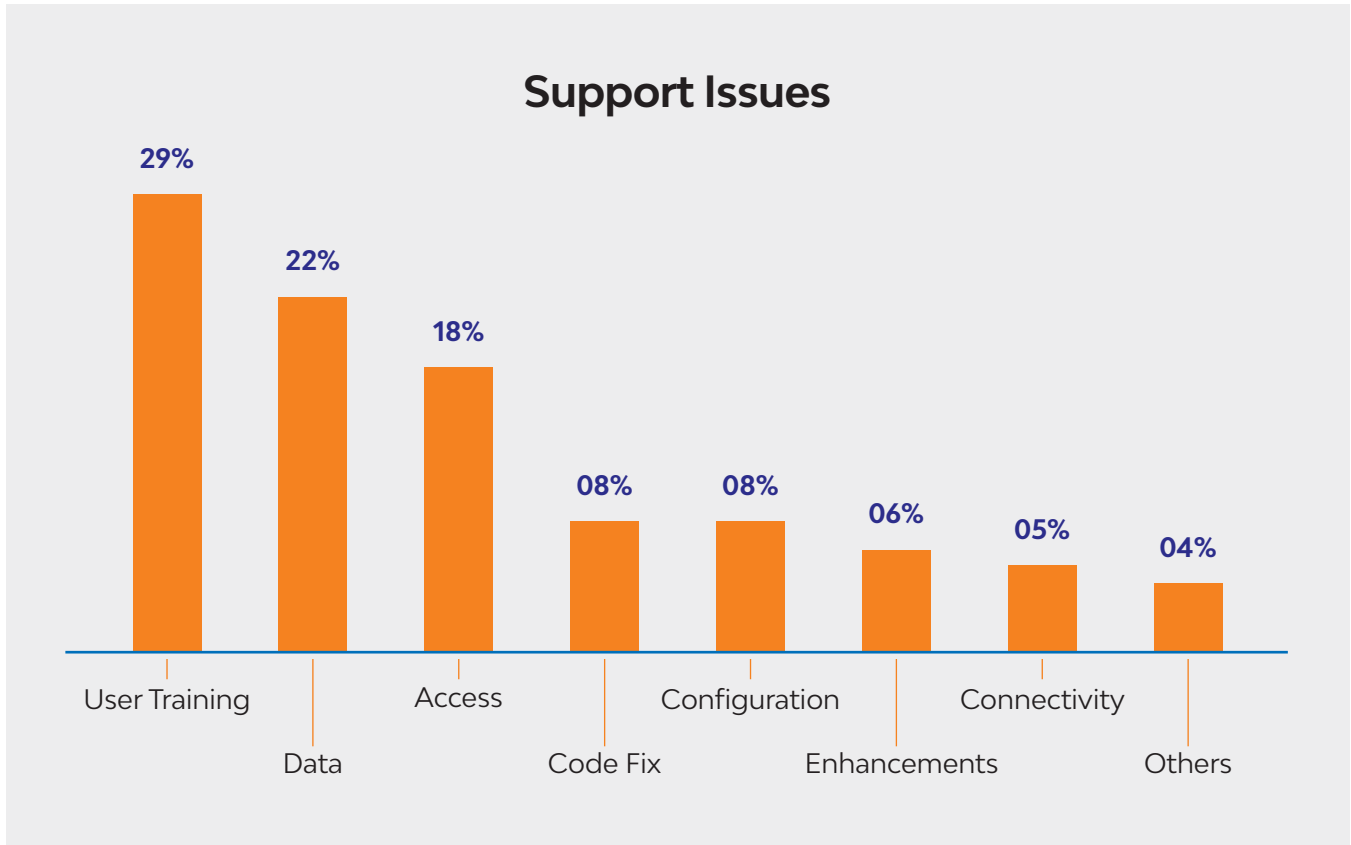
The inspiring idea in the 2018 book by HBS professors Karim Lakhani and Marco Iansiti, "Competing in the Age of AI," is an operating model where value is delivered absolutely or primarily by a virtual workforce. The operating model, in this case, refers to the way by which organizations provide value to the consumer. For example, most banks' operating model for processing a loan involves verifying the client credentials and credit score, arriving at a decision, disbursing the loan, and executing a system for collecting monthly installments. Similarly, the operating model for purchasing goods (say, groceries) is locating a store, browsing the inventory with guidance by sales, selection, and finally making the purchase.

Lakhani and Iansiti argue that companies such as Ant Financial and Amazon have fundamentally transformed these operating models with a data-powered virtual workforce—with a human crew operating at the edge (picking an oddly shaped item or dealing with an HNWI). This strategy offers tremendous scale, scope, and learning opportunities that traditional operating

models fail to capitalize on, thus, helping the Airbnbs and Netflices of the world to disrupt the worlds of the Marriotts and Blockbusters.

Closer home, a similar analysis of the operating model for production support reveals a host of activities such as incident management, service requests, and response to alerts or emails.

A breakdown of incidents over 90 days in a typical multi-application support desk would look like this: The user training-related issues (29%) would typically require sending a step-by-step instruction by mail and occasionally need deskside or conferencing support. Most data-related problems can be traced to data discrepancies that require managers to add or revise data (they also point to the existence of data silos that don't agree with each other). Access-related issues (18%) have a set protocol to be followed every time. Typically, 27% of problems require L2 or L3 support, with approximately 8% requiring a code fix. This is a general observation and may vary widely based on application type and maturity.



At the same time, the support desk is involved in a host of repetitive tasks such as ticket routing, follow-ups, monitoring (for high priority incidents), communication (with both users and higher-level support desks), updating work notes, SLA tracking, application health and jobs monitoring (via mailbox), and SOP creation. While some of these tasks require human intervention, most are repetitive in nature. The APDC model stands to augment the human workforce with a virtual workforce who can

process the manual, repetitive tasks while humans focus on the really complicated issues machines can't tackle. Before we jump into the APDC operating model though, we must recognize the pillars they stand on.

- 1. What gets measured gets managed.**
- 2. Why automate when you can eliminate?**
- 3. An AI-powered and data-centered operating model.**

1. What gets measured gets managed

One of the fundamental questions that rarely gets answered is: "What is the trend of ticket type? Are we getting more standard issues (such as access or training) or code fixes? Do we have excessive incidents, and are those related to connectivity, code, or integrations?" Most of the time, it's an intuitive rather than factual take limited by our memory. I have yet to come across a dashboard that says, "Application XYZ had 50% code fixes or 70% access-related issues".

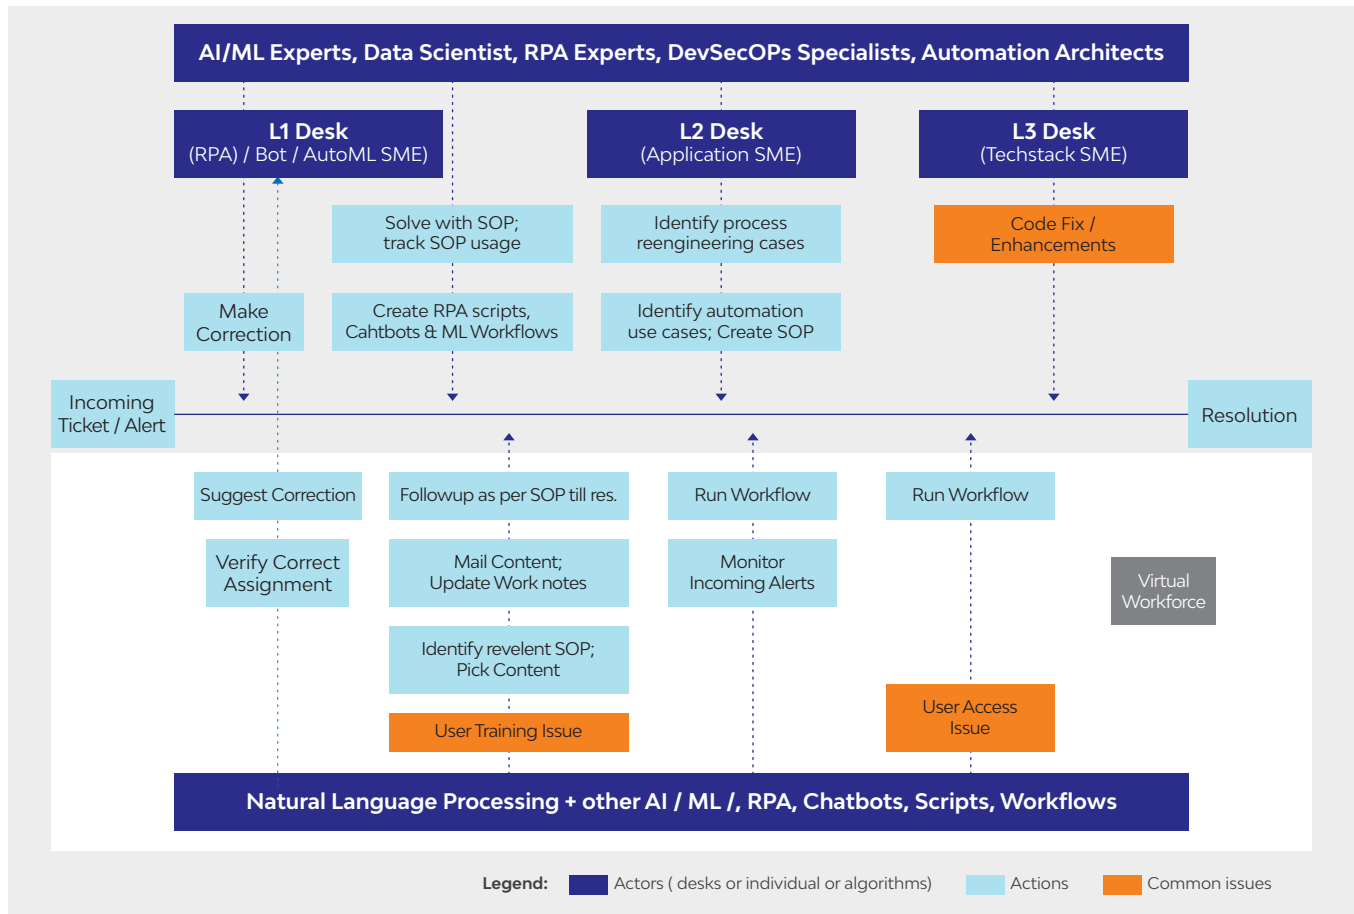
One of the best ways to measure issues is with the help of knowledge articles. I believe all incidents closed should reference a knowledge article unless it's a code fix or enhancement (that means we'll have to create those knowledge articles in the first place diligently). And this knowledge article should be referenced every time a similar incident is closed. Tallying the usage numbers give a pretty good idea of the trend. As we shall see ahead, this humble knowledge article forms the foundation of chatbot and AI-powered auto-resolutions for common issues.

2. Why automate when you can eliminate

Chatbots and RPAs are mainstream these days. It is common knowledge that repetitive workflow should be automated. The question we often miss is, "Can this volume of incidents be eliminated?" We saw north of 100 tickets a month for a particular application for a whole quarter in a previous support engagement—mostly around user training and access issues. For some time, we contemplated using chatbots to cater to training issues while diverting access issues to the workforce. However, the final solution—using an SSO in this case—changed the workflow and eliminated the volume.

In the day-to-day rush to close incidents, we fail to take a good look at redundant processes and workflow (complex approvals processes are a good example) that may have made sense in the past but not so much now. Still, working on a sizeable volume of such tickets is a drain on resources that can easily be minimized and even eliminated. There is no easy way to say it—reducing count may reduce workforce and, therefore, billing. But, if you don't do it, someone else will.

An AI-powered and data-centered operating model.



Application support is a cost center. The budget is fixed, and there is always pressure to reduce costs year on year (or in a downturn). In contrast, revenue-centered applications can ever hope for more funds as there is a prospect for clear and measurable cashflow and therefore, ROI.

While support teams, through incremental efficiency and productivity gains, deliver cost reductions, they can only do so much,

especially after years of optimization. Adopting advances in technology can provide a way out.

Fig II shows an AI-powered and data-centric operating model for application support. The lower section depicts a virtual workforce (composed of AI/ ML algorithms and automation) that supports the human workforce. Some common tasks (there are others as well) are allocated to each; the

purpose of the human workforce is two fold: to develop and maintain the virtual workforce and carry out tasks that a virtual workforce can't handle (e.g., code fix). It follows therefore that the skill sets of a support desk will have to change.

For example, the L1 or L1.5 desk would have to understand trends and implement RPA, chatbots, and ML workflows, but they won't have to reinvent the wheel. They will be supplemented by a team of AI/ ML experts, data scientists, RPA experts, DevSecOps specialists, and automation architects who will generate the base artifacts (or templates, if you will) that can be tweaked across different engagements. Most importantly, the new team will need to have an automation mindset to identify opportunities and be willing to implement automation wherever possible.

The artifacts will have to be modular, and most likely, we will have to pick and stick with platforms such as Mosaic or Azure (for AI/ ML), ServiceNow, or HPALM (for ITSM tools). Replicating the artifacts for different platforms will hamper scale and speed to market.

This APDC operating model offers scale benefits, i.e., the marginal cost to scale for another client is negligible (though not zero) compared to what the traditional model offers. However, the marginal cost will depend on the percentage of automation achieved (a critical KPI could be: 1% of workforce required relative to a traditional model). The higher the usage, the better the quality of models will be, and the lower will be the marginal cost. And the team of experts can be spread out across multiple engagements like a Centre of Excellence is. However, the organization must bear the initial start-up cost that will likely exceed that of a traditional model, particularly in the first year. Subsequent years will offer more automation benefits as both the model and the human workforce mature.

Before we delve into the technical specifics, it is essential to highlight that endeavors such as O2T can fall under the ambit of the new model. O2T offers benefits in pockets, while the new model recommends a fundamental change in the operating model driven by a clear vision.

The nuts and bolts

This section will outline some use cases and methods to transform those using the APDC operating model.



Ticket routing at L1 desk

Ticket routing is bread and butter for every IT support, yet rarely done right. Tickets are frequently misrouted, leading to SLA breaches. Ticket routing is essentially a natural language processing and classification problem. Most importantly, support desks will have well-labeled data making this use case a low-hanging fruit that should be plucked right away. (reference examples can be found here:

<https://www.analyticsinsight.net/it-ticket-classification/> and

<https://towardsdatascience.com/it-support-ticket-classification-and-deployment-using-machine-learning-and-aws-lambda-8ef8b82643b6>)

Additional attributes such as urgency and sentiment can be deciphered from the text. The underlying issues behind the occurrence of similar incidents can also be detected. Similarly, the models can read directly from emails and social media mentions (such as Workplace).



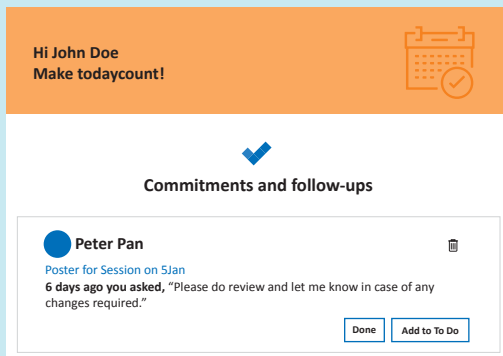
Suggest routing corrections

Some tickets will still be acted on by a human workforce, but AI can trigger suggestions when the models point otherwise.



Resolution at L1.5 desk

The humble KA that we proposed earlier makes its mark here. L1.5 desk is utilized for standard, high volume issues (say, user access and password reset) that can be resolved with an SOP and without application or technical knowledge. Ideally, the L2 or L3 team would identify such issues and produce KAs. The L1.5 desk would leverage these KAs to solve problems or pass them on to L2 or L3 desk. AI can help match incoming issues to KAs and reproduce the relevant text in a mail sent to the end-user. This is quite similar to the message we receive every day from Cortana (reproduced here).

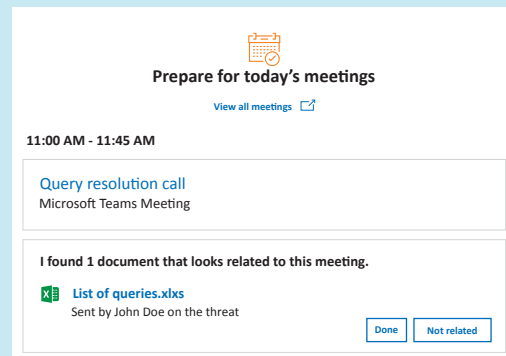


Hi John Doe
Make todaycount!

Commitments and follow-ups

Peter Pan
Poster for Session on 5Jan
6 days ago you asked, "Please do review and let me know in case of any changes required."

Done Add to To Do



Prepare for today's meetings
View all meetings

11:00 AM - 11:45 AM

Query resolution call
Microsoft Teams Meeting

I found 1 document that looks related to this meeting.

List of queries.xlsx
Sent by John Doe on the threat

Done Not related



Monitoring

AI can help monitor the trend of incoming incidents that may point to larger (high priority) issues and run preset workflows. Such significant problems, especially those coming outside of business hours, can be escalated to the right people. Here we can leverage several automation tools (such as Ulpath) to run a preset workflow. However, the trigger would be an AI rather than a human workforce.

Similarly, AI can monitor job failures and alert Point of Contacts (or run other preset workflows)



Chatbots

This tool can be leveraged for standard, scripted Q&As and with recent advances, non-scripted as well.



Workflows

Here, automation will be used extensively. The O2T initiative has generated a good volume of use cases that can be leveraged across BU and LTI aided by AI triggers. Other use cases, for example, could be closing tickets based on business rules (say, close after two follow-ups)



Identify use cases

Mosaic Discovery analysis is an excellent tool for helping agents identify trends using textual data. A similar analysis can be generated on demand.



Fulfill standard requests

Standard requests such as that for software, devices, and access can be automated to interact with backend workflows to secure permissions and initiate processes automatically. These catalogs can be built with tools such as ServiceNow.

The use cases mentioned above are not all there is. As the model matures, use cases can be generated incrementally. The idea is business enablement—technology needs to enable business, which includes optimizing application support (cost optimization rather than fanatic automation). A systems view coupled with cost-benefit analysis at every step can point out which items to automate and which to leave out.

As stated before, the APDC operating model will require a new team model—one composed of AI/ ML experts, data scientists, RPA experts, DevSecOPs specialists, and automation architects in addition to the application and tech stack SMEs. At the same time, it will have to bring disparate technologies- Chatbots, RPA, scripts, macros, batch programs, NLP, speech recognition, computer vision -under one roof to create value for clients

The primary reasons for adoptions are improved CX from better turnaround time (faster closure of low priority and straightforward issues) and service quality, reduced cost, and better employee fulfillment from engaging in creative work, not to mention the cost reduction from employing a lean, AI- enabled workforce.

It's high time, organizations moved away from pockets of automation and AI enablement for discrete clients and projects and develop a domain- agnostic solution leveraging a multitude of technologies- a productized application support solution. The imperative, therefore, is to be agile. Organizations that move fast will gain the first-movers advantage and capture or drive the market; everybody else will need to adapt fast or go home.

About the Author



Sinjan M Kashyap

Specialist- Software Engineering | Manufacturing

Sinjan has 10+ years of delivery and consulting experience in the Manufacturing and Technology domain and is currently engaged in delivering data driven business outcomes for global brands across sales, services, and the supply chain. He holds an MBA from IIM Indore.

LTI (NSE: LTI) is a global technology consulting and digital solutions company helping more than 400 clients succeed in a converging world. With operations in 31 countries, we go the extra mile for our clients and accelerate their digital transformation with LTI's Mosaic platform enabling their mobile, social, analytics, IoT and cloud journeys. Founded in 1997 as a subsidiary of Larsen & Toubro Limited, our unique heritage gives us unrivalled real-world expertise to solve the most complex challenges of enterprises across all industries. Each day, our team of more than 33,000 LTItes enable our clients to improve the effectiveness of their business and technology operations and deliver value to their customers, employees and shareholders. Find more at <http://www.Ltinfotech.com> or follow us at @LTI_Global.