AI & ML:
Driving the Next Generation of Innovation in DevOps and Workload Automation

By Stephen D. Hendrick
An ENTERPRISE MANAGEMENT ASSOCIATES® (EMA™) Research Report
November 2020

Sponsored by:

bmc

PagerDuty

EMA™ IT & DATA MANAGEMENT RESEARCH, INDUSTRY ANALYSIS & CONSULTING

AI & ML:
Driving the Next Generation of Innovation in DevOps and Workload Automation

Table of Contents

EMA

AI & ML:
Driving the Next Generation of Innovation in DevOps and Workload Automation

Executive Summary

The reality of artificial intelligence and machine learning (AI/ML) is being made possible by the evolution of technologies associated with big data, analytics, advanced analytics, operations research, data science, cloud computing, microservices, and decision science. AI and ML are just the most recent chapters in the amalgamation of computer science, mathematics, data management, and decision science.

This research demonstrates that most large enterprises see AI/ML as important priorities and strategies are already in place to leverage AL/ML capabilities. AI/ML expectations at large enterprises are lofty, which could lead to inflated expectations followed by a trough of disillusionment. EMA does not see this as likely because much of the heavy lifting in AI/ML is now in the past. EMA sees AI/ML unfolding in DevOps and workload automation first as machine learning avenues to provide feedback on decisions made and actions taken. This feedback loop provides an opportunity for the continuous improvement of which decisions and actions are taken, thereby driving a wide array of improvements in quality.

AI will be seen as a way to bring automation to the feedback loop supported by ML. This automation will begin as a collection of objectives that will provide a context for the actions being taken by the enterprise. An AI engine will initially guide humans by recommending decisions and actions to be taken and transition over time to higher and higher levels of automation based on experience, quantified benefits, and risk management objectives.

Leading products in workload automation today provide policy-based automation where automation engineers can set rules based on analytics from the workload automation tools own meta data. For example, a policy for times of peak processing where compute utilization exceeds a threshold would have the workload automation product interface with an infrastructure automation solution to provision additional compute and process workload there and then terminate the provisioned resources when processing is complete. The next evolution of this would be to have AI/ML algorithms with workload automation solutions that can self-learn such patterns and then be able to recommend remedial actions taking the system closer to a self-healing concept.

Many DevOps and workload automation tools are already sporting ML-based capabilities, so within the next five years, there could be two waves of innovation: the first based on ML and the second on AI. The research from this study on AI/ML in large enterprises focuses primarily on the impact of DevOps and workload automation. The key findings from this research are as follows:

- **AI/ML benefits will spring from improved quality and automation.** ML will drive improvements in quality, which will drive wide-ranging benefits on application performance, reliability, OpEx efficiencies, cost savings, and customer satisfaction. AI will follow as a second wave of innovation built on ML that will enable expansive automation. This automation can mean improved speed in delivering new applications and updating applications, and improved employee productivity—all of which reduce SDLC cycle time and leave more time for addressing higher-order development work.

- **Enterprises must build a foundational understanding of AI/ML.** Vendors and large enterprises are aware that AI/ML will drive innovation, but it is challenging for enterprises to wrap their heads around what AI/ML is, how it delivers value, and how they should leverage this value. Senior management must therefore understand the principles behind AI and ML so they know what AI/ML can deliver and how it differs yet complements advanced analytics (see a discussion of this in the Conclusions section).
- **Companies should develop a Center of Excellence for AI/ML.** AI/ML is driving the next phase of innovation in DevOps. Develop a CoE for AI/ML that includes subject matter experts and senior IT and business leaders. Based on this survey data, enterprises already see AI/ML as priorities and have a strategy, but challenges introduce doubt in how to operationalize AI/ML. The wide range of benefits from AI/ML identified in this survey suggests that AI/ML is expected to have a pervasive impact on DevOps and Workload Automation.
- **Users will determine where to use AI/ML to the greatest effect.** Based on your DevOps maturity, determine how best to phase in DevOps tools and where to leverage AI/ML to maximum effect. Where do you really need AI/ML, and do you have the staff in place to achieve your objectives?
- **Customers need to demand more of their software vendors.** Software vendors could be doing far more than they are today in explaining their AI/ML capabilities and how they add value. Software vendors need to build stronger relationships with their customers involving training, education, and consulting to ensure that customers derive the greatest benefit from their products and are successful using them. AI/ML capabilities from vendors will evolve and change quickly. Task vendors with keeping you up to date on their AI/ML strategy, roadmap, and support in optimizing use of their tools. View this as a joint opportunity between your enterprise and its software vendors and as a way for both of you to be more successful.
- **Workload automation Tools are embracing AI/ML.** The mission critical aspects of workload automation are beginning to leverage AI and ML. A combination of predictive analytics in support of improved decision making, pattern identification through machine learning, and optimization to scale resources and align with workload needs are actively being pursued.
- **Don't delay.** Large enterprises have already prioritized their development AI/ML strategies. Their larger IT resources also mean they can assimilate AI/ML technology at a faster rate. However, AI/ML give smaller/medium enterprises the ability to compete more effectively with large enterprises. Therefore, don't make AI/ML a lost opportunity because of delays in understanding what it can do for you or delays in implementing the technology.

Methodology

The DevOps AI empirical research project was fielded in August 2020. The primary objective of this research was to understand the impact that AI/ML is having on the enterprise, DevOps, and workload automation. Specially, this report will address the following:

- The impact that AI/ML is having on DevOps, workload automation, and the enterprise
- What is the state of workload automation in the enterprise today?
- Identify best practices in workload automation

A secondary objective of this research will be to see how DevOps maturity influences AI/ML. The intent is to understand whether there is correlation between enterprise size, DevOps maturity, IT budget, and commitment to AI/ML.

Survey Parameters, Screening Criteria, and Data Analysis

EMA's DevOps AI sample was recruited through a third-party enterprise panel using an email invite to an online survey. The sample size was 289, with 45% coming from North America, 28% from Europe, and 27% from Asia Pacific. The survey went into the field in August 2019 and targeted IT professionals with deep experience in AI/ML, workload automation, and event and incident management. EMA targeted larger enterprises with 1,000 or more employees, but the average enterprise size was 14,000 employees and 25% of the enterprises had 20,000 employees or more. A wide cross-section of vertical industries with material involvement in application development and DevOps was included. EMA did not include enterprises involved in education or information technology in the sample because these vertical industries consume DevOps tools and technology in unique ways that are not always consistent with end-user enterprises. EMA's survey was directed at staff engaged with pipeline activities, including developers, DevOps managers, IT operations management, and IT executive management.

Screening criteria included enterprise size, involvement in application development, role, area of responsibility, vertical industry, and deep subject matter experience.

The standard method of displaying data for each chart in this report is as a single response question and the percentage responding from the entire sample of 289 responses. EMA will indicate in the chart or text describing the chart when the data displayed is from a multiple-response question, represents valid cases from the sample, or shows some other non-standard approach involved in preparing the data.

Based on the sample size, the margin of error for the survey at 90% confidence is +/- 4.8%.

The State of AI/ML in Large Enterprises

Large Enterprises Show DevOps Maturity

One of the first questions that EMA asked was for the respondent to assess the DevOps maturity of their enterprise. DevOps has been around for more than a decade and shows software revenues over $8 billion. It stands to reason that large enterprises, with their correspondingly bigger IT budgets, would spend more on DevOps and therefore have a higher level of DevOps maturity. In fact, that's exactly what happened. Figure 1 shows that 32% of large enterprises have automated nearly all of their DevOps activities, 48% have automated many DevOps activities, and 16% of enterprises have automated some of their DevOps activities. Taken together, this means that 96% of large enterprises are deeply engaged in their DevOps journey. This contrasts with the 4% of the sample was just beginning their DevOps journey.

Figure 1: DevOps Maturity in Large Enterprises

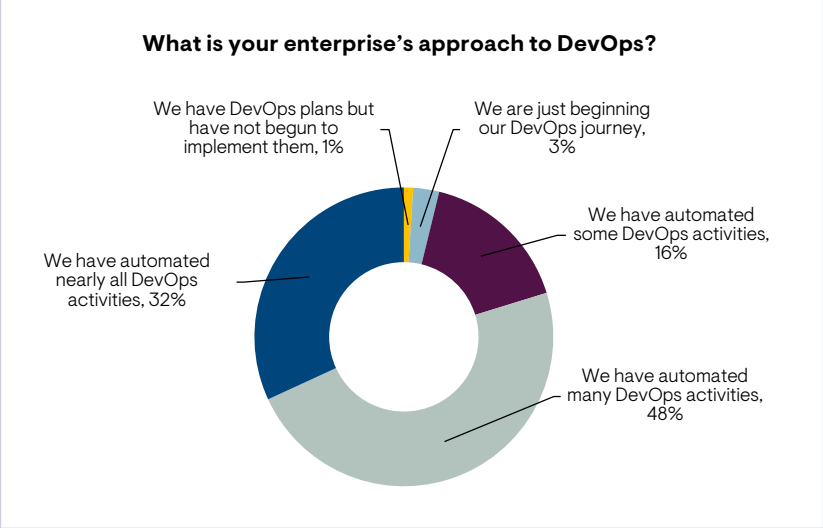

The important point regarding DevOps maturity is that larger enterprises have larger IT groups with bigger budgets and have been investing in DevOps for a longer period of time. This is likely to give large, mature DevOps shops an advantage in productivity, speed, and efficiency. For the 4% to 20% of large enterprises in the survey who are early in their DevOps journey, they need to accelerate the expansion of their toolchain for risk and increasing gap in their IT competitive advantage.

EMA elected to show DevOps maturity as the initial figure in this report because data will be used as a segmentation variable to drill deeper into other questions. Using DevOps maturity to provide a more granular view into other questions shows splits by enterprises that have automated some, many, and nearly all DevOps activities, since this accounts for 96% of the sample.

IT Budgets Correlate with DevOps Maturity

Because IT is important to every enterprise, larger enterprises will typically have larger IT budgets. DevOps has been used for over 10 years, giving larger enterprises plenty of opportunities to experiment, adopt, and mature their DevOps activities.

Figure 2 shows the 2020 IT budget for enterprises in the sample and segments these enterprises by level of DevOps maturity. Automation of DevOps activities (some, many, or nearly all) is the proxy for DevOps maturity. Each line in Figure 2 shows 100% of the enterprises in the segment distributed by IT budget category. Each curve has bimodal characteristics, and the peaks of these bimodal curves shift to the right (meaning larger IT budgets) as DevOps maturity increases.

Figure 2: 2020 IT Budgets by DevOps Maturity

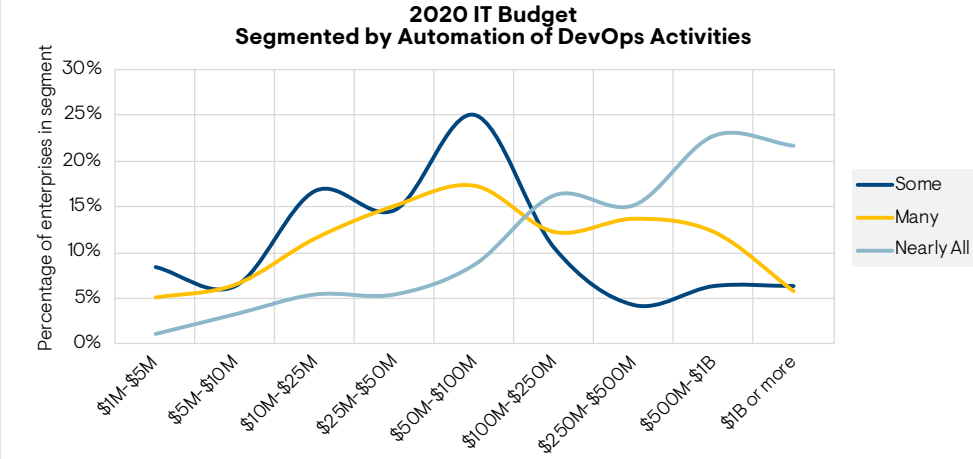

Enterprises that have automated some DevOps activities show 17% of its segment with $10M-$25M budgets and 25% of its segment with $50M-$100M budgets. Enterprises that have automated many DevOps activities show 17% of its segment with $50M-$100M budgets and 14% of its segment with $250M-$500M budgets. The significant difference comes with enterprises that have automated nearly all DevOps activities, showing 16% of its segment with $100M-$250M budgets and 23% of its segment with $500M-$1B budgets.

This data simply confirms EMA's suspicion that the larger the enterprise, the larger the IT budget and the more mature their DevOps practice is likely to be. This is important because more mature DevOps practices may be early adopters of AI/ML capabilities, thereby potentially accelerating their ability to improve how IT supports the needs of the business.

Large Enterprises See AI/ML as High Priorities

Figure 3 summarizes how large enterprises view AI/ML. To begin with, 87% of large enterprises say that AI/ML are high or extremely high priorities. At the same time 86% of large enterprises have an AI/ML strategy that is mostly or extremely well defined. The reason for this is because AI/ML can systematically improve the real-time quality of decision-making and actions taken, then bring automation to these decisions and actions. The impact of this can be profound and can include faster time to market, improved quality and reliability, reduced risk, reduced cost, increased customer satisfaction, and better application performance.

Figure 3: Enterprise Perspectives on AI/ML

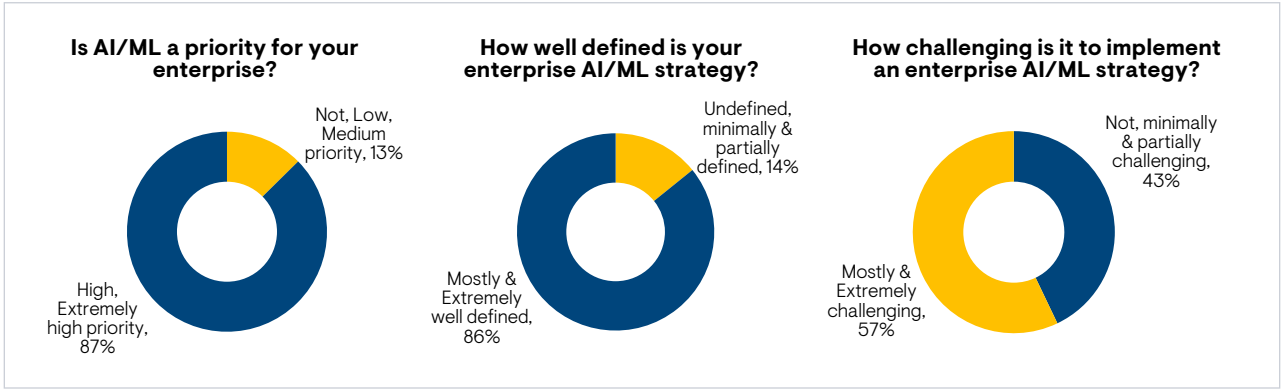

However, 57% of large enterprises also say that implementing an AI/ML strategy can be mostly or extremely challenging. Herein lies the challenge of AI/ML. While large enterprises are on one hand very bullish about the capabilities of AI/ML, they also understand that leveraging AI/ML capabilities is not easy.

Who Sees AI/ML as More of a Priority?

Figure 4 looks at the priority that enterprises attach to AI/ML, segmented by IT budget and DevOps maturity.

The demographics of IT budget size, which are shown in Figure 20, reveal 33% of the sample have budgets less than $50 million, 29% have budgets between $50 million and $250 million, and 38% have IT budgets between $250 million and $2 billion. These budget categories are fairly evenly distributed from the standpoint of number of enterprises in each category, but the budget ranges are exponential, with budget midpoints at $25 million, $150 million, and $875 million.

Figure 4: How the Importance of AI/ML Changes by IT Budget and DevOps Maturity

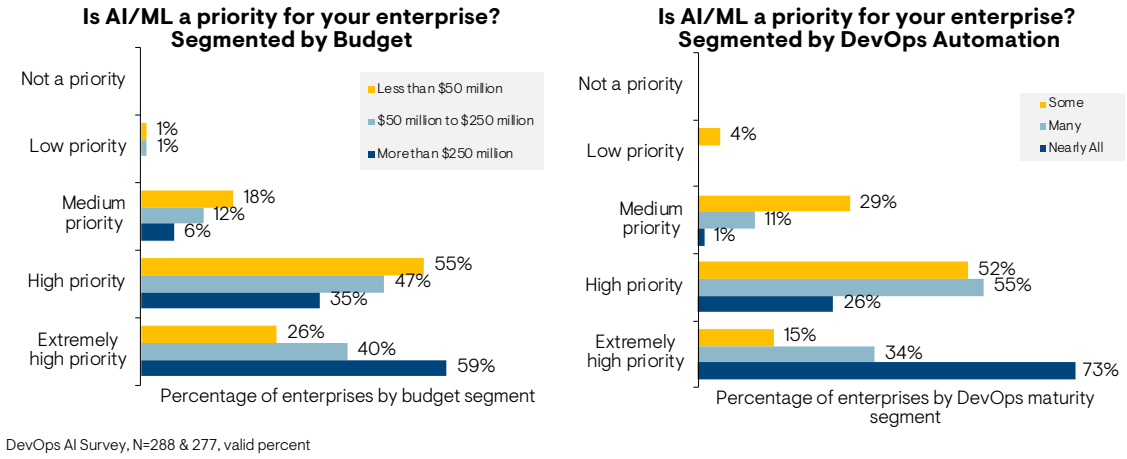

DevOps AI Survey, N=288 & 277, valid percent

Both graphics show a strong prioritization of AI/ML by nearly all enterprises. However, 55% of enterprises with IT budgets below $50 million and 47% of enterprises with budgets between $50 million and $250 million predominately see AI/ML as a high priority. For these two segments, high priority is the peak of their respective distributions. For enterprises with budgets above $250 million, 59% see AI/ML as an extremely high priority. This contrasts with just 40% of enterprises with budgets between $50 million and $250 million, and 26% of enterprises with budgets less than $50 million that see AI/ML as an extremely high priority. Because of the exponential difference in IT budget size classes, enterprises at the top of the IT budget range have significantly more IT economic power than the firms below them which, combined with the great priority they attached to AI/ML, means they are likely to increase their competitive difference over enterprises in the other two segments. Down-segment enterprises can mitigate this risk by accelerating their AI/ML journey and by being more strategic in their AI/ML investments.

The second graphic in Figure 4 shows the priority of AI/ML segmented by DevOps maturity. The same characteristics described in the prior IT budget analysis carry over to DevOps maturity, but are amplified even more. For enterprises that have automated nearly all of their DevOps activities, 73% see AI/ML as an extremely high priority. This contrasts with just 34% of enterprises that have automated many activities and 15% of enterprises that have automated some DevOps activities. Because DevOps represents the degree of sophistication of an enterprise's IT business, it is hard to improve on activities, processes, and automations that don't exist. However, enterprises that are still very much in the midst of their DevOps journey can gain competitive ground by filling DevOps gaps with tools providing advanced AI/ML capabilities and leapfrog an entire generation of tool functionality.

Expected DevOps Benefits From AI/ML by DevOps Maturity

Figure 5 shows the expected DevOps benefits from using AI/ML. These benefits are shown in descending order of enterprise utility. The responses highlighted with a gray background are those where the overall response rate is 35% or greater. This data is further segmented by the level of DevOps activity automation within the enterprise.

Figure 5: Expected DevOps Benefits From AI/ML by DevOps Maturity

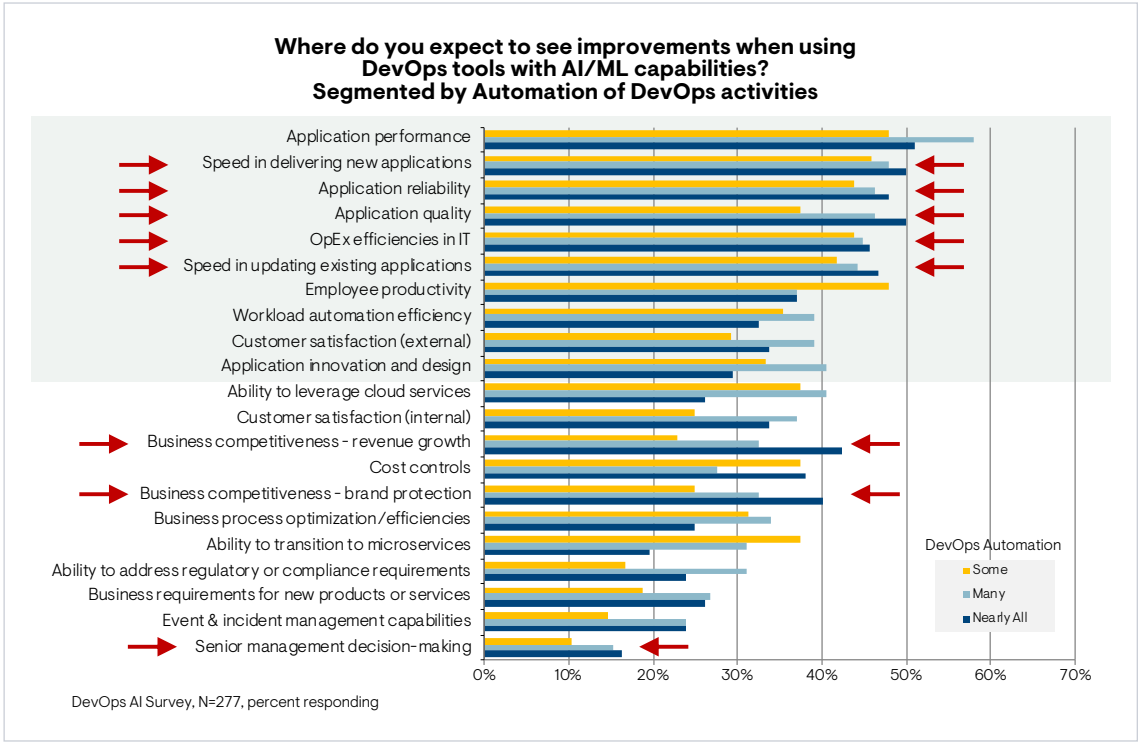

The top 10 benefits thematically fall into the two AI/ML categories: quality and automation. Quality will largely stem from machine learning combined with advanced analytics. Machine learning uses feedback from human response data, experiences, and actions to improve future machine-based recommendations and decision-making. The result is improved application quality, which has a downstream positive impact on application performance, reliability, OpEx efficiencies, and customer satisfaction.

In some cases, machine learning will be embedded in products and not user-configurable. This approach is useful for optimizing the performance of well-bounded analytic activities, but the higher-value use case for machine learning is as a recommendation engine for decision-making that initially runs in supervised mode. With experience, it can transition over time to running in an unsupervised capacity. The transition from unsupervised to supervised leads to automation. This automation means improved speed in delivering new applications and updating applications, and improved employee productivity—all of which reduce SDLC cycle time and leave more time for addressing higher-order development work.

The left-facing arrows on the slide call out responses where the utility of the response increases with DevOps maturity. There are two cases where these findings are significant based on the survey's margin of error. Both have to do with business competitiveness. The first involves revenue growth and the second involves brand protection. Figure 5 shows that revenue growth was seen as a benefit by 42% of enterprises with high levels of DevOps automation, compared to 33% with medium levels of automation and 23% with low levels of automation. This suggests that as a DevOps practice matures, it becomes more closely align with how to best support the needs of the business because of increased speed, agility, and productivity. These same characteristics hold true for brand protection and confirm how mature DevOps practices are better integrated into the business fabric of the enterprise.

Expected DevOps Benefits From AI/ML by IT Role

Figure 6 is another view into expected DevOps benefits from using AI/ML, except that this view is segmented by IT role instead of DevOps maturity. The rank order and overall utility remain the same as the prior slide, but there are some unmistakable areas that were important to IT operations. Areas where IT Ops has significantly higher AI/ML expectations in DevOps include speed and updating existing applications, application innovation and design, revenue growth, and business process optimization.

Figure 6: Expected DevOps Benefits From AI/ML by IT Role

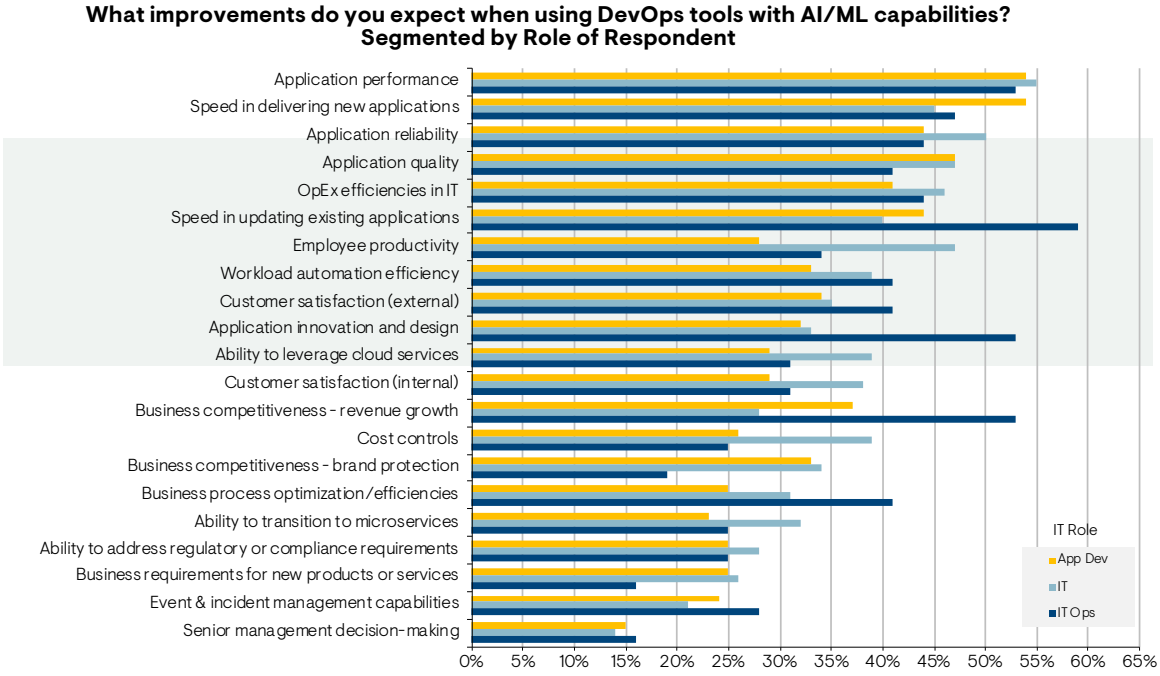

What is interesting is that these areas show a keen interest by IT Ops on internal operations, as well as key activities in application development and business competitiveness. This is happening because IT operation activities are inherently data-driven, relying on events and telemetry. This also shows how IT Ops is now more integrated into the fabric of a company's digital persona. IT Ops is participating more at the frontend of the lifecycle (application innovation and design) to ensure the supportability of new production apps. IT Ops also cares about revenue, indicating a transition in focus from IT project to product. The growing complexity of IT Ops and ITSM means a more formal focus on business processes, governance, and the goal of optimization across the lifecycle.

The Challenges in Using DevOps Tools with AI/ML Capabilities

Fifty-seven percent of large enterprises find implementing an enterprise AI/ML strategy either mostly or extremely challenging. Figure 7 details what those challenges are. The responses highlighted with a gray background are those in which the overall response rate is 35% or greater. This data is further segmented by the level of DevOps activity automation within the enterprise.

Figure 7: AI/ML DevOps Challenges by DevOps Maturity

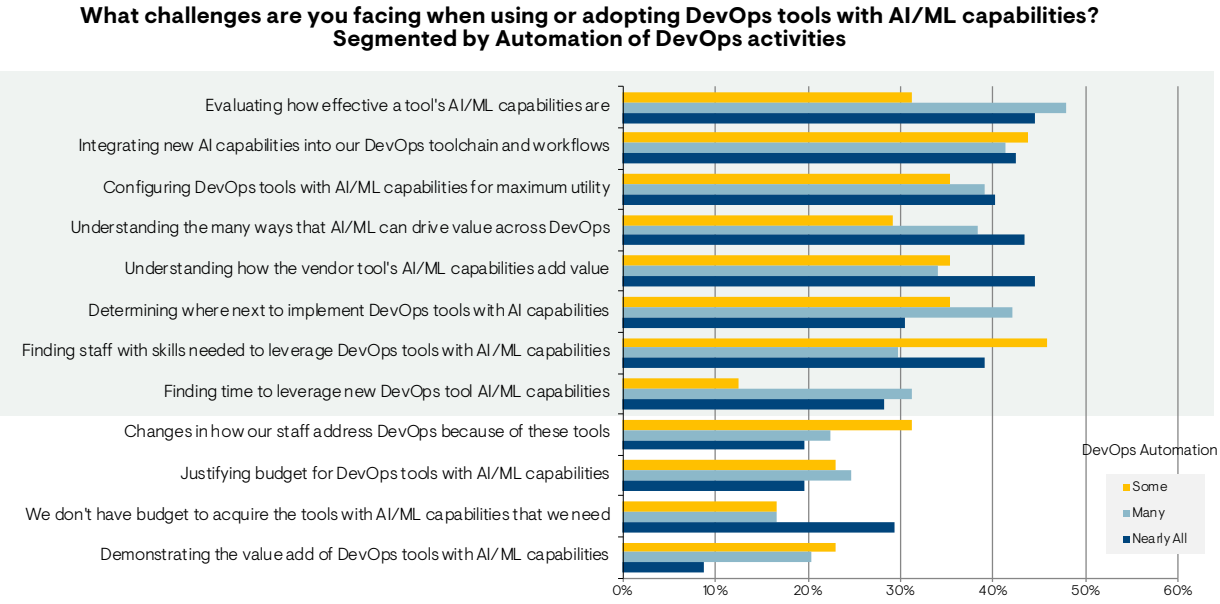

To summarize the challenges in Figure 7, enterprises are having difficulty understanding:

- How AI/ML adds value to DevOps tools and DevOps processes
- How to evaluate AI/ML capabilities
- How to integrate and configure AI/ML capabilities
- Where to implement these capabilities
- How to find staff with the right AI/ML skills

It's also true that enterprises with mature DevOps practices were the leaders in stating these challenges. These challenges exist because of two reasons:

1. Enterprises don't understand what AI/ML are, so they are not equipped to fully leverage AI/ML capabilities unless the capabilities are deeply embedded and don't require use involvement.
2. As enterprises get more involved with AI/ML, they begin to know what they don't know and adopt a more coherent approach to building out AI/ML expertise.

What this means is that implementing AI/ML capabilities is difficult. It's difficult because AI/ML introduces continuous improvement and therefore continuous change. The problem is that change is a dirty word in IT. In most cases, IT problems are because something has changed. At the same time, there is also confusion between what AI/ML are and what advanced analytics are. Advanced analytics, such as correlation, predictive analytics, and optimization, are not AI/ML. However, they are often combined with AI/ML to build AI/ML capabilities. Therefore, developing a firm understanding of what AI/ML are needs to be the first step in an enterprise's AI/ML strategy.

The Impact of AI/ML on Workload Automation/Management

Is Workload Automation a Part of Application Development and DevOps?

In Figure 8, workload automation's role in application development is shown on the left and its relationship to DevOps is shown on the right. Looking at application development first, 73% of the sample says they have a formal process for addressing workload automation during application development. This contrasts with the 9% of large enterprises that address workload automation on more of an ad hoc basis. The 18% in the middle that informally address workload automation depending on what's being developed poses a conundrum. Do these enterprises always address workload automation during application development when it's needed, or do they simply not have a rigorous approach to addressing workload automation? The safe approach is to always have a formal process for addressing workload automation during application development. If this enables companies to claim that the informal responses can be grouped with the formal responses, then 90% of the sample is including workload automation as a part of their application development process.

Figure 8: Workload Automation's Relationship to Application Development and DevOps

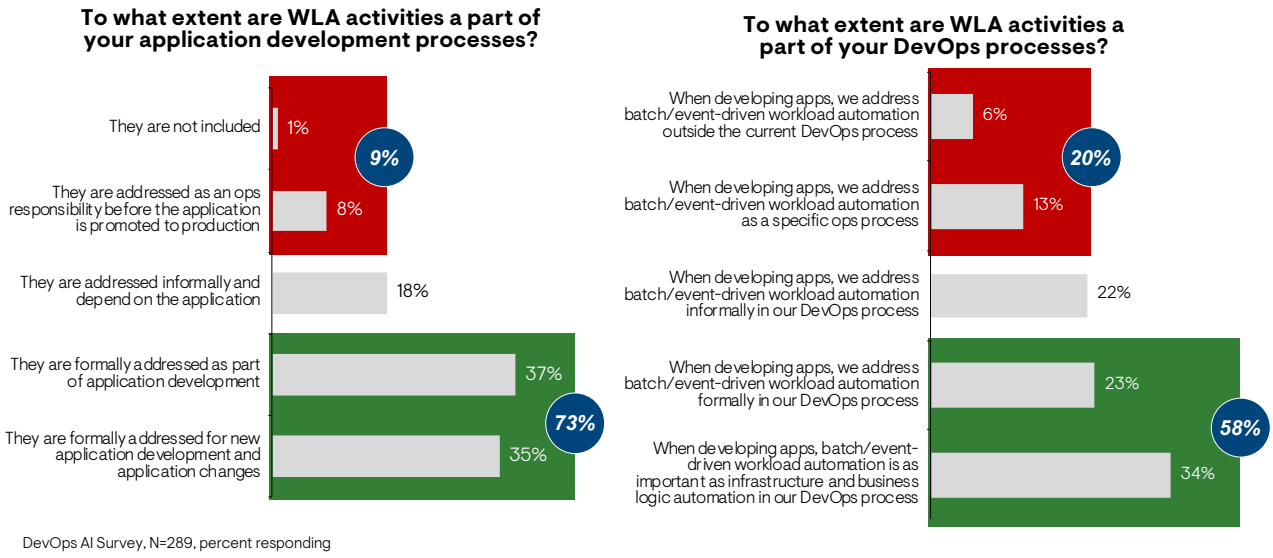

On the right side of the slide, there is a similar pattern. Fifty-eight percent of the sample have a formal DevOps process for workload automation, whereas 20% approach workload automation on an ad hoc basis or as part of the Ops workflow. This leaves 22% who address workload automation as an informal part of DevOps. EMA is convinced that most enterprises are on a DevOps journey that increases their DevOps maturity from year to year. EMA believes that enterprises with an informal workload automation approach align better with those that have a formal approach, meaning 80% of enterprises have a coherent approach to workload automation.

EMA looked at how DevOps capabilities have been adopted over time. EMA's longitudinal DevOps adoption data shows that the early emphasis in DevOps was on application development and establishing an effective approach to CI/CD. In the last several years, the emphasis in DevOps has shifted to operations with the rise of AIOps and increases in observability, enabling operations to become more data-driven and keep pace with the DevOps improvements already in place for application development.

Workload automation is a mission-critical activity that is now formally a part of application development and DevOps at the majority of large enterprises and is in the process of becoming a pervasive element of nearly all large enterprises.

How Enterprise Workload Automation Strategies Vary by DevOps Maturity

Figure 9 shows three curves in a chart corresponding to EMA's DevOps maturity question asking how many activities are automated in their DevOps pipeline (some, many, or nearly all). Each curve is then cross-tabbed against the five stages of workload automation maturity. Enterprises that have automated some DevOps activities are the blue curve, which peaks at 35% but then declines to 12%. Enterprises that have automated many DevOps activities are the green curve, which shows greater maturity in how workload automation is addressed in DevOps, with 25% in the final stage. Then, enterprises that have automated nearly all DevOps activities (the orange curve) show as 62% in the final stage of how workload automation is addressed in DevOps.

Figure 9: Workload Automation Strategy Segmented by DevOps Maturity

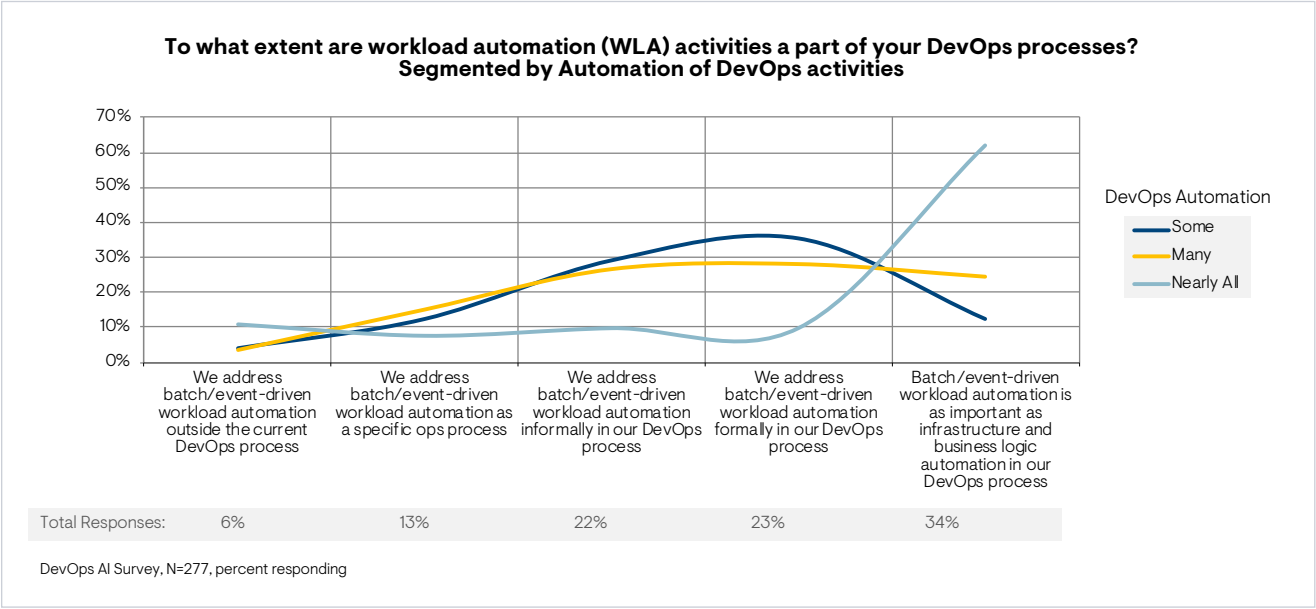

There is a strong correlation between enterprise DevOps maturity and workload automation DevOps maturity. The more mature a company is in DevOps, the more mature the approach to having workload automation be a part of application development and DevOps. Because workload automation is a mission-critical activity, enterprises of any size can avoid risk, lost productivity, and time to market delays by bringing workload automation into their DevOps pipeline for application development.

AI/ML Top 3 Business Benefits for Workload Management

Figure 10 provides a list of workload management business benefits. Workload management encompasses IT operations and workload automation. EMA first asked large enterprises which workload management business benefits AI/ML would drive. For those responses selected, EMA next asked enterprises to rank order their top three business benefits. Figure 10 shows the percentage of first, second, and third selections based on total mentions for the response.

Figure 10: Top 3 AI/ML Business Benefits for Workload Management

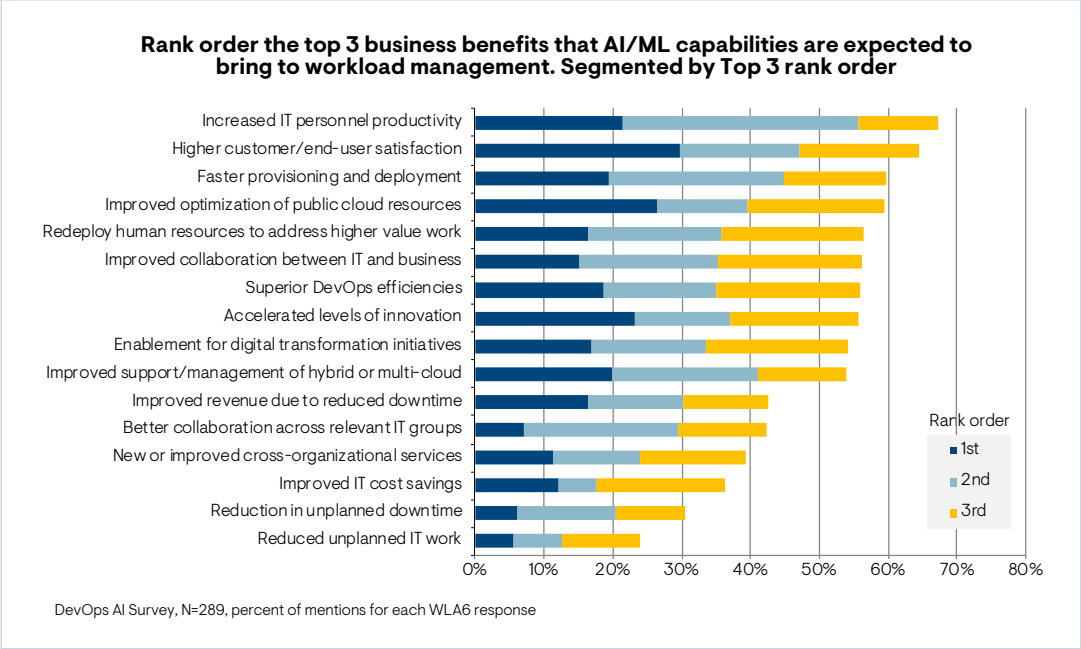

DevOps AI Survey, N=289, percent of mentions for each WLA6 response

Increased IT productivity and higher customer/user satisfaction are the two overall leading business benefits. These two responses also have a distinctly broader scope than the other responses, which aligns with the strategic importance that enterprises attach to AI/ML for workload management.

Increased IT productivity, the #1 overall business benefit, will likely key off ML to initially improve workload management and AI to achieve WLA objectives through an evaluation of risk and economics to automate and optimize operational activities.

Higher customer and user satisfaction and improved optimization of public cloud resources are the leading #1 and #2 responses. Both of these are likely to leverage ML as a way to improve and automate problem remediation, problem avoidance, and resource cost minimization or optimization.

EMA also looked at the importance that enterprises attach to the improved support and management of hybrid or multi-cloud environments based on the percentage of #1 and #2 choices. It becomes the fifth-ranked choice, rather than tenth. Support for hybrid and multi-cloud environments is an important topic today, and AI/ML could provide a way to help optimize workload management across the many environmental choices that exist.

AI/ML Top 3 Functional Benefits for Workload Management

Figure 11 shows that ML in support of auto-remediating job processing issues is the leading choice and overall functional benefit. This is because ML can provide a far more granular approach to evaluating job processing issues, determining potential paths forward, implementing the appropriate solution based on objectives, and continuously learning/improving from actions taken.

Figure 11: Top 3 AI/ML Functional Benefits for Workload Management

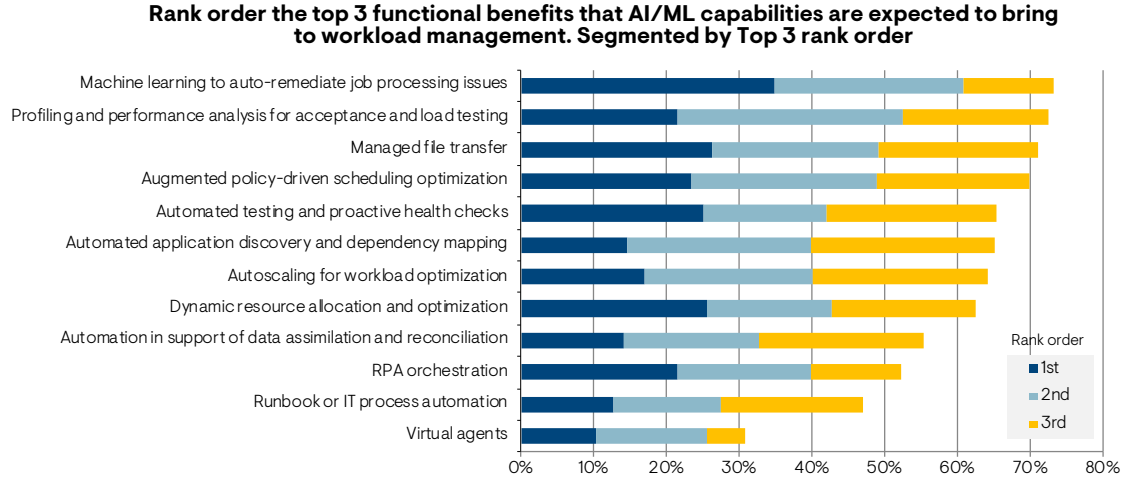

DevOps AI Survey, N=289, percent of mentions for each WLA8 response

Performance management pre/post workload deployment is the #2 overall expected benefit from AI/ML. While today's observability tools address pre/post workload performance, ML can be leveraged to optimize workload performance while accounting for changing resource availability and enterprise priorities.

The next three benefits were managed file transfer (MFT), policy-driven scheduling, and proactive health checks. They can likewise be addressed through a combination of today's existing tools and ML to bring a more comprehensive and coherent framework for evaluating availability and performance, understanding key objectives and SLAs, recommending, and initiating actions, and driving continuous improvement.

AI/ML Best Practices for DevOps and Workload Management

The priority that large enterprises give to AI/ML, the investments they have already made in AI/ML strategy, and the many AI/ML-driven businesses are collectively driving the next generation of innovation in DevOps and workload automation. Based on the data in this report, the best practices for large enterprises are as follows:

- **Enterprises must build a foundational understanding of AI/ML.** Vendors and large enterprises are aware that AI/ML will drive innovation, but it will be challenging for enterprises to wrap their heads around it. So, understand the principles behind AI and ML so you know what AI/ML can deliver and how it differs yet complements advanced analytics.
- **Companies should develop a Center of Excellence for AI/ML.** AI/ML is driving the next phase of innovation in DevOps. Develop a CoE for AI/ML that includes domain experts and senior IT and business leaders. Based on this survey data, enterprises already see AI/ML as a priority and have a strategy, but challenges introduce doubt in how to operationalize AI/ML.
- **Businesses will use AI/ML to drive benefits in quality and automation.** Machine learning is already driving benefits across DevOps and workload management. Research how ML builds on today's existing tools to bring a more comprehensive and coherent framework for evaluating availability and performance, to identify key objectives and SLAs, recommend and initiate actions, and drive continuous improvement. These are not insignificant gains and in light of hybrid-cloud directions, containerization, and enterprise needs, ML can optimize workload management regardless of environmental complexity—a huge enterprise advantage.
- **Users will determine where to use AI/ML to greatest effect.** Based on your DevOps maturity, determine how best to phase in DevOps tools and leverage AI/ML to maximum effect. Where do you really need AI/ML, and do you have the staff in place to achieve your objectives?
- **Customers need to demand more of their software vendors.** Software vendors could be doing far more than they are today in explaining their AI/ML capabilities and how they add value. Software vendors need to build stronger relationships with their customers involving training, education, and consulting to ensure that customers derive the greatest benefit from their products and are successful using them. AI/ML capabilities from vendors will evolve and change quickly. Task vendors with keeping you up to date on their AI/ML strategy, roadmap, and support in optimizing use of their tools. View this as a joint opportunity between your enterprise and its software vendors and as a way for both of you to be more successful.
- **Don't delay.** Large enterprises have already prioritized their developing AI/ML strategies. Their larger IT resources also mean they can assimilate AI/ML technology at a faster rate. However, AI/ML gives smaller/medium enterprises the ability to compete more effectively with large enterprises. Therefore, don't make AI/ML a lost opportunity because of delays in understanding what it can do for you or delays in implementing the technology.

Conclusion

The prioritization and investment that enterprises have in AI/ML are consistent with the high expectations that exist around the business and functional benefits that enterprises are anxiously anticipating. Could this be a case of inflated expectations that will lead to a trough of disillusionment in DevOps and workload management? EMA doesn't think so. Much of the focus of ML technology today is as a feedback loop based on decisions made, actions taken, and the impact of results. This feedback is then used to reinforce or weaken decision models so that future actions will improve and benefit from a continuous stream of new knowledge gained. ML is being embedded in many DevOps tools today, and the impact can broadly be described as improved tool quality. Quality can mean many things. It can drive better actions, reduce cost, increase performance, increase customer satisfaction, and strengthen productivity. The results of this survey showed many different benefits driven from improved quality.

Today's IT environment, especially in the areas of event, problem, and operations management, is more data-driven than ever. This has given rise to an emphasis on big data, analytics, and advanced analytics. Advanced analytics, such as predictive analytics, mathematical optimization, and Bayesian/neural networks, are all great approaches to decision-making and well understood by operations research professionals and data scientists. Advanced analytics therefore is effective at driving decision-making and actions taken. The union of advanced analytics and ML provides a sound way to drive actions and use the impact of these actions to improve what actions to take in the future.

While ML is very effective at driving improvements to quality, AI will drive another round of innovation by introducing an automation overlay on the advanced analytics/ML decision/action/feedback loop. This automation will begin as a recommendation engine supervised by humans and transitions based on accuracy and confidence to be partially, then fully, automated. The automation that AI brings will deliver another round of benefits, including reduced operational costs, increased speed in delivering new applications and updating applications, and improved employee productivity, all of which reduce SDLC cycle time and leave more time for addressing higher-order development work.

Appendix A: Demographics

This section includes selected demographic data from the DevOps AI survey.

Figure 12

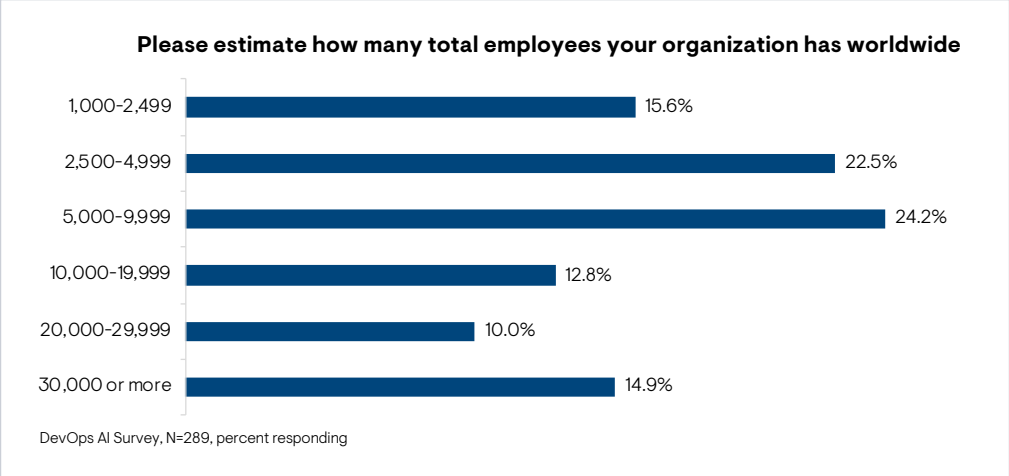

Figure 13

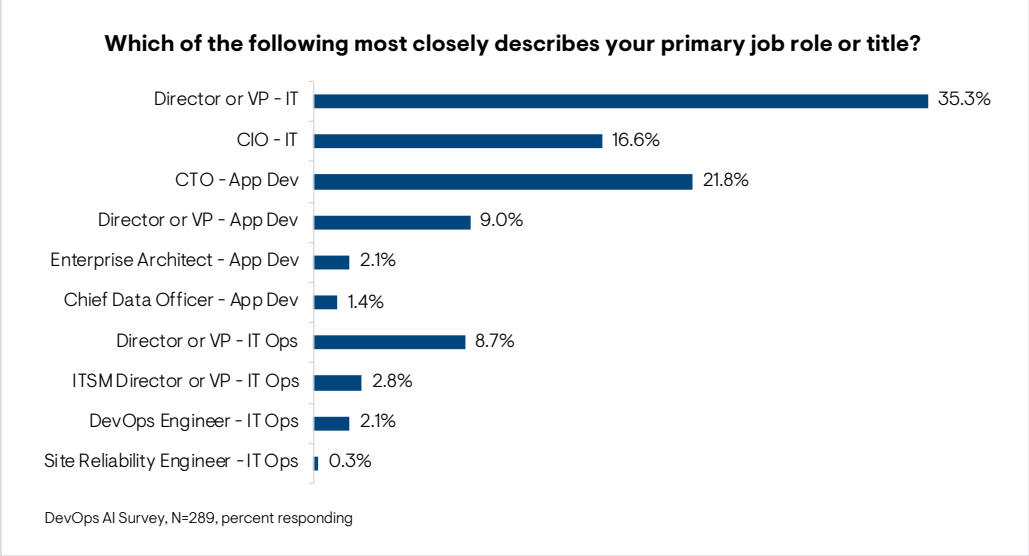

Figure 14

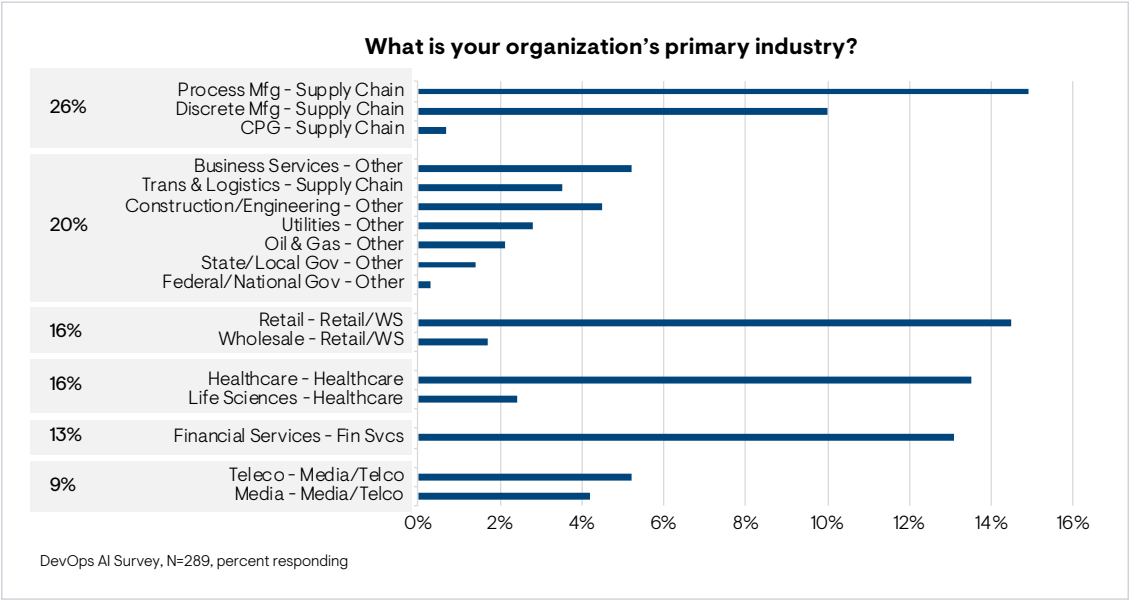

Figure 15

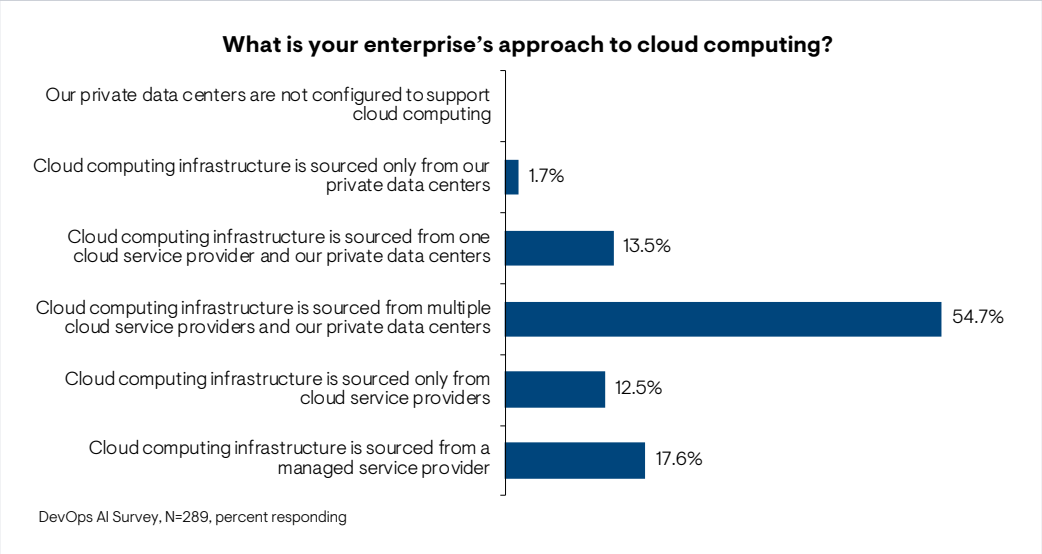

Figure 16

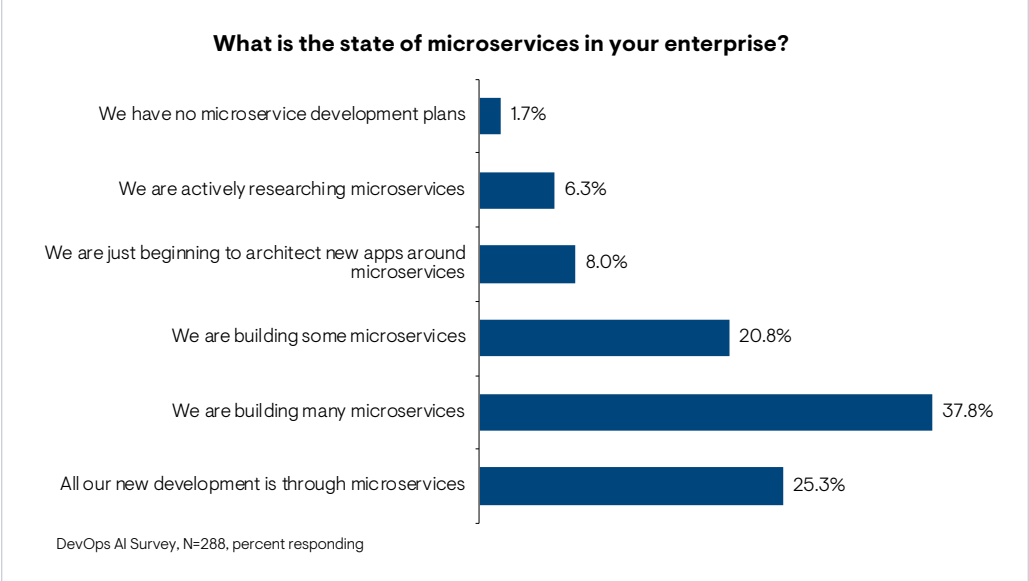

Figure 17

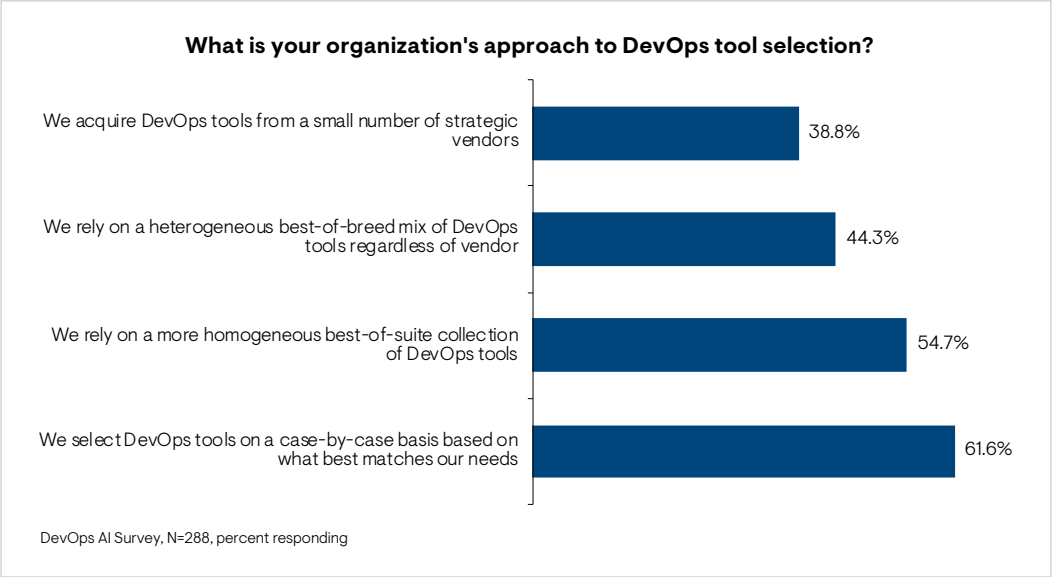

Figure 18

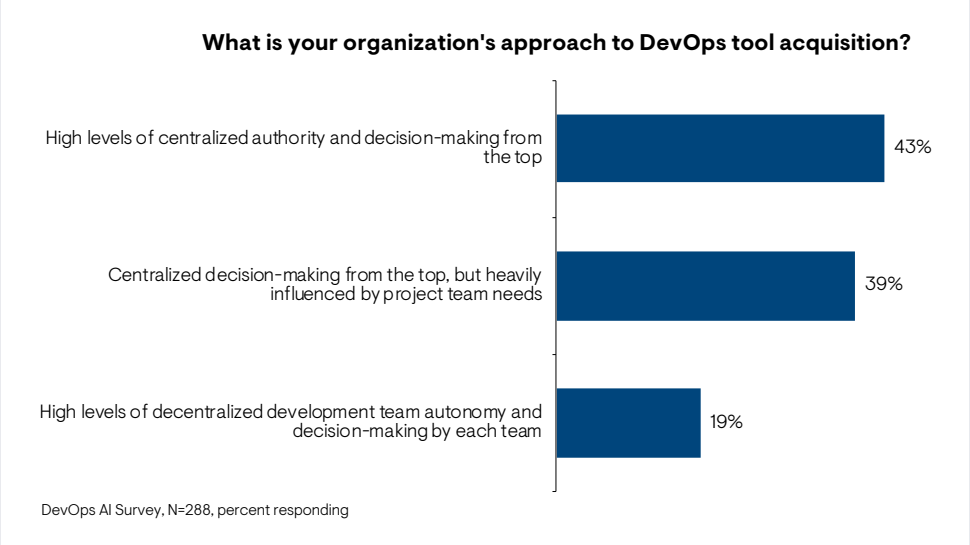

Figure 19

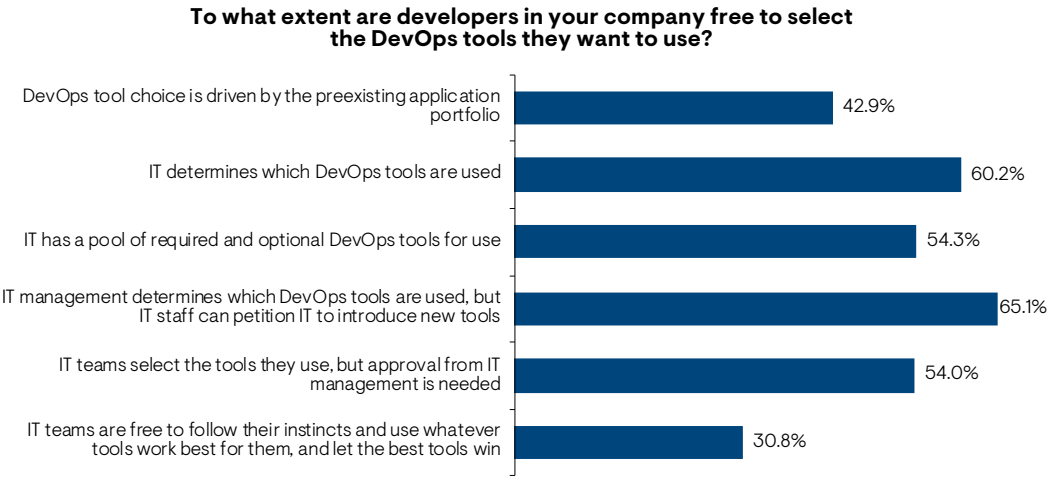

Figure 20

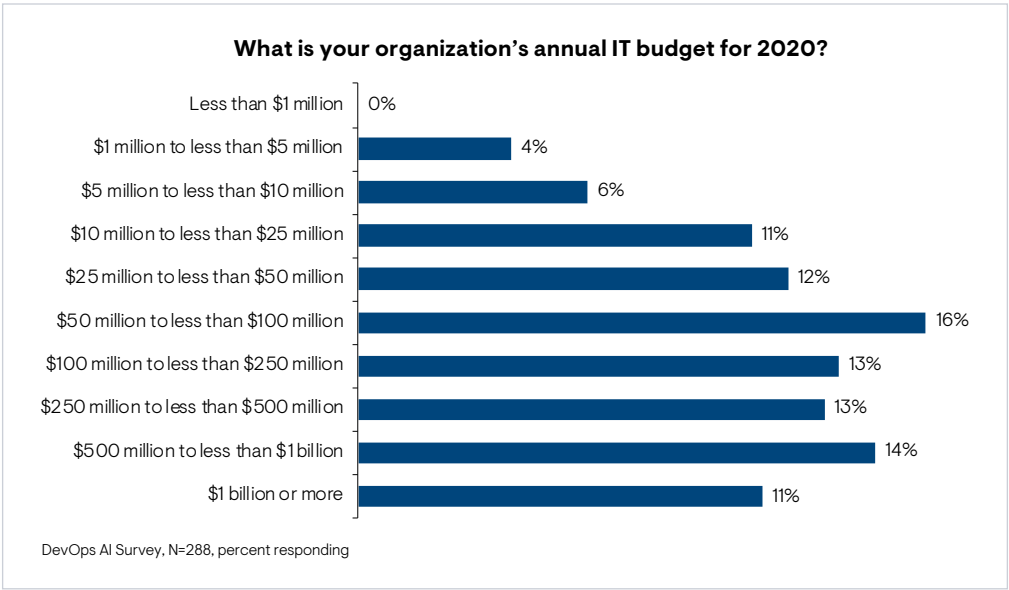

Figure 21

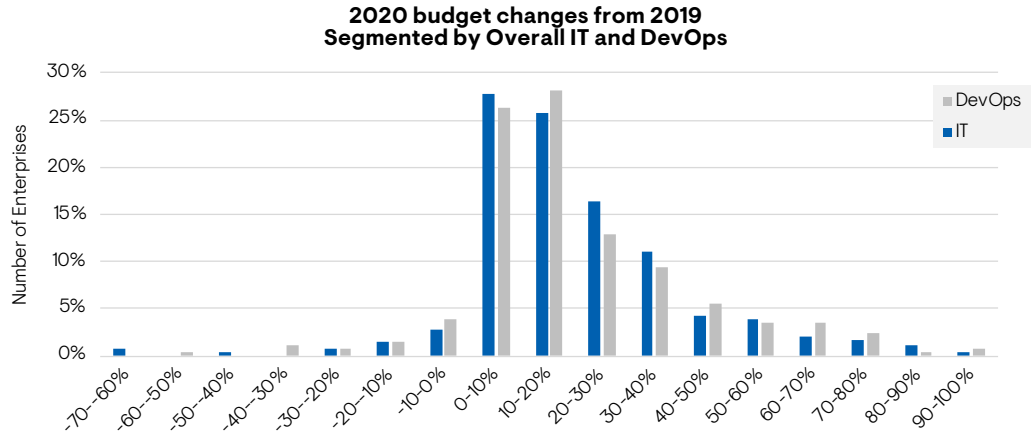

DevOps AI Survey, N=288, valid percent

About Enterprise Management Associates, Inc.

Founded in 1996, Enterprise Management Associates (EMA) is a leading industry analyst firm that provides deep insight across the full spectrum of IT and data management technologies. EMA analysts leverage a unique combination of practical experience, insight into industry best practices, and in-depth knowledge of current and planned vendor solutions to help EMA's clients achieve their goals. Learn more about EMA research, analysis, and consulting services for enterprise line of business users, IT professionals and IT vendors at www.enterprisemanagement.com or blogs.enterprisemanagement.com. You can also follow EMA on Twitter, Facebook or LinkedIn.

Corporate Headquarters:
1995 North 57th Court, Suite 120
Boulder, CO 80301
Phone: +1 303.543.9500
www.enterprisemanagement.com

4050.110420