White Paper

Why AI Governance Is a Business Imperative for Scaling Enterprise Artificial Intelligence

Sponsored by: IBM Corporation

Ritu Jyoti
September 2023

SITUATION OVERVIEW

According to IDC's Worldwide Semiannual Artificial Intelligence Systems Spending Guide, version 1, August 2023 — which tracks artificial intelligence (AI) software, hardware, and services across industries and use cases — enterprises worldwide are expected to invest $151 billion on AI solutions in 2023 and the worldwide AI IT spend will exceed $308 billion by 2026. AI technologies and software are helping companies innovate, transform experiences, build efficiency through automation, and contribute to the improved top line, bottom line, and green line across industries. The advent of generative AI has the potential to revolutionize myriad industries. Businesses that can effectively leverage AI are likely to gain a significant competitive advantage.

Today, AI-powered contact center solutions accelerate time to resolution and improve the customer experience. Deep learning algorithms accelerate diagnosis and treatment of serious illnesses and support personalized medicine. By using AI, industrial IoT can now predict when a machine will break down and recommend preventive maintenance thus avoiding any potential downtime. Artificial intelligence/machine learning (ML) is helping financial institutions improve loan underwriting and reduce risk. AI can also help lessen financial crime through advanced fraud detection and spotting anomalous activity. AI is playing a critical role in cybersecurity today. By improving organizations' ability to anticipate and thwart breaches, protecting the proliferating number of threat surfaces with zero trust security frameworks, and making passwords obsolete, AI is essential to securing the perimeters of any business. Generative AI can empower sales and marketing professionals to produce blogs, social media posts, web copy, sales emails, ads, and other types of customer-facing content. Knowledge management applications are expected to hold the most promise of generative AI for enterprises as the labor intensiveness involved in creating structured knowledge bases has been extremely difficult for many large companies. Generative AI is also poised to transform code generation for software engineers and infrastructure professionals.

AI brings unprecedented advantages to businesses but also an incredible responsibility. Organizations are confronted with potential negative business impact (e.g., damage to brand reputation, reduced public trust, revenue loss, criminal investigations, regulatory backlash, customer privacy loss, and hidden costs) if the AI/ML business risks (e.g., data privacy, discrimination, black box, noncompliance, and safety) are not mitigated appropriately. Adopting AI governance for responsible AI is now a critical business imperative for scaling enterprise AI initiatives.

Lack of AI governance and risk management solutions was identified as a top barrier to AI adoption, second only to cost in IDC's 2022 *AI StrategiesView Survey* of 2,053 organizations (see Figure 1).

FIGURE 1

Top Barriers to AI Adoption

Q. What have been your top barriers to adopting AI?

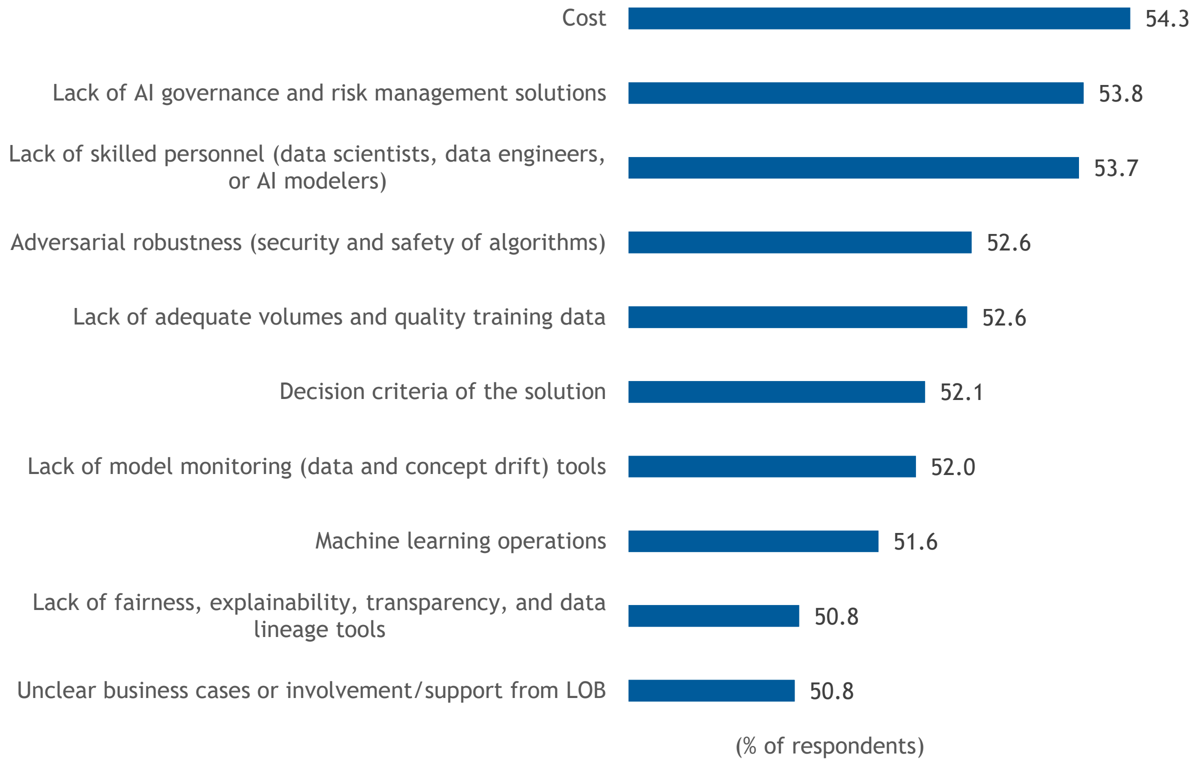

n = 2,053

Source: IDC's *AI StrategiesView Survey,* 2022

Before we delve further on the current state of AI governance and associated challenges, let us level set on definitions discussed in the sections that follow.

Responsible AI

Responsible AI (RAI) is AI that is aligned with human-centered values, fair, explainable, and focused on reducing the unintended consequences of AI by ensuring that the system's intent and use are aligned with the norms and values of the users it aims to serve. RAI aims to empower employees and businesses and fairly impact customers and society – allowing companies to engender trust and scale AI with confidence.

AI Governance

AI governance is essentially the set of processes, policies, and tools that bring together diverse stakeholders across data science, engineering, compliance, legal, and business teams to ensure that AI systems are built, deployed, used, and managed to maximize benefits and prevent harm. AI governance allows organizations to align their AI systems with business, legal, and ethical requirements throughout every stage of the ML life cycle.

AI GOVERNANCE CRITICALITY

IDC estimates that most G2000 companies have more than 100 machine learning models in production, with some having thousands. The impact of these models is increasing too. IDC surveys suggest that at least a quarter of G2000 companies credit their AI capabilities with contributing to more than 5% of their earnings. It is no surprise, then, that AI investments are likely to continue to grow as around half of G2000 firms plan to increase AI spending, viewing it as a strategic imperative.

The growing breadth, complexity, and scale of enterprise AI is beyond the ability of any one team of data specialists to manage. DataOps, MLOps, DevOps, and even AI life-cycle tools, often on cloud, are increasingly being used to support the increasing range of activities associated with deploying complex models. These activities range from data engineering to data science to DevOps to model management and compliance. They usually involve many differently skilled specialists with different agendas working together across various parts of the enterprise. But this is not enough to ensure that management, customers, and regulators trust model outcomes to satisfy important requirements such as accuracy, consistency, fairness, and explainability. To add to it, it is interesting to note that many organizations resort to manual AI governance leveraging spreadsheets, which is time consuming, prone to errors, and lacks scalability.

Model risk management and governance are partly driven by an increasing number of compliance requirements such as the European Union (EU)-based GDPR as well as management's fears of potential legal liabilities/adverse publicity associated with some models. They also reflect an attempt to manage and optimize this complexity of AI models. Risk management is usually separated from those that originate the risks (separation of duties). Traditionally, in many firms, this places risk management in centralized functions such as treasury or when the risks are less pervasive in operations. As AI becomes more strategic and critical to the business, we believe it will become a strategic concern for the ultimate stakeholder regarding the risks of the enterprise. The CFO/CRO/CPO will insist on taking control.

Over the past several years, regulators across the globe have started passing legislation or providing regulatory guidance on AI systems. Regulations governing AI vary across countries and are rapidly evolving. The European Commission is widely viewed as leading these efforts through its attempt to pass a comprehensive, cross-sectoral AI regulation. As per our research, EU AI Act is noted as the number 1 regulation critical for organizations' AI implementations globally. In addition, regulators in Hong Kong, Singapore, the Netherlands, and the United States — among many others — have been outspoken on the need for appropriate corporate governance to address AI-related risks, including risks relating to bias, model drift, privacy, cybersecurity, transparency, and operational failures.

One notable feature of several emerging regulatory pronouncements, particularly in the financial sector, is their express focus on the importance of board-level oversight of AI risks. For example:

- The U.K. Financial Conduct Authority and Bank for International Settlements have both recently underscored that boards and senior management will need to tackle some of the key issues emerging from AI because that is where the responsibility for AI risk will ultimately reside.
- The Monetary Authority of Singapore has suggested that firms should set approving levels for highly material AI decisions at the chief executive officer or the board level and should periodically update the board on the use of AI within the company so that the board maintains a central view of all material AI-driven decisions.
- The NY DFS recently required each New York domestic insurer to designate one or more members of its board and its senior management to be responsible for oversight of the insurer's management of climate risks, and it is likely that similar regulatory requirements for AI risks are coming.

These are some of the recent examples of the coming wave of regulators demanding board-level responsibility for overseeing the regulatory, operational, and reputational risks of AI.

Customer privacy loss, revenue loss, damage to brand reputation, and hidden costs are the top negative impacts that organizations are concerned about due to the lack of use of RAI (see Figure 2). Customer privacy loss issues have gone from compliance to negative business headlines. Regulators are looking to the C-suite and the boardroom to demonstrate leadership that will promote innovation and consumer trust. AI governance is paramount in realizing RAI and plays a critical role in advancing the business value of AI.

FIGURE 2

Top Negative Impacts Due to the Lack of Responsible AI

Q. What negative impacts due to the lack of use of responsible AI are you most concerned about within your organization?

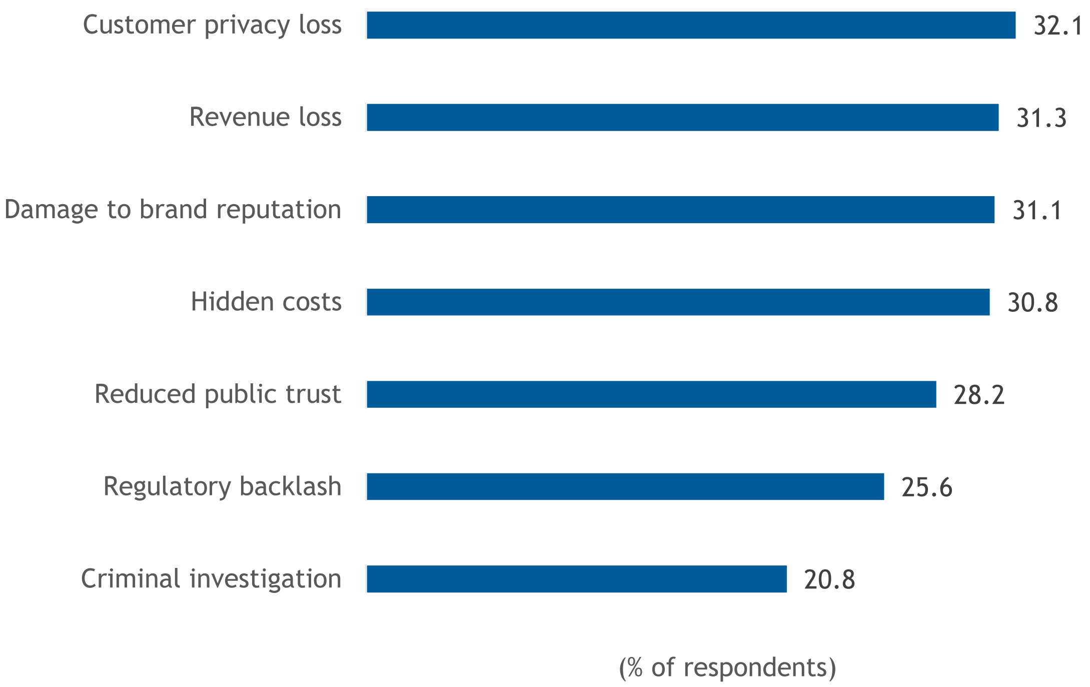

n = 2,017

Source: IDC's *AI StrategiesView Survey,* May 2022

FOUNDATIONAL ELEMENTS OF AI GOVERNANCE

Given the market sentiment and requirements around RAI and understanding that trust in AI can help drive business growth and innovation, enterprises will need to ensure the following components in their AI governance:

- **AI/ML life-cycle governance:** While data scientists and machine learning engineers oversee the machine learning operation, model validators and approvers are responsible for monitoring the model in production for fairness, drift, quality, and explainability and automated collection of metadata about model development and experiments including training data and model validation/approval metadata collection.
- **Collaborative risk management and compliance:** Risk management is an integral part of corporate governance, risk, and compliance (GRC). Companies must ensure that AI models adhere to their corporate business rules to avoid any revenue loss or unintended negative consequences as well as AI regulations and that a policy engine can define and enforce

automated AI governance policies and rules. Essentially, GRC standards should be applied to all AI projects, and there should be repeat monitoring and certification every time the model changes.

- **Regulatory excellence:** Regulatory excellence for AI comes through digital collaboration across an organization and by exploiting leading-edge tools and technologies that can help organizations do proactive impact assessment of regulations and do model audit and reporting.

AI GOVERNANCE PLATFORM

AI disrupters are responsible by design and understand the importance of incorporating RAI into their AI strategy from the start, right from the ideation of an AI initiative (i.e., use case identification). They operate an RAI approach and ensure it with AI governance across the complete life cycle of all their models (in-house or third party), enabling the organizations to engender trust and scale AI with confidence while better mitigating business risks, meeting regulatory requirements, and creating sustainable value for themselves and their stakeholders. To be responsible by design, we believe that organizations need to move from a reactive compliance strategy to the proactive development of mature, responsible AI capabilities. With the foundations in place to support the responsible use of AI across the enterprise, it becomes easier to adapt as new regulations emerge. That way, businesses can focus more on performance and competitive advantage.

AI disrupters have a holistic AI governance framework that will ensure that the right processes and technologies are employed. Training the entire workforce in AI governance is a critical component. As we know, RAI is not just about establishing the appropriate governance structures. It is also important to translate those ethical and legal frameworks into statistical concepts that can be unambiguously represented in software. It is equally important to understand the synergistic role of MLOps and AI governance and how RAI platforms complement the MLOps platforms and vice versa.

MLOps provides critical inputs to AI governance through its infrastructure and tools. Effective governance requires visibility into end-to-end AI system behavior. Without MLOps tooling, getting your AI governance team the insights they need to monitor, understand, and mitigate AI system risk is challenging at best. Having MLOps infrastructure and processes makes the technical evaluation of AI systems against governance-driven requirements much easier and more scalable.

To address one of the biggest barriers in scaling AI responsibly (i.e., complexity) — and an undertaking that involves multiple stakeholders and cuts across the entire enterprise and ecosystem — AI disrupters use/are planning to use an RAI governance platform with critical capabilities that can help them unlock superior business benefits (see Figure 3).

FIGURE 3

Critical Capabilities of an RAI or AI Governance Platform

Q. What do you think are the critical capabilities of an RAI or AI governance platform?

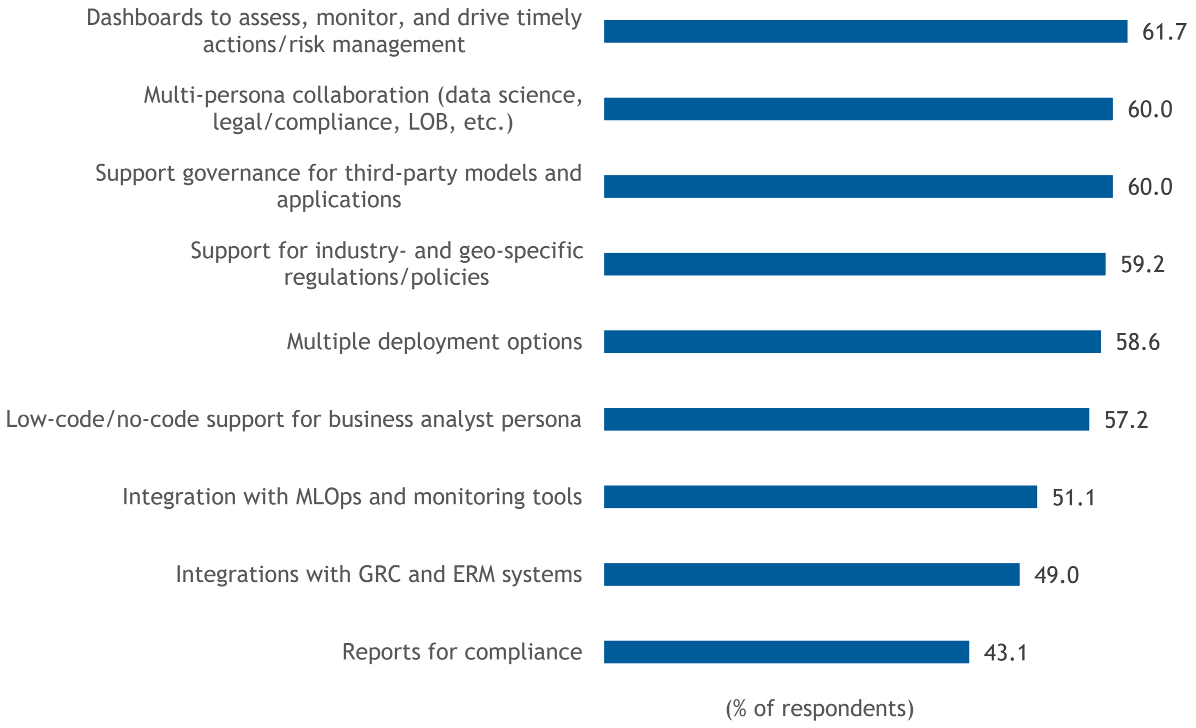

n = 1,403

Source: IDC's *AI StrategiesView Survey,* May 2022

RAI or AI governance platforms help organizations standardize governance, reduce overhead burden, help ensure regulatory compliance, and increase stakeholders and customer trust through model explainability and transparency. They bring together the diverse stakeholders involved in developing and using AI systems, using the right tools to collect the right data, and making better decisions about how these systems should be built, deployed, used, and managed to maximize benefits and prevent harm. Essentially, they help maximize AI success. To support multi-stakeholder collaboration, organizations are looking for a common dashboard that can help them assess, monitor, and drive timely actions. With newer regulations emerging rapidly, support for geo- and industry-specific regulations in the platform is critical so that they can exercise appropriate due diligence from the start (even before initiating an AI project) and at every stage of the life cycle. AI governance platforms should seamlessly integrate with an organization's standard governance, risk, and compliance and enterprise risk management (ERM) solutions.

VOICE OF THE CUSTOMER

As part of this research, IDC interviewed multiple enterprise customers that have been using AI responsibly.

The following list outlines example customer scenarios and highlights insights from their AI governance adoption lessons learned and derived business benefits:

- A leading custom application development company and IT consulting agency are using AI for HR, customer service, and IT operations. The company is using a commercial AI governance solution that has custom integration with its corporate GRC system. The CIO of the organization owns the overall AI budget, and AI governance tools and technologies are included in it. It is around 15% of its total AI budget. Although, the company embraced AI governance as an afterthought for the first few initiatives, it is rapidly moving to embrace it as a forethought for the company's new initiatives. The CIO collaborates with the head of legal/compliance and LOB heads to ensure alignment to business rules and requirements. According to the CIO of the organization, "Using AI governance platform helps engender trust in AI and helps us be proactive and make more informed decisions in a faster and easier way."
- A retail data science, insights, and media company is leveraging AI to gather insights for CPGs, pricing optimizations, supply chain management, and personalization of customer experiences by recommendations of products, placement, coupons, and so forth on digital channels. The company has been focused on building a robust strategy around governing the data pipelines and automating the testing and deployment of models into production, as well as measuring the drift post production. This is done to ensure that there is no price gouging. A significant part of its AI budget is spent on tools to ensure end-to-end AI governance as its current landscape of governance, risk, and compliance is quite fragmented and incomplete. According to the SVP, Engineering, of the organization, "AI governance tools are doing a great job of shortening the time to earn stakeholders' trust because of the transparency around the process, the data, and the outcome." She also noted that the company is exercising due diligence and experimenting with generative AI, as it holds enormous potential to transform its operations and customer experience.
- A global financial service organization in the capital market space is using AI for anti-money laundering and fraud detection. The solution is deployed across multiple public clouds, and the company is using commercial AI governance solution that has been custom integrated with its GRC system. Around 25% of its overall AI budget is spent on AI governance tools. According to the CIO cum CISO of the organization, "The business benefits of AI along with AI governance are pretty apparent. We are seeing rapid ROI in terms of avoiding business losses and productivity gains."

CONCLUSION AND IDC GUIDANCE

We have now entered the world of AI-augmented work and decision making across all the functional areas of a business, from front to back office. AI, machine learning, and natural language processing are changing brands around the globe across multiple industry sectors. AI disrupters will scale AI initiatives, drive better customer engagements, and have faster rates of innovation, higher competitiveness, higher margins, and superior employee experiences. Organizations worldwide must evaluate their vision and transform their people, processes, technology, business models, and data readiness to unleash the power of AI and thrive in the digital era.

While many organizations have taken the first step and defined AI principles, translating these into practice is far from easy, especially with few standards or regulations to guide them. Responsible AI implementation continues to be a major challenge. Taking a systematic approach from the start while addressing the associated challenges in parallel can be beneficial. A systematic approach requires proven tools, frameworks, and methodologies, enabling organizations to move from principles to practice with confidence and supporting the professionalization of AI.

An organization that wishes to accelerate the AI adoption and time to value with responsible AI should:

- Evolve from being reactive to being proactive.
- Ensure that responsible AI governance starts at the executive level. The C-suite has a critical role in establishing the organization's policies, charter, and accountability.
- Be data driven with a focus on mitigating bias and improving data quality.
- Establish a disciplined model development, management, and monitoring process for all AI models.
- Create organizational transparency with strong ongoing AI governance methodologies.
- Develop risk metrics that can be incorporated into larger enterprise risk frameworks. Create a cross-functional group of AI experts to proactively address AI-specific risks and biases aggressively. This group should include IT as well as those in business and compliance functions.
- Coordinate the drivers for change for responsible AI inside and outside of the organization. Governance guidelines should clearly address why they are in place and leave little room for interpretation. Teams across distinct functions, such as leadership, data science, and legal, must understand the imperatives and incentives of each.
- Become *responsible by design* to help scale AI with confidence. By shifting from a reactive AI compliance strategy to the proactive development of mature responsible AI capabilities, organizations will have the foundations in place to adapt as new regulations and guidance emerge. This way, businesses can focus more on performance and competitive advantage and deliver business value with social and moral responsibility.
- Partner with a trusted and innovative technology supplier and professional services firm like IBM. IBM is focused on helping customers transform responsible AI from theory to practice.

MESSAGE FROM THE SPONSOR

Why IBM for AI Governance

IBM watsonx.governance is a one-stop automated software toolkit built on the IBM watsonx platform and designed to direct, manage and monitor the AI activities of an organization. Operationalizing AI helps to streamline processes, minimizing time consuming and costly human errors while driving the ability to responsibly scale model production.

There are three components to watsonx.governance. The first is the ability to monitor, catalog and govern across the AI lifecycle, increasing model transparency and predictive accuracy. The second is risk management which includes proactive identification and mitigation of bias/drift while automating facts and workflows to help comply to business standards. Lastly, watsonx.governance helps organizations adhere to compliance by automating the translation of external AI regulation into policies for automatic enforcement. User based customizable dashboards, dynamic reports and collaborative tools streamline and hasten processes while aligning the growing number of an organization's AI stakeholders.

IBM Expert Labs team offers the tools, assets and partnership customers need to expedite implementation. Working across all stages of the AI lifecycle, from planning to build, deploy to operate. Expert Labs help deliver trusted AI solutions at scale and speed.

Book a meeting with an IBM expert

About IDC

International Data Corporation (IDC) is the premier global provider of market intelligence, advisory services, and events for the information technology, telecommunications and consumer technology markets. IDC helps IT professionals, business executives, and the investment community make fact-based decisions on technology purchases and business strategy. More than 1,100 IDC analysts provide global, regional, and local expertise on technology and industry opportunities and trends in over 110 countries worldwide. For 50 years, IDC has provided strategic insights to help our clients achieve their key business objectives. IDC is a subsidiary of IDG, the world's leading technology media, research, and events company.

Global Headquarters

140 Kendrick Street
Building B
Needham, MA 02494
USA
508.872.8200
Twitter: @IDC
blogs.idc.com
www.idc.com