

"Introduction to AI in Business - The 5th Industrial Revolution" is a 3-part series presented by Certified Information Security's founding Principal and content author, Allen Keele. Understand the critical aspects of AI accountability, safety, and risk management in business settings. Discover key frameworks such as ISO 42001 AI Management Systems, ISO 23894 AI Risk Management, and NIST AI Risk Management Framework 1.0. Uncover the compliance mandates of the EU AI Act of 2024. Explore other sessions in this 3-part series:

Session 1 - Embark on a journey through the basics of AI in simple terms. Dive into the 2025 AI trends and its integration into business operations, all supported by well-recognized and reputable sources.

Session 2: Delve into the reasons behind organizations incorporating AI and its practical applications across various industries like healthcare, retail, services, manufacturing, quality management, and supply-chain management.

For insights on establishing and overseeing AI within your organization, visit <http://www.certifiedinfosec.com>.

Welcome to this short course on artificial intelligence in business, the fifth Industrial Revolution. I'm Allen Keele. Thanks for joining me.

We'll begin the course by getting a basic understanding of artificial intelligence in the business context. We'll then expand upon that by going into AI adoption and trends globally in 2024. We'll then look at how to benefit from and leverage AI in an organization. We'll then get more specific examples of how organizations are benefiting from integrating AI in various industry sectors. We'll then move on to managing AI risk, AI, accountability, safety and risk. We'll then learn more about how we can leverage existing newly emerging standardized frameworks to plan, integrate, manage and improve AI in business. And then finally, we'll go into late-breaking newly emerging AI legislation and regulation that governs how AI can be used in business.

Moving from technology to smart technology can bring a lot of wonderful benefits. But it can also bring some risk that we need to consider as well. We're ready to start the AI accountability and safety. Risk module to learn more about how we manage the risk that comes along with the benefits of AI. In 2016, theoretical physicist and mathematician Doctor Stephen Hawking said that the potential benefits of creating intelligence are huge. That every aspect of our lives will be transformed in short success in creating AI could be the biggest event in the history of our civilization. But it could also be the last, unless we learn how to

avoid the risks. In short, the rise of powerful AI will be either the best or the worst thing ever to happen to humanity. We do not yet know which. Achieving the benefits of AI by integrating AI and leveraging AI in our business processes.

Phil requires us to do so responsibly, which means managing risk of AI. Risk management can enable AI developers and users to understand impacts and account for the inherent limitations and uncertain. In their models and systems, which in turn can improve the overall system performance and trustworthiness, and the likelihood that AI technologies will be used in ways that are beneficial. While some a high risk and benefits are well known, it can be challenging to assess negative impacts and the degree of impact. One example of a risk that we may not be familiar with would be AI hallucinations, which results from AI and misinformation. The problem is artificial intelligence hallucinations are when AI invents false information. That don't reflect real data and the truth. An example of this would be fictitious quotes or inaccurate company figures, or flat out bogus facts. The causes would include outdated or incorrect training data, incorrect training, data categorization. Errors, contradictions and distortions in the training data. Insufficient context from users. Or difficulty. Recognizing human sarcasm and slang. So then this results in risks. One of the biggest risks when using AI systems is trusting AI and its accuracy. We've talked about this before. Hallucinations that blend misinformation with facts. Can spread on the Internet at breakneck speed. Only careful and responsible fact checking can prevent the spread of AI hallucinations. So let's look at some examples of existing problems today, along with tools that can help us address these issues. For example, deep fakes. Deep fakes are AI generated media such as video. Audio text that appear to be realistic but actually depict false or manipulated content. Tools such as anti fake or fake catcher can help detect this content. We also have AI generated text AI already imitates human writing patterns and styles quite well. That's why tools that recognize the difference between. Humans and machines are becoming essential detect ChatGPT is a tool. That helps detect AI generated text. You can imagine the challenge. With professors who are checking term papers of students who may be using. Generative AI to actually write the entire term a. We also have potential copyright infringements that may not be purposeful on the human user who is depending upon AI for text generation not realizing the copyright infringement that came along with it. Many AI systems use training data content to generate their own content. And we can also wonder if that new content is considered original content and potentially copyrightable. One proposed solution is to use watermarks on AI generated content. So for example, when AI generates an image. Copyright to that image. Facial recognition is also becoming increasingly popular, but not every organization uses this information responsibly. Tools such as Fox Protect data privacy and canned outsmart facial recognition systems. We also have various dangerous AI behavior. Given their rapid development, AI systems, ethical standards also need to be scrutinized, and companies such as DeepMind are developing tools that can detect and classify. Harmful AI behavior. So as we've learned, AI is not a collection of discrete processes that happen independently away from our operational and organizational processes. Is actually integrated into the organizational processes. Or perhaps even replacing processes. So we need to manage the risk of this along with these processes, not separate from. So AI risk management needs to be integrated. Upgraded and incorporated into broader enterprise risk management strategies

and processes. Treating AI risk along with other critical risks such as cybersecurity risk and privacy risk. Will yield more integrated outcomes and organizational efficiencies. We've learned that leveraging and integrating AI can bring huge and phenomenal benefits and competitive advantage to our organizations. But if we don't manage the risk of AI. It can hurt. Our organizations and business equally so, so. We need to recognize managing risk better will help us to avoid risk better and therefore maximize the positive return from AI well. How do we manage risk better, it turns out. There are existing newly released AI risk management frameworks and management systems that give us a road map to doing this better. Why not leverage these frameworks to manage risk? Better as we integrate AI into our organizations. Let's look at existing frameworks that have been recently released in 2023, including ISO 42, Double 01, AI management system frameworks, as well as the ISO 23894 guidance and AI risk management. Along with. The NIST AI Risk Management Framework 1.0 that was developed and released by the National Institute for Standards and Technology in the United States in 2020. T3. We'll begin by looking at the ISO 42 double 01 standard ISO being the International Organization of Standardization. You may recognize them as the organization behind ISO 9001 quality, ISO 27, Double 01, Information Security. ISO 22301, business continuity management. Well, they have a management system, framework standard and certifiable specification, ISO 42, double 01. To help us establish a program for AI management, O again I like to. This you can't manage your finances with accounting if you haven't even established your accounting program or your finance department in a similar way, how do we integrate, manage and maintain? AI in the organization, if we haven't even established a platform for governing it, planning it, implementing it, monitoring it and improving it. That's what ISO 42 double 01 can help us do. In today's rapidly evolving technological landscape, businesses are constantly seeking ways to stay ahead of the curve. One of the most promising avenues, as we've seen for achieving this is through the implementation of artificial intelligence. However, with great power, as we mentioned, comes great responsibility and this is where ISO 42 double 01 comes into play. This ISO international standard provides a comprehensive framework for managing AI systems. Ensuring that AI is used responsibly and effectively to drive business success. And as we said, we need to practice AI integration responsibly by managing the risk that comes along. With it an ISO. 2 double. One incorporates that principle as well. AI risk management is integrated into the organization's broader enterprise risk management. So we could look at it graphically as this we have enterprise risk management as well. Holistic high level framework for managing any flavor of risk, whether it be occupational health and safety risk, environmental risk, compliance risk. Cyber security risk AI risk. Can we have a program that will basically integrate and streamline our approach to managing all those different flavors of risk? Having one set way of doing things that they all have in common? But still, allowing each specialty risk program to consider the unique nuances of AI risk. Versus environmental health and safety risk, for example. This is what enterprise risk management can give us and we have a standard for that called ISO 31,000 enterprise risk management. So an organization establishes. Consistent management of risk throughout the enterprise to guide and harmonize these niche types of risk, such As for AI, cybersecurity, privacy and compliance risk. So within that broad framework, we can sharpen it. Can tune it for AI. Risk

management. And AI process management with an AI management system. So ISO 4201 would be used to extend ISO 31,000 risk management to managing artificial intelligence. So we begin with enterprise risk management and then within that we could have ISO 42, double O 1 entitled, Information Technology, Artificial Intelligence Management System. To then establish a formal system strategy and policy, complete with roles and responsibilities for leveraging effective and responsible use of AI throughout the enterprise. And then within that, we can manage the risk of AI with. Another ISO standard 23894 or the NIST AI risk management framework to address risk a potential desired and undesired outcomes in the development and use of artificial intelligence throughout the enterprise. ISO standard 42, double O 1 is a management system standard for organizations working with artificial intelligence. It supports organizations in developing responsible and trustworthy AI systems. The standard has requirements for understanding and managing AI risks. Now it is supported by another guidance document called ISO 23894 entitled Information Technology. Artificial intelligence guidance on risk management. Which is the ISO AI risk management framework especially tailored for identifying, analyzing, evaluating and controlling AI risk within an overall. AI management system that is actually dictated by ISO 42001 OK Certified Information security of course does provide training and competence validation, which is professional examination and credentialing. For ISO 4201 systems development, implementation and management. And of course, you can visit certified information security to learn more about this training and certification program. Again, ISO 23894 is there to help us fine tune our risk management. To understand how to better manage the extra nuances of artificial. So ISO 23894 was released in 2023. To provide guidance on how organizations that utilize artificial intelligence to develop, produce, deploy or use products, systems and services. Can manage risk specifically related to AI. The guidance also aims to assist organizations for integrating risk management into their AI related activities and functions, as well as describing processes for the effective implementation and integration of AI risk management. The application of the guidance can be customized to any organization and its context. And of course, certified information. Also provides training and confidence validation for ISO 23894 as part of its ISO 31,000 enterprise risk management training and certification program. And of course you can go to certified information security to learn more about this training and certification program. Most countries have a standards Bureau that well. Provides the basic understanding of what standards will be used in that particular country, whether it be for electricity or technology or products in the United States, the National Institute of Standards and Technology provides a service. Was actually called the National Standards Bureau up until as recently as 1988. At any rate, NIST was tasked with creating. The NIST Risk management framework. To help guide the responsible development and implementation of AI in the United States, so as directed by the national Artificial Intelligence Initiative Act of 2020. The goal of the AIRMF or AI risk management framework is to offer a resource to the organizations designing, developing, deploy, or using AI systems and to help them manage the many risks that they are. I and two then promote trustworthy and responsible development and use of AI systems. The NIST AIMF is intended for voluntary use and to provide the ability to incorporate trustworthiness considerations into the design, development, use and evaluation of AI products, services and systems. Again, this is

essentially a standards Bureau, not a regulator. Later, O they create standards for organizations in government to use at their own discretion. Since NIST isn't a regulator, they don't force the issue with anybody. Having said that, that doesn't prohibit. Regulatory authorities and licensing organizations from adopting NIST standards as their own specifications for enforcement. So for example, we may have a nest standard that underwriters laboratories adopts as its own requirement. For specifications in electrical equipment and appliances, as an example, O again when we say that the framework is designed to be voluntary, well, that's voluntary. According to NIST. It could be that we have another. Organization that uses an is standard as a mandatory required specification away from NIST. So the framework designed by NIS is intended to be rights preserving, meaning it's intended to respect individual rights and freedoms. It is intended to be non sector specific. It's applicable across different industries and sectors. Therefore, you can be assured that this will help any organization in any sector to manage AI risk better. It is also use case agnostic. Is intended to be adapted to various AI use cases, applications and contexts. So the AI RMF is structured to support organizations in increasing trustworthiness of AI systems and. Fostering responsible practices over time. It's a living document. It's expected to evolve with feedback from the AI community and changes in AI technology. Currently we're working with NIST AIRMF version 1.0. NIST never intended the AI RMF to be the end. All be all of AI risk management that it was to be used. Free standing and exclusion of any other risk management practices or frameworks. The AIRMF may be utilized along with related guidance and frameworks for managing AI system risk or broader enterprise risk. In other words. A IRMF is perfectly at home working within an ISO 31,000 enterprise risk management. Ystem or within and alongside ISO standards for AI management and risk management. Some risk related to AI systems are common across other types of software development and deployment. Examples of overlapping risk include privacy concerns. Related to the use of underlying data to train AI systems. The energy and environmental implications associated with resource heavy computing demands of AI, cybersecurity concerns related to the confidentiality, integrity, and availability of the system and its training and output data. And general security of the underlying software and hardware of AI systems. Now, while ISO 23894 provides a nice accompaniment to ISO 31,000 to essentially give ISO 31,000, the extra considerations for AI risk that ISO wants us to consider. For ISO 42, double 01 compliance. NIST Aimf goes far, far deeper. So while it takes five days to properly teach NIST aimf, I can give you a nice high level overview in a couple of slides. Now NIST AI rmf. Has. Four functions in its core and its core is essentially the framework and the process of the program all combined. It starts off with the inception or the foundation of the program. The strategic direction of the program, which is in governance after all, if you haven't made decisions in governance. On what it is you want to manage and how you want to manage it, then you have no business practicing any process to fulfill what you don't know that you want. So it starts off with a lot of govern requirements for governing AI risk. This is where senior leadership has a lot of decisions to make, and This is why NIST AIRMF can't be simply delegated off to middle management and techies. There are a lot of governance decisions and input that needs to be accomplished. OK, so once we've established the framework, the program and the preferred process for practicing AI risk management, then we go on to

actually identifying AI risk and doing some initial. Analysis of AI vulnerabilities and threats. Identifying them, assessing them for credibility, assessing them for probability, and then what we will do is progress on to trying to then quantify whether it be in a quantified risk assessment or a qualified risk assessment. What kind of level of AI risk that we actually have and whether that's beyond what we can tolerate? This is aligning with the risk assessment process of clause 6. In ISO 31. So This is why if you understand both frameworks, you can see how one could fit into the other. The dilemma is, is they kind of use different terminology and they didn't really map. The The four NIST AI RMF functions. To the framework and process of ISO 31,000. If you understand both, you can understand how to correlate them, which is what I kind of do in class. I crosswalk one against the other, but at least they're following the same basic approach. So in the measure function, what we're doing there is again trying to analyze what we know of the risk to determine how much risk is there and whether it's beyond what we can tolerate and if so. We move on to managing the risk and that's actually risk treatment is what we would call it in ISO 31,000. So this is where risks are then prioritized and then treated or acted upon based on projected impact. Now of course, as I've mentioned, certified information Security does provide training and competence validation for NIST AIRMF, and you can go to certified Information Security's website to learn more. Well, you heard me mention that NIST AI RMF is more extensive than what ISO 23890. Gives us. Well, how much so? Be aware of what you ask for. NIST AIRMF has 19 subject matter categories. That actually encompass 76 desired outcomes or objectives that they call subcategories and then following that they support these 76 desired outcomes or objectives with 460 recommended implementation actions. And reporting processes. So if you really wanted full detail on how to assess and manage. AI risk. I think that this framework is an excellent accompaniment to what we do in ISO 23894 within an ISO 42 double 01 management system. In Europe, we have actual law governing. Use and implementation development of artificial intelligence. We have the new EU artificial intelligence. Act. President of the European Commission, Ursula von der Leyen actually explained that the EU AI act is the first ever comprehensive legal framework on artificial intelligence worldwide. So enacting this new artificial intelligence legislation is indeed a global historic moment. They need to control and constrain artificial intelligence, and robotics is not a fresh concern. In fact, famous physicist Isaac Asimov released and published his three laws of robotics in 1942. He explained in the first law that a robot may not injure a human being or, through inaction, allow a human being to come to harm. Secondly, a robot must obey the orders given it by human beings. Except where such orders would conflict with the first law, and then finally, a robot must protect its own existence. U to the limitation of making sure that such rotection does not conflict with the first or second law. Advancing to today, the EU Parliament's Committee on Legal Affairs put the EU Artificial Intelligence Act into force in August 2024, the first comprehensive regulation on AI by a major regulator anywhere. This act assigns applications of AI into three risk categories. Firstly, applications and systems that create an unacceptable risk, such as government run, social scoring of the type already used in China, must be banned. Secondly, high risk applications such as CV or resume scanning tools that rank job applicants must be subject to specific legal requirements. And then finally, applications not explicitly banned or listed as high risk. Will then be largely left unregulated. The Euai Act did

not simply spontaneously come upon. It's been a development for quite some time. It began in April 21, 2021, where the EU Commission proposed the AI Act. Then it went through various stages of consideration until February 2, 2024, where the EU Council of Ministers unanimously approved the draft law. So that draft law went further through consideration, publication more consideration until finally March 13, 2024 EU. Approved the AI act. So that means that the AI act has been approved. When does it go into force? Nine months after entry into force, code of Conduct is applied 12 months after entry into force. Governance rules and obligations for general purpose AI become applicable 24 months after entry into force. The start of application of the Euai Act for AI Systems comes into play. And then finally, 36 months, three years after the entry came into force back in March 2024, the application of the entire AI act is enforced for all risk levels. Let's look at the 8 fundamental objectives of the EUAI Act. Firstly, AI systems should be safe and transparent for consumers. Investments in the AI sector should be encouraged, not discouraged. AI systems must be classified into risk levels that are actually detailed in the AI act itself. Also, security, health and fundamental rights must be protected and the environment should also be protected from negative consequences. Of AI use. Oversight for AI systems must be in place to avoid discrimination. Companies and other deployers of AI should be able to use AI in a regulated capacity. And finally, AI must be well understood by everybody concerned. So who is affected by the Euai act well? You may be affected by the AI act whether you operate within the EU or even if you are outside of the EU in terms of geographic area, the Act applies to artificial intelligence within the European Union, however. AI outside of the EU can still be covered by the AI Act if the citizens of the European Union can access and use the AI. So this means relevant parties, the EU act for AI applies to providers. Product manufacturers, importers, distributors and providers of AI systems, which pretty well covers any organization involved in AI in any way. There is no fixed turnover or user threshold for the applicability of the act, and that translates to large or small the AI act applies if your organization. Is involved in. Providing manufacturing, importing, distributing. AI exceptions. Well, we have AI systems that are excluded from the scope of law. Such as systems used exclusively for scientific research. Systems used purely for household activities or systems used exclusively for defense and military purposes. The Euai Act classifies risk into four different risk levels, beginning with minimal risk. This is where there are no obligations. Examples of this would be AI based spam filters or AI using computer games. Limited risk applications. Well, that would have. Transparency obligations, meaning the AI systems that interact directly with people and chat bots must be transparent of what the AI is doing and how the AI is impacting processes. Then we have high risk application of AI and this requires conformity assessment. Examples of this would be AI systems used in critical infrastructure management or perhaps in HR management. And then finally we have unacceptable risk. And this means. This type of use of AI is strictly prohibited and. Examples of this again would be social scoring, biometric identification in public spaces, as well as potential behavioral manipulation. Our first requirements for compliance with EUAI Act begin at the minimal risk level, which you may remember has transparency obligations. So let's take a look at these transparency obligations first. Obligations when it comes to protecting against generating Deep fakes, deployers of AI systems that generate or manipulate image, audio or video content constituting a deep fake

shall disclose that the content has been artificially generated or manipulated. Then we have transparency obligations for artistic works. If the AI content is part of an artistic, satirical or fictional work, disclosure is limited to a manner that does not impair the enjoyment of that work. Public information. Anyone who uses AI generated or manipulated content to inform the public about matters of public interest must disclose that the information was AI generated or manipulated. An exception for criminal prosecution cases do not have to be disclosed if the content is used to detect, prevent or investigate criminal offences. Another exception is for review. There's also no obligation to disclose if the AI generated content is subject to human review and responsibility for the content is born. Complying with the EUAI Act can be expensive, but expensive is a relative term relative to what? How expensive is it to fail to comply? Well, companies that do not comply with the rules will be fined and can be fined severely. Let's look at the most severe case non compliance with prohibited AI practices is punishable by fines up to €35 million or up to 7% of annual global turnover, whichever is higher. Noncompliance with other obligations is punishable by fines up to €15 million or up to 3% of annual global turnover. And then finally, even supplying misleading or false information. Is punishable by fines up to seven and a half million euro or up to 1 1/2 percent of annual global turnover. Again, whichever is higher. Well, we understand that protection of privacy rights is a critical grave concern of the EUAI Act. But that's also a primary concern of GDPR. The general Data Protection Regulation. So the EU doesn't intend for the AI Act and GDPR to necessarily be freestanding. They compliment each other. The AI act is a building block of the EUS overall digital ethics strategy. It's intended to be a supplement to the GDPR. The AI Act builds upon the principles of data protection that were previously established by the GDPR. Expands the GDPR to include requirements for AI systems that process personal data. We now have to consider that the AI act extends the protection against automated decisions by AI systems to require transparency, accuracy, and fairness. That we need to incorporate human oversight because the AI Act emphasizes the importance of human oversight and AI system deployment to promote responsible decision making. O again, the AI act complements existing EU regulations. It's part of the EU strategy for digital ethics and the protection of citizens rights. It should not be viewed as an isolated measure, but to be used in conjunction with other EU regulations. Such as the general Data Protection regulation. Let's take a deeper look at how AI overlaps with information and privacy concerns. Because data is usually processed on the provider servers, an AI provider can also be considered a responsible party. Along with a company that actually uses the AI system. And in the case of personal data processing, there's an obligation to obtain a declaration of consent from data subjects similar to GDPR. And if personal data is involved, data subjects must be informed about the scope and intended outcome of using AI similar to GDPR. And finally, before processing data, companies have the application to classify the data and analyze possible consequences and risks. So providing this brief overview of leveraging AI in business, the fifth Industrial Revolution is not so brief after all. Put together all of this content comes to over 2 hours and 16 minutes, but we've covered a lot. Now understand the basics of what artificial intelligence is. We understand where it's been and where it seems to be headed. We now understand the value. And also the risk of integrating AI into an organization. And we've seen examples of this. We delved deeply into AI accountability, safety and risk and extended that



to learn how we might leverage. Standardized frameworks to plan, integrate, manage and improve AI in business, and then finally we learned that responsible use of AI development deployment. And usage is not an. It's not a should, it's a shell. We now have. Regulation. And legal requirements that now go with organizations that use, develop or provide AI. This concludes certified information securities presentation on artificial intelligence in business, the fifth Industrial Revolution. I'm Allen Keele. For joining me.