

ServiceNow AI Control Tower Demo

68 MIN · YOUTUBE · [HTTPS://WWW.YOUTUBE.COM/WATCH?V=SS9EULAM8SO](https://www.youtube.com/watch?v=SS9EULAM8SO)

<https://www.youtube.com/watch?v=ss9eUlaM8So>

SUMMARY

The presentation provides an in-depth overview of ServiceNow's AI Control Tower (AICT), emphasizing its role in AI governance, risk management, and operational efficiency. The AICT serves as a centralized platform for organizations to manage their AI assets, ensuring compliance with regulations and ethical standards while facilitating innovation.

- AICT offers a structured framework for AI governance, helping organizations mitigate risks like bias and privacy breaches.*
- It provides a centralized view of AI systems, enabling proactive identification of risks and compliance with regulations such as the EU AI Act.*
- The platform integrates various governance tools into a single interface, allowing for unified visibility and control across AI initiatives.*
- Key personas involved include the Chief AI Officer (CAIO), AI Center of Excellence (COE), product owners, and risk and compliance teams, all of whom collaborate to manage AI lifecycle and governance.*
- The AI asset inventory tracks models, datasets, and prompts, providing a comprehensive overview of AI systems across the enterprise.*
- Built-in workflows facilitate risk assessments, compliance checks, and case management for AI-related incidents.*
- The platform supports continuous improvement through real-time monitoring and performance metrics, ensuring responsible AI development and deployment.*
- ServiceNow's AICT is designed to evolve with changing regulations and organizational needs, promoting safe and strategic AI adoption.*

Welcome everybody. So we're going to take a look at the ServiceNow AI control tower. It seems to be a very big hot topic. So I thought I'd spend some time putting some pieces together so we can look at a focus sort of foundational overview, what its purpose is, an overview what it's comprised of, review of the sort of AI asset inventory and life cycle, maybe some potential use cases. And then I want to give sort of an insight into the um just looking at some of the tables, some of the main components.

And then we'll do a couple of use cases just to highlight some of the the capability. Very quickly, my name's Stuart Ship. I'm a senior advisory solution consultant for ServiceNow. I'm based in New Zealand, in Auckland. And I'm currently working in the GPC component of ServiceNow, which is our global partners and channels. Quick safe harbor disclaimer. Right, so it's just here to remind everybody that any forward-looking statements we make today are based on the current assumptions and expectations, which of course may change.

And actual results may differ due to risks or uncertainties. And of course, no commitment to delivery on any

future features or timelines. Basically, it's just being transparent and clarifying that today's content is for informational purposes only. If you are interested in this product, please make sure to reach out to your ServiceNow account executive and solution consultant or your partner team or your partner solution consultants. And make sure you have the relevant information that you need when you're making your decisions.

So yeah, quick quick agenda. We'll do an overview of you know, what's the purpose of AI control tower, look at the various components, the life cycle, some of the capabilities, some sample use cases, some of the packaging and the skew licensing. Saying obviously again, most of that will go through your ServiceNow team or your partner team for actual pricing, but it just give you an idea of the skew. And then an insight demo. Great. So without further ado, let's jump into a quick introduction to AI governance in general. So ServiceNow recognizes the importance of AI governance in managing legal, ethical, operational risks all associated with AI. The AI control tower or AICT, if you hear me saying that, that's what I'm referring to, offers a structured framework to ensure AI systems are transparent, fair, accountable and aligned with organizational values. And of course, this is looking to help reduce risks like bias or privacy breaches and security threats, while at the same time enabling safe innovation.

So why are organizations worrying about AI? Well, this slide is actually showing some real-world headlines that show why AI without governance is risky and unpredictable. So from recruitment tools that show gender bias to chatbots providing misinformation, healthcare algorithms that failed vulnerable groups, you know, these are not just technical failures. They're failures in accountability, fairness, and safety. And I guess these headlines sort of reinforce a common theme that without clear oversight, AI systems can cause legal, ethical, and reputational damage.

And if anything, it underscores why governance frameworks like ServiceNow AI control tower are critical for organizations that want to innovate responsibly while protecting customers, employees, and their brand. So there are three key imperatives that are pushing organizations toward AI governance. Typically, it's your risk mitigation. With gen AI adoption accelerating, businesses need a structured risk-based approach to ensure model accuracy and responsible use. And of course, that then helps prevent reputational damage, third-party exposure, or any sort of costly compliance failures.

And then the second one typically is regulatory compliance. So we've now got global frameworks like the EU AI Act, which are pushing enterprises to adopt governance tools to meet standards and respond quickly to audits or regulatory inquiries. And then the third one typically is operational efficiency. So customers want proof that AI is reliable and accountable. So governance frameworks provide that assurance while reducing manual oversight, and that of course is driving streamlined AI operations.

So the AI control tower gives organizations a centralized enterprise-wide view of the AI systems. It acts as the backbone for governance, compliance, and value tracking. So with that risk mitigation, you know, we're helping teams proactively identify AI risks, track issues, and of course avoid those costly legal reputational fallouts, especially if third-party models are involved. With that regulatory compliance, you know, like the EU AI Act,

the control tower enables organizations to quickly respond to audit requests and prove alignment with those compliance standards.

And when it comes to operational efficiency, by showing trusted AI usage across departments, not just in those isolated teams, it provides the transparency needed to measure performance and drive responsible productive gains. So it's not just a dashboard, right? It's it's a governance nerve center for enterprise AI. And then from a ServiceNow perspective, you know, we have a long history of helping customers manage risk from the launch of the our first risk product offering back in 2016 to today's AI control tower. So this solution is basically like a natural evolution of our of our platform strategy. We you know, we're already supporting governance in areas like security, privacy, asset and project management with our tool sets and and different applications. And now just seamlessly, we're sort of extending that to AI, which is obviously that rapidly growing area of concern for our for our customers. But what makes our solution unique is that it consolidates governance tools for models, data, ITM, etc. and more into that single integrated AI platform.

And of course, that enables unified visibility, control, and compliance. And because it's built on the Now platform, customers can connect it seamlessly with those other applications that make sense to have it, right? So if you've got SPM, you've got security, you've got risk, all of those pieces come into it and help you proactively manage your AI threats and regulatory complexity at scale. So you know, ultimately, customers are choosing ServiceNow because they know they have the depth of expertise and platform breadth to help them stay ahead of the AI risk curve.

Okay, so let's let's have a look at what it's what it's all about. So ServiceNow's AI governance approach is designed for the for the real world, right? Where organizations rely on a mix of in-house, embedded, and consumer AI technologies. So we're providing that unified oversight across typically three key AI sources. Firstly, your custom-built AI. So that could be using frameworks like Azure or AWS, OpenAI, and others.

Then you've got your your third-party apps. So you know, it could be embedded gen AI. It could be SAP, Workday, or Salesforce. And then you have consumer tools like Gemini or Cloud that users can access independently. As well of course as as the ServiceNow pieces. So this multi-vendor governance ensures AI um is then deployed ethically, transparently, and securely regardless of where it's built or embedded. And we at ServiceNow are also collaborating with you know, global experts and regulators to ensure that the governance practice align with evolving policies and enterprise risk expectations.

So let's have a look at what we have to do. So to to manage AI risk at scale, organizations must break down silos between business, technology, and risk functions. So the chief AI officer, the CAIO, if you hear me say that, the chief AI officer, is critical to coordinating the competing AI interests across the organization. And the AI COE, center of excellence, team works with the CAIO to shape the strategy, performance, governance, and reputation of AI. And of course, depending on the maturity of the organization, there may not even be at this point a CAIO or an AI COE. But essentially, there'll be some sort of ad hoc cross-functional team that are making those decisions.

So on the left, and this is a bit we're sort of trying to join up using those personas, is the product teams. You know, they're focused and sorry, these are more like the the the value driving teams, right? So you got the product teams focused on innovation, business users, you know, pursuing efficiency, GTM go-to-market teams driving revenue and CX. But without alignment with the right-hand side functions, AI risk can escalate. So on the right, you know, we've got risk and compliance teams. So we need to coordinate with them to ensure AI meets internal policies and external regulations. We've got our technology, privacy, and audit teams, which are responsible for securing AI assets, governing data, and providing oversight.

And then, you know, personas like our enterprise architects making sure that AI aligns with broader architecture and strategy. So ServiceNow AI control tower helps unify these perspectives, providing a shared system of record to govern AI responsibility across the across the enterprise. So ServiceNow has been working behind the scenes to give organizations the tools they need to oversee AI AI governance throughout the enterprise. So in Feb, we released our AI inventory as part of our CMDB. In May, we delivered a GA product called AI Control Tower AICT with the ability to manage AI governance at scale across the enterprise. And we have many more hearts and to be delivered in quarterly releases.

Obviously, that includes the the second half of 2025 again safe harbor around these dates, but it gives you a good example of how we're doing that. And we're also working with our partners to build other capabilities that improve AI discovery, evaluation, and monitoring assets beyond 2025. So what am I trying to say here? As AI evolves, you know, as with all of our ServiceNow products, you know, expect our AICT product to evolve and its capabilities to evolve as well.

So the AI Control Tower begins with an in-platform AI inventory, basically a CMDB for AI, that brings together models, applications, and data from across platforms. And it serves two core groups. So your AI leaders like the CAIO or the COE, who need a unified workspace to manage strategy, life cycle, performance, and governance of AI use cases. You know, how to use AI capabilities, do they need to be consumed or implemented?

What is their performance? What is their value? And the second part of managing AI is collaborating with other functions. So our risk and compliance stakeholders, who must assess and manage AI-related risks. You know, and this is spanning security, privacy, audit, third-party oversight. And then we can move into what we kind of have inside of the AI Control package. So within it, we've got the AI Control Tower workspace.

We've got the AI Discovery and Inventory sort of data models that are coming into that. We have our AI Asset Life Cycle Management and Risk Management components. We have our AI Case Management components. And there's also built-in content that is related to organizations and major frameworks like NIST AI RMF and the EU AI Act. So together, these products create a complete system of record and action for enterprise AI governance.

And again, sort of reach out to your relevant partner, ServiceNow account executive, ServiceNow partner SEs, et cetera, to get the information that you need to help you make the right um the right details. And then the bit

that I left at the top, and this is on the ServiceNow AI Control Tower page that you can go and browse for yourselves. But it's important to remember that AI governance doesn't necessarily, right, exist in isolation. And it tends to probably work best when integrated with adjacent capabilities across the enterprise. So what am I talking about? As I will think about, yes, you can have AICT, but you may want to make sure that you've actually integrated this with your integrated risk management, used to be called GRC, obviously now IRM. And maybe your IT asset management components. And then that strategic portfolio management piece could also come come into that.

So just stepping through a couple of these capability, and I'll come to the skew in a second. But yeah, as as with all of our workspaces, we're looking to centralize an AI strategy and operations and oversight, and that's the key part of the AI Control Tower workspace. Connecting risk, compliance, business, and technical teams. So not to listen to any of our workspaces is providing a unified dashboard in this case to track AI assets, open cases, compliance gaps, performance metrics. And it's leveraging that AI inventory CMDB-based components to which you which you add in from AICT to manage models, data sets, prompts, and their life cycle. So from onboarding to decommissioning.

And it embeds risk and compliance workflows including sort of native AI impact assessments to identify vulnerabilities and and and apply appropriate controls. And of course, it supports that real-time monitoring. And in this case for emerging AI risk use and trends and ROI measurement. We scared a couple of times, so I just highlighted it here that this AICT comes with an AI Discovery and Inventory data model.

You can go on to the ServiceNow store and have a look at these. But basically, it brings in all the classes to support those AI systems, those AI models, and the connections that go go with it. So if you're in if you know ServiceNow quite well, remember to go into your `sys_db_object.list`. Have a look at all the tables and columns, and I'll and I'll show a bit of this during the demo so you get a feel for it to understand, you know, how this has been expanded natively using the capability in the tools. I'm just adding particular relevant sort of classes and updates that can be done for your AI assets and AI asset life cycle.

And you can also have the AI Risk and Compliance Management, again something that's available from the store. And this provides a centralized life cycle governance of those AI systems, models, and data sets. So enabling that ethical AI use case, risk mitigation, and regulatory compliance. And this has a life cycle management. We've got impact and risk assessments. You've got an AI case management so you can log, investigate, and resolve AI-related instance. Of course, it's then leveraging those compliance frameworks. So the EU AI Act and NIST AI AI RMF content packs.

And then you've got the 360 view, internal collaboration, et cetera, that you can use that component for. And impact assessments and again, I'll show you this in the the demo. They're a cool part of responsible AI governance. And of course, as you would imagine, ServiceNow is making them easy to configure and operationalize. So you can go in there and assessment covers dimensions like, you know, non-discrimination, data protection, transparency, human oversight, and other other ethical risks.

And then once completed, the system can then automatically generate tasks aligned to the AI life cycle phase, whether it's designed, development, deployment, we'll come to that in a second. Helping teams mitigate the identified risks early. So again, this isn't this isn't showing that governance isn't just a check box. It's embedding directly into the workflows and making accountability part of how AI is developed and deployed. And just to layer that a little bit more, so to govern AI effectively, you know, you've got to make sure that we've got our case management or our issue component that comes with that. So in case management, teams can log and manage a range of AI-related events, inquiries, incidents, risk alerts, compliance escalations.

And each case tracks its status, breach potential, priority level, root cause, helping sort of triage and prioritize our responses. And then with issue management, when risks or non-compliance are confirmed, this component of the tool lets you track remediation. So you can assign ownership, monitor progress, and ensure accountability across the across the teams. And within this, you have a few drill-downs. So here's an example of drilling into an AI case.

So it's allowing you to do what you would expect, log and triage an AI-related case such as system flaws, fairness concerns, compliance breaches. You can see and map the full impact context. So affected users, regulations that are triggered, related issues and consequences. And assign remediation tasks, linking them to AI life cycle phases to ensure timely resolution. So again, helping you meet those regulatory obligations because you're able to maintain a full audit trail of case status, milestones, actions taken, which of course is critical for compliance reporting and internal oversight.

Mentioned a couple of times, there's these content packs. I've put a screenshot in there. And again, you can go into the ServiceNow store and have a look at these and bring them down. And these currently are for the AI frameworks, the EU AI Act and the NIST AI RMF. EU is the European Union. NIST is the US National Institute of Standards and Technology, I hope. So they're just two very large accountable frameworks that are developed to promote trustworthy AI. And in these frameworks, you know, you'll get agencies, authority documents, citations, control objectives, risk statements, AI assessment templates, and AI risk assessment methodologies.

So just a quick recap, um, the core value of the AI Control Tower is that it's delivering across the enterprise, right? We're spanning governance, visibility, risk, and business alignment. So, we've got our centralized governance that's unifying that control of all AI initiatives, you know, whether they're built in-house, third party, embedded in SaaS tools, or powered by AI agents. We have our visibility and control, um, so this allows leaders the oversight needed to scale AI safely and confidently, breaking down those silos across the teams.

We've got, um, AI alignment to business goals, so it's linking AI projects to strategic outcomes, helping prioritize the efforts that deliver value, while again accounting for any sort of regulatory, uh, impact. Our AI lifecycles, so supporting AI from onboarding to retirement with end-to-end workflows that ensure responsible development, uh, and collaboration across roles. Proactive risk uh, mitigation, embedding risk and compliance into the AI lifecycle, reducing reputation operational exposure, um, while enabling faster, safer, uh, innovation. Uh, and then platform as a foundation to scale, so because we're built natively on ServiceNow, it integrates seamlessly

with IT, risk, security, and compliance workflows, uh, and that's creating that unified control fabric across the business. Um, so yeah, not just about oversight, it's about enabling safe, strategic, uh, and scalable AI adoption.

So, let's have a little look into, um, the AI asset inventory and lifecycle, which is kind of at the, uh, the core of that. So, the foundation of, um, the AI Control Tower is that AI inventory, which connects all aspects into that one single data model running on the platform. Um, and of course it runs on existing CMDB if there's one there, or in some cases CSDM, um, and there there five types of, um, AI assets in the implementation of, uh, of any AI system, or that's the way that, um, ServiceNow has set that up and adopted it.

So, what we refer to as AI systems are primarily the use cases that are being built for specific outcomes, uh, and their users. Um, and these can include classic AI, um, generative AI, agentic AI, but they're all referred to as AI systems. Um, next we need AI models, which are the foundation for the AI system, so they are the programs that are trained to derive the needed content without human intervention.

And multiple datasets can be used to train a model, so that the AI system can then perform the right level of activities and make specific decisions in the targeted area. Then come prompts, and these are specific instructions to the models and the AI systems to generate the right level of output from the AI systems. And then in addition, other inputs and outputs may be present that represent the type of, um, expected configuration for the AI systems.

So, let's take a a quick example of how this works. So, here's a just a a use case, right? It's an AI-based credit scoring for loan approvals, um, which provides scoring based on loan approvals or loan applications. So, the, um, the AI system itself consists of, um, user interfaces, um, monitoring tools, different models, decision rules, APIs, etc. And it uses different models like the Now LLM, OpenAI, um, and Azure, but the system is solving one particular use case, right? So, that's the credit scoring for, um, loan approvals.

So, all of these models are trained with specific datasets. So, for this use case, models need to be trained with financial information, um, public records, and financial behaviors, so that the AI can make the right level of decisions and perform the activity that it's intended for. So, this is kind of an example of how, uh, an AI system model, uh, and datasets come together in a simple GenAI, um, use case. So, we have, um, three primary key personas that play, um, vital roles in managing AI throughout the enterprise, um, using AI CT, AI Control Tower.

So, we have the CAIO, the Chief AI Officer, um, who is responsible for the holistic AI governance process within an organization, along with the, uh, AI COE, Center of Excellence. So, the AI COE team is dedicated to driving and supporting the adoption, implementation, and best practices of AI across an organization. And the AI Control Tower will provide a workspace for these stakeholders to understand the AI use case, uh, use case, inventory, um, apply a lifecycle approach, uh, track adoption, and measure AI value. Uh, and AI case management can handle basic questions related to AI or related issues such as AI violations and AI risk and compliance requirements.

And then our second persona, the the product owners, so they build AI capabilities into their operations to help address the business, uh, and operational needs of customers and employees. So, with the AI Control Tower, they can see, um, the entire flow of the AI strategy, demand, and roadmap, and potential portfolio projects to make sure they track all of the capabilities in, uh, an organized way. And the the third persona, uh, a critical group, is the risk and compliance team.

So, they're responsible for ensuring that AI capabilities are being built responsibly and ethically, uh, and the risk and compliance teams collaborate on how to, um, define the right policies, uh, risk assessments, third-party assessments, um, understanding the privacy and regulatory aspects and associated frameworks, understanding, uh, deployment strategies, and manage the content and workflows also in a single place with the AI CT. So, the AI Control lifecycle, um, applies to product owners, Chief AI Officers, and risk and compliance, security, legal, and architecture teams, etc.

And product owners must follow this lifecycle when building products. The first phase of the lifecycle is the demand. Product owners or anyone in the organization can submit an intake for an AI use case or issue from their workspace or employee center. Once it's submitted, um, the AI steward, the product owner, and the risk and compliance team will start their assessment. And this means talking with the product owner and risk and compliance teams. And then we have legal, security, third-party, and data governance teams who may also be involved, uh, and then following the assessment approval, the product owner can build the entire capability by following the guidelines and getting approvals.

And this stage includes, uh, maybe a pre-deployment review to ensure compliance with the guardrails in place. And then the final phase is deploying the AI capability, monitoring it, and then analyzing the value AI is delivering. And that, of course, will then foster, um, continuous improvement. So, this slide looks to illustrate how, um, AI CT, AI Control Tower supports the entire AI governance lifecycle from demands through value recognition.

And the boxes break down the specific capabilities AI CT provides during the, uh, AI governance lifecycle. So, boxes with, um, the blue border are essentially related to AI lifecycle governance. Boxes that have a blue border and are shaded blue show, uh, a multiple task view, so they're performed at the platform level to support shared tasks across multiple functions with, with, um, within an organizations. So, boxes with a yellow border are part of the risk and compliance use case, uh, and reflect capabilities related to, um, the risk and compliance team. Uh, and those boxes with AI CT designations in green, um, signify the features that are contained in the AI CT SKU.

Uh, those with obviously SPM, uh, are not part of the AI CT SKU, um, but those features are available to, um, customers who have SPM in their instance. Uh, so for example, intake tasks are shared by IRM and SPM at the platform level, while AI and risk assessments and risk and control identification, issue management, um, continuous monitoring are handled by the, uh, the IRM workflows. Um, Oh, I just quickly, uh, what was it?

FRIA is in here somewhere. FRIA, so FRIA is the fundamental rights impact assessment, which is essentially, uh, a structured evaluation of how, um, an AI system model or dataset could affect, uh, individuals' fundamental

rights. And I mentioned a couple times, so we're shipping two main content packs, uh, one for the EU AI Act, and the second to support the NIST AI Risk Management Framework. Uh, both of those include relevant citations, which are mapped to, um, control objectives.

Um, and then you can obviously see, uh, associated, um, products within that, descriptions, authority documents, and so on. And then, of course, we'll look to update these content packs, uh, as we go forward with each release, uh, and you don't require a separate subscription for these. So, while not mandatory, AI Control Tower aligns with CSDM V5 to offer a structured lifecycle view of AI systems, models, datasets, and prompts across ideation to consumption. So, organizations that are already using the CSDM or CMDBs can easily integrate AI inventory into their system of action, linking AI systems to services, business assets, and digital products.

And this model supports both locally hosted AI applications and cloud-delivered AI functions. And these are all tracked and consumed as part of the service instances. So, giving business applications a reliable way to consume those AI capabilities. So, for example, a summarization AI use case might use GPT-4o and its datasets to generate insights, which can then be deployed on infrastructure and tracked by CMDB relationships in the in the Now platform.

So, yeah, we've got AI Control Tower as a whole providing that unified workspace to manage strategy, risk, and operations for AI across the enterprise with built-in governance aligned to evolving regulations. And it integrates AI assets into CMDB or CSDM, enabling full visibility into systems, models, datasets, and prompts. We've got lifecycle features that support assessment playbooks, upgrade coordination, and alignment to defined AI capability lifecycles, risk and compliance capabilities, which will help assess internal third-party AI risks, complying with those frameworks, the EU AI Act and the NIST AI RMF, monitor for early signs of non-compliance.

And of course, we've got our case management, which is centralizing our AI-related incidents and policy enforcement. While value and engagement metrics help track adoption, usage, and security posture for generative AI workflows. So, let's look at some of these typical use cases just as an example, right? So, I think before I just jump into this and this is a very simple one, but the AI Control Tower is not an industry- or geo-specific solution because AI is agnostic, so is AI CT, right? It's ideal for any organization adopting AI in its daily operations and getting the most value from its from its deployment.

So, just some a few use cases just to be put together, right? To help understand maybe a couple of business applications that may use AI-driven governance. So, the first one is a personal finance advisor. So, it's a tool that maybe is using AI to analyze user spending habits and providing personalized budgeting and saving recommendations. And that may have, you know, maybe a risk score that's that that is fairly minimal. But again, you still need to track and and run that through. You might also have a bot for L1 customer support, right? So, a chatbot that uses natural language processing NLP to interact with customers, answer queries, and deliver support for basic financial services.

And maybe that's got something that's along kind of a limited risk. Maybe there's some transparency in data privacy and security that you've got to think about. Another one could be a credit scoring for loan approvals. So, an AI-powered credit scoring system that evaluates individuals' creditworthiness for loan approvals, right? So, now we're moving into something that actually got a little bit more of a higher risk. So, is there potential for discrimination? We've got to think about the accuracy and transparency as we move in something like that.

And then something on the far end side, right? Financial eligibility scoring system. So, maybe we've got an AI system that scores individuals based on their financial behaviors, their social media activities, and personal data to determine their eligibility for financial products. So, obviously this might have some critical risk. So, we may have violations of fundamental rights, excuse me, a lack of transparency, there could be social and ethical impacts we need to do. So, just some of those use cases, but in all these cases that we need to track those. We need to understand those. We need to understand what they comprise of and make sure that we're governing them and able to see how we're governing them and what we're running against them.

So, I will quickly mention I just want to talk about the patching availability. Again, for this section, I'm not putting any pricing in here or anything like that. I just wanted to give people an understanding that there is a separate SKU for AI CT. Always talk to your ServiceNow representatives or your partner representatives, who in turn will speak to ServiceNow representative, to make sure that you've got the right information when it comes to licensing and how that works. But I did want to kind of just show that there is an AI Control Tower SKU.

You do get some of this SKU if you have a Pro Plus or an Enterprise Plus SKU. But obviously, just looking at this, you can see all the extra needed elements that you get from AI CT with the the SKU that comes onto it. And just a quick one around it's actually a simple a simple meter. So, you basically when you're putting this in, it's just a total count of your AI assets that work towards that that SKU. But again, make sure you're talking to the relevant people to get the relevant information to base off any decisions that that you're making.

So, yeah, thanks for this. I think next we're going to jump into just an Insight demo. Please note in the demo, you may see things where there's usage integration to things like SPM and IRM, which might have separate licenses, etc. But in general, we're going to stick to the AI CT. I'll go through the demo showing just some standard kind of if you're looking at it from an implementation side, what are the components it's made out of. And then I'll do two kind of walk-through scenarios just kind of cement that feel that goes into it.

Great. So, yeah, people who know me, I come from technology side of the business before I kind of moved into various roles. So, I like to look at sort of components that helps me visualize. So, the first part of this demo, we're just going to go through and look at some of the components in detail, talk to those, and then we'll do some use case stories to take a look at that. So, first of all, I always like to go into the system and understand any data model, if that's what I'm looking at, get an idea of what that data model is, how it's built, what the forms are. That helps me visualize and understand any components, particularly in ServiceNow as it grows quite a lot and new applications come out a lot.

So, no different for for this one. So, typically what I'll do when I come into a new system, I'll go and have a look at the the database or the tables, look at those operational pieces, and how they how they all fit together. So, don't forget that you can do your sys_db. sys_db. Do impersonate this person, which is why I'm getting the wrong thing. There we go. Just bear with me.

There we go. Technical There we go. Bear with me. So, yeah, I like to come in and have a look at these tables. This should now work. So, I'm going to go to my sys_db. And this little object list. Oh, yes. And obviously, anything you see inside of the system, this is basically just a demo system with demo data, demo pieces, and obviously not related to any individual or company if that is perceived that way. So, it's entirely entirely pre-demo stuff.

So, yeah, maybe I just want to look at the the AI applications. So, let's just do AI Control Tower. Come in here. And then yeah, we're going to look at some of these tables. So, I'll pick which one should I pick? Probably the AI Asset Governance Details cuz that's quite a cool one. So, yeah, you can go into each of these tables, take a look at, you know, what are the columns. I'll just put it into You could see, you know, what are the what are the related rules or client scripts UI policies. And I just And this is really good way of kind of getting to understand any components and capabilities that goes into that. And obviously, like I said, tables to go on through it. And something else as well, you can jump into the schema map. So, you can also do this by going to the to the nav all, excuse system definition tables and columns.

And then opening schema map from here, but I find this easier, right? Cuz I'm looking at tables, and then I can jump into the schema map, and look at just some of those, you know, reference pieces. And this allows you to then kind of drill out of, you know, what are all the columns in there. So, it's just a very useful mapping exercise that I find is a is a good way to kind of start when you're looking at a new product and the capabilities that surround it and how it kind of relates and and sits together. So, let's just do one more of those just so you get an idea what this look at the the AI asset life cycle. So again, gives you sort of the columns and any kind of relationships that sit into that, and then I can obviously build that out and then look at kind of the reference the reference tables. So yeah, if you if you don't do things like this, you know, fair enough, you don't have to, but I find if I'm kind of moving into a new capability, a new component, this is something that's a useful exercise to go to go through.

Anyway, right. So let's go back into the system. So that gives you a good sort of introduction to, you know, what are the new tables, what are the new capabilities that sit underneath that part. Just going to impersonate a user. Mainly we'll use this user in the the run-throughs, but they have better demo data than my admin does. And we'll go and look at some of these some of these components. So yeah, the workspace and you can get to that from here. We're just going into the AI control tower. And we'll just talk through some of these. So essentially this is basically the central interface for the AICT.

And offers that unified view of AI governance across the enterprise and it's looking to integrate, you know, intake and risk, compliance and monitoring workflows in in a single location. And we can go through those. And we do have things like top items to review. So it's kind of obviously a dynamic task summary for immediate

actions. So we can surface key AI governance tasks that require attention, such as you know, AI systems that are pending review or assessments.

We can click into those. And it's essentially just making it easy for governance as you as you go through that. And then the state of AI the overview. So here we have a high-level KPIs showing how many AI systems are registered in the in the organization, right? And their current sort of life cycle stages. And and how many are in ideation, how many in build and test, how many deployed. And then, you know, we've got an AI systems by type by type as an example.

So classification of AI technologies that are in use. It could be ML models or LLMs. And this is just helping teams understand where AI is being applied, whether it's generative AI, right? It's machine learning or genetic or rule-based that goes into that. We've got our risk classification, so distribution of AI risks where that sits. So, you know, we classify AI systems using those risk frameworks like NIST and and the EU one. And this gauge can help indicate how many systems are high, medium or low risk based on criteria.

And we can drill into all of these as well. So then again, that's kind of essential for assigning the right controls and oversight. And then as we come down, AI systems by provider. So, you know, which vendors' tools are being used, right? Is it OpenAI? Is it ServiceNow? Is it Anthropic? So then we can understand, you know, this gives us vendor-specific governance, risk profiling, integration visibility across the enterprise.

So the AI systems trend. So how many AI systems are being onboarded over time, right? Or deployed. And this will allow us to track adoption velocity, informs of resource planning, monitor the rate of change in AI usage. And then just down here we got a compliance, so the audit or assessment status of systems using that EU AI and NIST AI frameworks that go into that.

And then of course if we have our case management, AI cases by priority. So active issues or risks sorted by severity that we can we can build out onto. And then we can jump into AI asset inventory. So essentially this is the total number of AI systems across the enterprise. So, you know, in this example we're currently tracking over and this is all just dummy data, 2,300 AI systems. And that includes everything from like I said LLMs to ML models. They're all in production. So this is basically giving us a real-time asset register.

A foundational sort of governance requirement that many organizations obviously don't have today and and they need that goes into that. And then we have our status breakdowns, AI systems by provider, AI systems by type. So again, you know, splitting those up to make them digestible. And then down the bottom, you know, here's an example of departmental distribution of AI systems, the AI asset inventory by department. So this this visual as an example is incredibly useful for understanding where AI is being developed or used across the enterprise by function, right? So which teams are using it, where they're using it, how are they using it.

And it can be useful for highlighting where systems may be stuck. So there may be ones that you can see for a particular team that are not progressing past the approval stage, right? So then that enables targeted engagement with particular departments to unblock that progress or apply sort of governance controls that goes

into it. And then a risk and compliance component. So this is showing risk-based distribution of the AI systems. And the risk classification chart can show how AI assets again are using frameworks like NIST or internal policy standard. And each system is assigned a risk level based on factors like data sensitivity, automation, sorry, use of automation, impact and decision-making that goes into that.

And then we have things like AI systems by risk classification, AI asset pieces. We can have compliance by authority documents and policies if they exist. And this view sort of benchmarks compliance against those frameworks, against those policies. So it gives you a policy level as an example, a policy level lens on AI governance. Not just where AI is being used, but whether it's actually complying with your documented obligations.

Okay, so let's go into the next tab in the AI control tower. So this is the sort of section where if you're not hopefully used to ServiceNow, you get these list views of the records you want. But we've got our AI asset inventory here, which is the foundational level of AI governance. So this is where all the known AI systems and models and prompts that I was talking about data sets can be tracked, reviewed, see which life cycle phase they're in, see what's deployed, what isn't deployed, all the relevant pieces that can be added into that.

And then we have approvals. So during that intake or risk and compliance components, there may be approvals that's all associated here. We have our life cycle components. So this section is that end-to-end journey of each AI system. So from ideation to decommissioning, providing oversight at each of the each of the stages. And those of course are split out of the box between ideate, build and test, deploy, retire, etc.

And then of course our cases. So that that section that provides the central place to sort of track incidents issues. And we'll there's another section a second, we'll come back to this. And that goes on to that. And then inquiries that that come from that side. There's a small configuration panel with just some easy click capabilities showing you which which templates you can use, some controls within it. You can activate these or deactivate them. So again, a lot of the latest applications that we're building in ServiceNow are building basic sort of configuration capability.

Health and safety comes to mind, they have a very similar thing inside of the the workspace. Again, trying to make those different components easier to handle. And then we have our AI case management. So here we're visualizing, you know, case progress. What's open, what's in review, helping identify backlogs and response velocity. So pretty kind of standard sort of use case that goes into that. But you can run your case management from this.

And of course there's inquiries, right? If we need to track those and their life cycle, there's also an out of box piece. And again, thinking about the the business units and and how that looks. So yeah, just a brief kind of overview of the different components and how that fits. And then there's some other pieces. So which I go now, if I jump into the portal, we can just look at some of the components that are up there. And you may have, you know, ESC Pro or whatever you wherever you're building, whatever your central portal is. As with all of these capabilities, you can put that in there. Remembering that AICT sits within the risk and compliance section. So

we just jump into risk and compliance.

And out of box you're you're getting the ability to report an AI use case, right? To request an AI use case. Request a data set, a request an AI model model. And these are all pre-formed forms. Of course, you can change them depending on your kind of needs and capability, but as with everything in ServiceNow, we try and give you an out-of-the-box capability that allows you to pretty much run out of the box and then just customize as you need. So, here's an example of request an AI use case. And we'll do this in a couple of the walkthroughs as we go through through the tool.

So, just taking that a little bit further, we can go and look at some of those uh Sorry, about this. Always does this. Just some of those components from a uh tool set. So, I've come back in as my admin user. Maybe I want to see um some of those record producers. I'm just trying to give you a visual right for all of the different uh components that um that get built out.

Actually, let's see that as a favorite. I did. Um give you an idea of all the different components. And again, this will make more sense to people that are quite well versed in ServiceNow. The the demos that I show later with the the storyline will make more sense to those people as also the people that don't necessarily sit in ServiceNow day-to-day. Um but I'm just going to the record producers um and let's find that uh uh Let's uh find the AI use case one as an example.

Cool. This going to be out of lock. Uh Let me try if I can put that in there. There we go. So, that was the one we were just looking at request an AI use case. So, these again, they'll be those that come out of the box. So, I'll just change the scope so that we can see this more. And it just gives you an example of you know, which table that would sit in any related sort of scripts.

And the reason I show you this as well is, right, if you if you need to build out this capability to something bespoke, you can always start with these components. Best practice is always clone. Don't change the out-of-the-box. But again, just gives you an idea of those components that they're built into into the portal. Um what's another good one we can look at? So, actually yeah, the the playbook component. Um So, we could go through and go into the AI control tower again. Just show you an example before I jump into that component.

And again, we'll go through this in the in the demo. components recent It's got a playbook associated, which you see it on the left when this record finally comes up. And that configuration component, of course, I picked one where it doesn't come up. Let me try another one. Try this one.

There we go. That's what I was looking for. The life cycle. So, this is pre-built AI asset life cycle playbook that you can use. And again, if you need to expand that, clone and then update as you need. But let's just take a quick look at uh Look at that. And then we've done some good things to the end.

Finally, we should bring up to the workflow studio. Yeah, there we go. Um we're just going to have a look at that

from from the back end. And again, I just want to show you all the various components. For the workspaces, that'll be UI builder, right? If you're trying to change those components. But hopefully, we can just do a quick search here for asset life cycle and this will be the one that we're using during our demonstration as we as we go into that.

As we go down that, I won't go into all of the stages, but just showing you the the playbook component that that also comes comes through that. Um so, let's quickly impersonate. What else can I show you? Maybe just a little bit more onto the case management piece. And I'm going to impersonate uh and I'm going to say that I have better data.

So, let's go to case management. And we can look at different cases. I'll just pick on these overdue cases and we'll pick one just to give you a view of kind of how that looks. So, cases have their own sort of life cycle. They have approvals, a timeline, a state that goes into that, description, action tasks. You can get the related sort of regulations. And this is obviously tied to the EU one. And it causes increments. So again, that ServiceNow capability of tying all the relevant components together and some nice features to do a 360 view.

So, this can be particularly important. Obviously, not going to load. This going to be particularly important if you have a lot of relationships, particularly when you move down to authority documents, maybe citations, controls, control objectives. This 360 view can be really really useful to get a good kind of overarching understanding of how that's how that works. Cool. So, that's that part of the demo done. I just wanted to show some of those components, capability pieces, how some of those bits tie together.

And now we're sort of do a story walkthrough scenario using some personas that will hopefully bring all of those pieces together, the the slide deck plus the quick demo I've shown you on on those pieces into kind of a real a real use case. So, we're going to log in as our user Mary Cruz. And Mary Cruz is essentially the AI product owner. So, let me just quickly do that.

Sorry, there is a bit of a lag that goes into it. And just while it's doing that, so Mary Cruz is going to come in and what she wants to do is propose a a new AI system. We're going to call that the AI benefits recommender. And this is essentially to help guide employees through personalized benefit selections. And that could be based on their job role, location, family status. Helps support the internal HR team in that space. So, starting point for that for Mary is going to be those record producers or the the requests, whatever whichever way you want to call them.

They're pre-predefined forms that they can fill in requesting that AI use case. And as with all the portal capability, I didn't jump through all those. You can create favorites, right? You go to the risk and compliance piece as well. But I'm just going to click on it from here for something slightly slightly different. So, I have made this sort of prefill based off of the name, hopefully, just to kind of speed up this demo section. You don't want to watch me type lots and lots of times. So, I'm going to give this a 9.0 cuz I've done a few I've done a few of these with this.

But you can start to see it, right? So, what is the state? Is it draft, deployed? Is there a model category? And you can build these out as a classic or a generative. A bit of a description. You can relate documentation, who's going to be the owner for that, who's going to be managed by, what particular team. And then you can build out your AI models, right? And your data set. So, in this example, we're going to use GPT 4.0.

And then there's a particular data set, the ServiceNow organization products and services that I want to use from that. So, just an example of how that's how that may transaction. So, that's going to be submitted by Mary. And then we can look into that. So, now that's come in, we are going to hit into another persona. I need the right clicks again. Sorry, too many tabs. It's where I need to be.

And then I'm going to now um impersonate a new persona who is the AI steward manager. Um Now, I probably saying this name wrong. I apologize if I do. Aleni? Elini? Whichever one you want. Aleni Reybek. And I'm going to impersonate this this user. And we're going to say it's Aleni. She spends her time in the AI control tower workspace. So, now we've run through these components. Hopefully, you've got a kind of a better understanding of that. So, what she's going to do is jump in to see, you know, what are the new systems that have been added or suggestions that have come in.

And I can see the record that Maria has just has just submitted. So, let's open that up. And because I'm in the right persona, so now I want to review this, right? I want to push this through the life cycle. We want to perform all the different tasks. We want to see if we can do it, our testing, any approvals that go into that. So, this is kind of just a simple overview of how that life cycle work, right? Maybe they reject it straight away. And of course I will be skipping through this very quickly. You would imagine or please have some artistic license that there's a lot more so going under the hood and lot lot of people are associated to that. But let's just start the the review.

So we move into that and the playbook that I showed you as well can come into effect. And as we mentioned there can be various tasks. Maybe you add more. Could be architecture, could be risk assessment or stakeholder review, impact assessment. But the key thing is we've got different tasks that are occurring inside of that life cycle as we as we go through it. So that means there's just there could be an impact assessment, right? So this could be related to a particular sort of questionnaire.

You could have a stakeholder review. So maybe an approval of you know, do we want to do this? Does it have the right value? We obviously need to sign something to our risk assessment team. You know, what is the risk of us doing this? You know, what do we have to take care of and make sure all that's related. There could be architectural review to our architects to go into it. And again, I've just simplified this. I'm not going to walk through each of these tasks.

I've just set up a script that will just set these all to complete. But again, hopefully you're getting the idea, right? Of in assess there's there's certain activities that we want to occur. And they can be assigned to different users, different capabilities. And then all of that collaboration, connectivity, all in one platform piece comes together. So then once we've done our assessment, we move into the build and test. And then we can move into our development plan. So do we want to create a controller test station, right?

Do we need to collect information on the deployment regions they are system. Again, there could be lots more tasks that goes on to this. And these can be associated to the relevant people. And again, I'm just set these up so they can move to completed hopefully within within a short amount of time. These again will be set to and we'll repeat that once completed. Relevant people and all the capabilities that sits underneath that. And once we've done our development plan, we can move into the pre-deployment assessments.

And again, it has its activity and how that bit works. So now we may do things like review issues or any policy exceptions, right? That comes out of that of that piece. What's our conformity assessment and how does how does that run? Just wait a few seconds for these to set to completed as well. Perfect. So just again, trying to show you that entry point, the intake. We're going through that asset life cycle and then we're moving into our deployment. But all as we're going through this, all of that governance capability is all being drawn in, right?

What do we need? Where do we need it? How is it controlled? What's our activity? What are our risk and compliance measures? How do we have related assets? Moving forward, how do we have all that connectivity and data that sits into that as it goes forward. So yeah, quick spin through, but hopefully that gives you sort of idea of kind of the intake and how that works through. So the next scenario I'll show will just kind of highlight more that assessment component that may that may come through. So I'll log back in Mary Cruz again.

This is going to be a little bit of a repetition, but I'm going to submit another request so we can use a record. But yeah, come back in as Mary Cruz who is our AI product owner. I'm hoping the demo gods are with me. There we go. So and we're going to go back to the portal. And again, I'm just using the AC portal as my computer.

I'm going to go and request another AI use case that we can use for this kind of assessment one. So this one's a little bit more kind of laborious. I'll say this is Well, I've done a few of these. We want development. Version doesn't really matter. Let's say generative AI. And we'll give this one a description where we're basically thinking about an AI system that analyzes customer transaction patterns. And the idea being to maybe uncover buying habits, season trends, product affinities. And then that can give us insight for targeted marketing and so on. So again, it's a simplified sort of use case. As I mentioned, you can add documentation.

So maybe you've done some research, you've got a proper planning around this. Who's going to be the provider? We're going to use ServiceNow for this one. I'll give this to Mary. If that's going to be who's going to manage it. And then this one's for IT. So it's not for it's not for HR this time. And here's the example, right? You can create all these different AI models. You could choose different ones depending on what you need to have in there and it could have various data sets and so on and so forth that you could put in. But again, so that was just kind of kind of more of a walk through.

We'll create that create that use case. So now we're going to go back in. And this time two minutes down to open. We will impersonate I'm going to say her name differently now, I'm sure. Alina Lenz. Let's impersonate Alini. Who again is our AI steward that comes into it. So at this point again, we're going to come into our new AI systems added.

There's notifications, right? Or understanding. Maybe they're looking at the criticality, how that comes in. And I can see the latest piece that's been added in is the AI customer transaction. So now I'm going to go into this and I can see that I need to start the start the review. And this will hopefully kick off so I can show the assessment which will I hope be set out to to Mary Cruz. So I'll just check that that's worked.

And again, hoping the demo gods are with me. Yes, excellent. There's one missing. We'll we'll look at the impact assessment. So let's jump back in Mary Cruz. We're going to come back into our AI control tower. Again, there's notifications. This could come in lots of different ways. I'm kind of jumping through a lot of hoops in one go. But for whatever whichever way that Mary's got that, she understands that there's an assessment that she needs to do to to push this Excuse me.

particular AI asset through its through its life cycle. So she can come into this impact assessment. It's very similar to the first one, but I'm actually diving into one of the tasks a little bit more detailed to showcase some of that capability. And you can see cuz it's assigned to me, I can actually take these assessments. And don't worry, I won't step through every single one of these. But I wanted to showcase this particular feature, right? Because it's very important. And this is where you would have maybe some of that regulatory compliance in terms of, you know, what's the privacy and data protection?

There's all these different questions that you have to answer around that, right? Are we ensuring we've got non-discrimination and fairness? Thinking about our transparency, our accountability, freedom of expression, human oversight, right? Make sure we're covering all of that compliance and needing capability within within the tool that can be contract and used from that. And then of course we could submit that and it would move into others. And then the life cycle would would continue.

And that impersonation. Just come back to it. So I think yeah, that that will that will be it for the demo. I just I hope you found the information very useful, the breakdown very useful. I know there's a lot of new features and all the capabilities. That's the wonderful thing about ServiceNow that comes out. And sometimes it can be hard to sort of get the uptake. So the aim of these is to give you a really good sort of foundational basis of how that works, how it looks. Give you some use cases and hopefully get your understanding.

Looking forward to seeing you in your service AI CT journey. Thank you very much.