

Progress Over Perfection

Episode 4: People, Policy, and Privacy

SURVIVING THE NEW ERA OF FEDERAL AI IMPLEMENTATION

Host: Jack Moore

Guest: Nick Schutt

Series: Progress Over Perfection

Introduction

Jack: Hi, everybody. Welcome back to Progress Over Perfection. I'm your host, Jack. In this episode, I'm joined by Nick Schutt.

Jack: Nick is a serial entrepreneur with a strong track record in government contracting, events, operational technology, and technology innovation. Among other roles, he's the founder and president of Artemis Human Capital Management and host of the Robots and Red Tape Podcast, where he's spent over two years having real conversations about AI with some of the smartest people in the space.

Jack: Nick doesn't just talk about AI. He uses it every day, learns from it, and pushes back on the hype. I wanted to bring him on to share what he's actually learned from his guests and what he really knows about using AI tools.

Jack: We get into why most organizations aren't really ready for AI, and it's not the reason you think. We talk about the biggest mistakes leaders make when rolling out AI tools, why prompting matters more than the tool you're using, and why you should never blindly trust the output.

Jack: Nick also shares real-world examples of where AI is delivering — from cutting military planning cycles from two weeks to five minutes, to reducing eighty hours of monthly compliance work down to an hour and a half. And if you're wondering where Nick says to start in your AI journey, his answer might surprise you. Let's get into it.

Jack:Hi there, Nick. Just want to welcome you to Progress Over Perfection.

Nick:Hey, Jack, thanks so much for having me, man. I really appreciate you considering me as one of your guests.

Jack:Absolutely. Nick's had his Robots and Red Tape podcast for two years, and I aspire to have the following you have someday. Why don't you share a little bit about why you created Robots and Red Tape, what it's focused on, and the type of guests you have.

Why Most Organizations Aren't Ready for AI

Nick:Yeah, so I started it a little over two years ago — it's a side project for me. I run a couple of companies and a non-profit, and I have Robots and Red Tape. I really believe we're at the early stages of AI fundamentally changing how we interact with the world, kind of like the internet and social media have done over the last thirty years.

Nick:My experience with social media specifically has taught me not to trust large technology companies blindly. You only need to look at what happened in 2010 — the ubiquity of the smartphone, the rise of Instagram, and the direct correlation between that combination and rising rates of mental health disorders among children — to see how these societally changing technology rollouts have the propensity to go wrong and have unintended consequences.

Nick:Folks in the government contracting space, or the broader white-collar space, are tired of hearing about AI at this point — everybody's talking about it. But when I talk to friends in warehouses, mechanics, or less technology-driven professions, they barely know how to spell AI.

Nick:I started this because I want to educate people about what AI is, what it's not, what's real, what's hype and marketing — and to have interesting conversations with people who are smarter than me about how AI is going to impact their lives moving forward.

Nick:Something I've learned through this podcast: when people say AI, they mean language models, LLMs, and agents. That's maybe ten percent of what AI actually is, or less — it's the shiny, flashy thing. But all of that sits on top of your data infrastructure.

Nick:I had a gentleman on in the fall, Dr. Gurbinder Dhillon. He's at a company called Sensing now, but before that he was at Dun & Bradstreet for almost twenty years, running data infrastructure for their biggest revenue source. According to him — and he's an absolute technical expert — very few companies out there are actually ready to implement AI because of how disorganized their data is. You don't know where it is. You don't know what data you actually need and what you don't. It's not tagged properly. It's not organized.

The Federal AI Tools Landscape

Jack: I want to dig in here. Most of my listeners are in the federal space or GovCon — federal program managers, division chiefs, people like that. They're all being provided some tools, the ones that are FedRAMP or agency-approved, and they range in quality. Perplexity for Government is out there. Some have access to Copilot, some to other semi-homegrown chatbots.

Jack: So if they've got access, and let's assume the data feeding them is pretty good, they run into a couple of different challenges. One of the big ones is the people side of it — getting people to understand how they might use it within the context of their job, how the outputs are going to look, whether the underlying data is good. If I'm running a division or an agency and I want to get my people using AI, what have you learned — from your own work and from your guests — about effectively deploying it?

The Most Overlooked Piece of Any AI Rollout: Your People

Nick: Technical piece aside, the most important thing and the most overlooked thing is the users. I'm sure you've seen countless tools deployed at different agencies over the years — have you ever seen an effective deployment that didn't properly engage and train users?

Jack: Not at all.

Nick: That's the thing. It doesn't matter how cool the hammer you're using is if the person doesn't know how to swing it. I'd start there — focus on the functions of these people, how they're actually doing their job now, and what you can look at as far as training them to integrate it into their workflows.

Nick: But most importantly, proper training on prompting is huge. The two major areas that really need focus are prompting and not trusting the output. These AI tools — whether homegrown, ChatGPT, Gemini, Copilot, or Claude — are designed the same way social media is designed, to drive user engagement. They're designed to appease you, to make you feel good about the answer. It doesn't matter what you ask it, it's going to give you an answer, whether it actually knows it or not.

Nick: You can't blindly trust these things to hand you a finished work product. I just saw a case where a judge reprimanded lawyers on both sides of a case because they'd used AI to generate their briefs and didn't read them, fact-check them, or proof them — the AI made up cases on both sides. That's embarrassing in court. But if you're doing something that could impact people's lives — strike guidance, intel analysis, drug approvals at FDA — these tools are no match for your people and their judgment.

Nick: The AI doesn't know more than you do. It may help you get through some of the data, but you cannot blindly accept whatever it gives you. Look at the work product, spend time with it, rework it. It's only going to get you maybe seventy-five to eighty percent of the way there, and that's only if you properly prompt it.

Prompting, Tool Comparisons, and Building Context with AI

Jack:In addition to prompting — I was talking with a colleague about this yesterday. She's a relatively new user of AI, and we did an assessment gauging originality: it compares what you come up with against a bunch of different gen-AI outputs, and tells you whether you're more of an augments — someone who takes ideas and adds to them — or a connector, who connects disparate ideas. She asked me what she should do with that. I told her: you bring a unique set of experiences and skills to your job, like everyone does. You actually have to teach the AI who you are and help it understand what you bring versus what it can find for you — how it can augment your skills and experience. Building that context, and helping whatever tool you're using get to know you, is one of the most important elements of this, right alongside properly engineering the prompt.

Nick:Completely agree. That's part of what custom instructions are for — the settings where you can give the AI a standing instruction telling it specifically who you are, how you want it to interact with you, how you want it to answer. Mine, for example, tells it: every single answer needs the bottom line up front. Why do I care about what you're telling me? What's the most important part of this? Then give me the additional detail below, and be as detailed as you can — but summarize first so I know what I'm looking at.

Nick:It's important to also play with different tools, because different AI services are better at different things. For my money right now, Claude is the best one out there overall, based on current performance. Gemini is great for internet searches — obviously, it's Google — and for image generation; I think it does a better job than the others there. ChatGPT is great too, just not quite as good as Claude for most things. I pay for and use all three.

Jack:Our team at the FAA has access to Perplexity for Government and Gemini — those are the two primary tools on their government-furnished laptops. The instructions you're describing are key — teaching it to tell you its source, to be skeptical.

Nick:Your earlier point about the comparison to social media — it's designed to appease you — is something I built directly into my custom instructions: stop appeasing me. Focus on the task I've given you, and if I've made a leap in my prompt, be skeptical of the assumption rather than automatically treating it as correct. Question everything.

Jack:I always think about that in prompting too. I never ask for the thing I want in the first prompt. I start by asking, do you know about this? Tell me what you know. Then I logically reason my way down, filtering information, so the model has the context it needs to understand what I'm actually looking for.

Jack:Have you seen good training the government has put out, or that's available online?

Nick:Personally, I just learned by doing — though I know not everybody learns that way. A good place to start is LinkedIn Learning. And something people don't often think about: depending on where you live, your local library usually has a lot of free online courses and training, though that's more for personal use than corporate.

Nick:Beyond that, what we've done — and what our clients do — is lean into it, start using it, and pay attention to what you put in versus what you get out. Develop your instructions based on what you're seeing from the tool. Then figure out how to augment what you bring to the table rather than trying to have it replace you — because it can't. Everyone has unique skills and experience. The question is how you direct it to surround you with the functionality that makes you better at what you do.

Pattern Matching, Not Understanding

Nick:One of my favorite points from a recent guest: people overestimate how intelligent these tools are, by far. At the most fundamental level, it's pattern matching at scale. It's impressive — I'm not saying it's not — but it's like someone who's read every book in the library and never actually had a conversation.

Jack:That's the line from *Good Will Hunting* — you can quote every book, but until you've lived it and used it in real life, you have no idea what it means.

Nick:Exactly, and that's how these tools operate. They don't have context, they don't actually have conversations, they don't use words — they operate in tokens, and a token is a piece of a word. At the most fundamental level, all these tools are a huge amount of training data trying to predict the next token. That's it. It doesn't actually know what math is — the only reason it knows two plus two equals four is because it's seen that pattern millions of times. If you trained it on data saying two plus two equals five, it would say five, because it's not thinking, it's predicting. And it doesn't even speak in numbers — it's a language model. Two plus two are tokens, not numbers.

Real Results: From Military Planning to Compliance Reporting

Jack:You made a good point earlier that a technology implementation is a technology implementation, whether it's AI or a new accounting system or an HRIS. You're asking people — you need to engage the ultimate users of the system to develop requirements and build a tool that helps them, and you have to sell them upfront on the fact that it's going to provide value.

Jack:We've taken a bit of a cynical view on AI so far. Where does it actually provide value? There are myths around productivity gains — maybe it helps with coding, helps people build their own applications faster. But what verifiable results have you heard from your guests — if you do this, it will help you, and specifically how?

Nick: There's a huge number of genuinely helpful use cases. One that stands out: I was talking with Brian Drake about a human trafficking investigation episode. Something I learned recording that episode is that in these prosecutions, someone has to review the evidence to verify it exists — and the average lifespan of people doing that job is about three months. AI can help go through huge volumes of that material and reduce the amount of time analysts have to spend on the hardest parts of it, extending their productive careers. That's a massive deal.

Nick: Another guest, Clint Alanis of Smack Technologies — a retired Marine officer who ran military planning cycles — has used AI, at both the strategic and tactical level, to take planning cycles from two weeks down to five minutes. Think about what a two-week planning cycle means at the speed of combat today — by the time you're done, it's overcome by events. Being able to do that in near real time changes everything.

Nick: A third example: a friend of mine, Chris Anderson, runs a wine company called Joyner. Shipping alcohol across state lines comes with enormous tax-reporting overhead — every sale has to be reported at the state and often the locality level. That's an incredibly manual process. His tool takes wineries from forty to eighty hours a month of reporting down to about an hour and a half — you drop in your sales reports, it spits out properly formatted filings for every locality. He's now building API connections into the state systems directly.

Jack: I like that example because it pushes back on the idea that AI is just about taking jobs. What we tell our clients and our own employees is that the plan isn't for AI to replace you — it's to remove the meticulous, annoying process, like reporting alcohol shipments across state lines, and free you up to do something more valuable.

Where We Really Are on the AI Adoption Curve

Nick: Right now, comparing where we are with AI to the dot-com boom and the early internet rollout, we're at about the mid-to-late nineties — still getting CDs in the mail from AOL with internet time on them. It's still dial-up. We're early. The AI you use today is the worst it's going to be for the rest of your life.

Nick: I talked about this with my last guest, John Hrastar — people use the phrase exponential growth loosely, usually as an exaggeration, but the way AI is advancing really is exponential. It's crazy how fast it's changing. Starting around February, OpenAI and then Anthropic put out new models that made real leaps forward — and now they're using AI to help build the next generation of models.

The One Thing to Do Monday Morning

Jack:I want to wrap up, and I warned you I'd ask this: if I'm a federal program manager or organizational leader trying to effectively deploy AI in my organization, what's the one thing I can do starting Monday morning to make things better?

Nick:Start with your people. It doesn't matter what kind of technology rollout or process implementation you're looking at — it's all people-based and people-driven. Understand that this is a tool to make your people's lives better and your organization more efficient. Sit with your people, understand what they're actually doing if you don't already, and look for ways to improve the process from there.

Nick:Too often in technology debates, people lose sight of that. It doesn't matter what you're using — email, text, AI — these are all tools meant to improve our lives, work and personal. Start with your people. Figure out how to empower them and make them more efficient.

Jack:That's a great lesson, because if you don't start there, you're just throwing something at people, and you'll create resistance — you haven't sat with them to understand what they think about it. It comes down to a basic level of understanding what AI is, how to use it, correct prompt engineering, and really understanding what your people do, what their challenges are, and how technology can help them do it more productively, if at all.

Jack:Well, I want to thank you, Nick, for joining me on this now fourth episode of Progress Over Perfection. And for those of you listening, make sure you check out Nick's podcast as well, Robots and Red Tape.

Progress Over Perfection is Evans' podcast for federal leaders making consequential decisions under uncertainty. New episodes explore how the best program managers, agency leaders, and modernization teams move from analysis to action.