

Episode 138 | From programmable networks to agentic AI (Chris Wade)
Released April 14, 2026

DR: [00:01] You guys have a great customer, Lumen. You took them from 16 automated workflows to over 350 on your platform, which is amazing. So how'd you guys do that?

Chris Wade: [00:10] Sure. So Greg Freeman, who's a VP of their network over there, he would invite these members from the field into a rotation in the automation team to say, "Hey, bring your idea. Bring your expertise." They're given the opportunity to automate their role and then take those lessons back out to the field. So a lot of times, especially folks like myself get stuck in the nuts and bolts of APIs and MCP, but getting the most people to participate in automation is really the goal here.

Announcer: [00:41] This is "Telco in 20," a podcast that helps telco execs achieve a competitive advantage with AI and the public cloud. It is hosted by Danielle Rios, also known as DR. Today, we're talking to Chris Wade, co-founder and CTO at Itential.

DR: [01:00] Hi, guys. I'm DR. The AI space is en fuego right now. In March alone, three frontier models were released. MCP just crossed 97 million installs. Claude Code went from a research preview to \$2.5 billion in revenue, and NVIDIA launched Rubin and cut inference costs by 10x. That's not a quarter. That's one month, and that pace is splitting the vendor world in two.

[01:28] For AI-first software companies, every new release is a gift. MCP connects AI to any system. Claude Code lets one engineer do the work of five. Agentic frameworks turn months of development into weeks. Every new AI release makes their products better. But for old-school incumbents, every new release makes the gap wider.

[01:51] For 40 years, telco vendors made their APIs and their data their intellectual property. That was their business model. That was their moat. AI just turned their moat into a prison, and they can't break out without breaking their business model. Today, I'm talking to someone who is thriving in the AI-first world. Chris Wade is the co-founder and CTO of Itential, a network automation company that makes infrastructure programmable.

[02:18] Last year, his team launched FlowAI, agents built on MCP that don't just monitor your network. They operate it.

We're going to dig into how Itential is using agentic AI on live production networks, how they layer AI reasoning on top of deterministic guardrails so operators actually trust it, and how Lumen scaled from 16 to 350 automated workflows by putting field engineers in the driver's seat. So let's take 20. Chris Wade is co-founder and CTO at Itential. Hi, Chris. Welcome to "Telco in 20."

Chris Wade: [02:54] Hey, DR. Thanks for having me.

DR: [02:56] Yeah. I'm so excited to have a fellow vendor in telco. We're going to talk about lots of great things, especially how hard it is to get a startup going in telco. You started Itential in 2014, because you saw that networking was about to become programmable, I guess like how compute and storage has become programmable. So how'd you get that idea and that insight?

Chris Wade: [03:19] So we saw networking behaving like mainframes in the 1970s. You buy a Cisco router. You bought an IBM mainframe with power cords and the applications preinstalled. And like you said, what we saw with compute and storage, it just seemed like we were headed there with networking, with all the cloud stuff you've done, like VPCs, and then all of a sudden, transit gateways come along, and you're like, "This is the model going forward."

[03:39] So we felt like routers, switches, firewalls, load balancers were going to be more of a cloud operating model. It had to become programmable. In 2014, we were doing OpenFlow and the first SDN controllers and that kind of stuff, but our focus has always been on programmability. We really feel like that unlocks all sorts of opportunities for building applications and changing how we operate infrastructure.

DR: [04:00] Yeah. I mean, definitely moving things out of hardware into software is, from operating a business, so much easier, because it's fewer things. Software has a higher margin, and it's easy to change. It's easier to update. And so, AI is huge everywhere. Every vendor at MWC was talking about agents and AI. You saw it painted on their booths, in their marketing literature. You guys recently launched FlowAI, agents that can actually operate infrastructure, not just answer questions about it. And so, how does that work? Tell me about FlowAI.

Chris Wade: [04:34] Sure. So our core platform is an orchestration platform. We don't come from a really heavy data background. So a lot of the early AI stuff was extensions of ML. It was chatbots. It was, "How do I basically allow AI to interface with humans?" So we were really focused on, "How do we build agentic platforms that can actually operate infrastructure?" So that's really what we're focused on.

[04:54] So adding that reasoning layer to determinism has been our focus. With FlowAI, we're really adding that agentic layer, but leaving the determinism in place between that agent and the infrastructure as people start to think about how they want to allow agents to touch what they consider critical infrastructure.

DR: [05:12] Yeah. And I think that's a double-edged sword a little bit, because I think what makes AI valuable and powerful is taking action, doing something, like you said, not just providing a recommendation that a human is still deciding what to do. And the reason I call it a double-edged sword is our customers, telco customers, get a little bit afraid that it's going to be wrong. It's going to make a wrong decision. And so, I guess that's what you're talking about in terms of the determinism and keeping the guardrails on.

Chris Wade: [05:40] Exactly. You started talking about 2014. So I started thinking in the early days of Itential. When we started adopting automation, I felt the same resistance and reluctance that somebody on a CLI can be replaced by scripts or Python or Ansible or whatever. So trying to have this agentic layer take action on the network, it's a super good parallel to what we went through for automation.

[06:00] I need a human-in-the-loop. I need people reviewing things. I want to see what audit's happening. So the nondeterminism obviously adds a next layer of issues in our head trying to reconcile how that works, but the nondeterministic reasoning and the deterministic layer have so many similarities that we like to look at it through that lens, and we're in the trust and confidence business. So getting customers comfortable with how this is going to work is really important for us.

DR: [06:23] Yeah. And so, this isn't just a bare LLM interacting with your agents. I think you guys are using and betting heavily on MCP, Model Context Protocol, as the reasoning layer between the agents and your product. And so, connecting data isn't the same as understanding the business itself. You've got to really give it context. And so, how do you guys capture your customers' business context?

Chris Wade: [06:47] So I think this is the part of the space that's moved a lot in the last six months. So we launched our MCP server in May of last year. So in our little telco corner of the universe, that's pretty early in the cycle. When we start talking about business context, the skills that have been added. The skills really allow you to embed domain knowledge into what the agent's going to do. So back with automation, we took MOP documents and design documents, and we tried to encode it in pipelines.

DR: [07:12] Right.

Chris Wade: [07:12] And now we're going back to, if I could give the agent the design documents and my methods of procedure, now it can reason through that. So it's really the combination of context and those skills. We started off with prompting and stuff. Now, you're seeing spec-driven development. You're seeing a more robust version of that prompting.

[07:29] But if you take the skills, which is that domain knowledge, and you combine it with MCP, which is the tools that you're going to give that agent, that's the secret sauce we're seeing right now, is the combination of those two things.

DR: [07:39] Yeah. And how are you keeping that all in sync? Because I don't really think people are going to have a small number of agents in their telco organization. They're going to have hundreds or maybe even thousands. And so, within your product, how are you sharing that across agents? Are you going in and updating every agent, or is there some sort of base layer that inherits this change?

Chris Wade: [07:59] Sure. So this might be a technical way to think about it, but we're really thinking about what we put in the context window each time. So there are system prompts that Intentional inserts that are an explanation of what it means to automate or orchestrate a service or to do maintenance in the network. So

you don't have to put that in each place, but each of the agents from a user prompt perspective that our customers would input is the tailoring of that business activity.

DR: [08:23] Got it.

Chris Wade: [08:24] So if you're running a telco network versus DOCSIS network versus enterprise IP services, you would put that context, and it would share it across each of those agents.

DR: [08:33] Got it. Everyone's really focused on, "How do I get the data from underlying legacy systems into my agents?" But there's this whole other side, which is context, and it could be processes, authority levels, rules, and that is reality. That is how telcos run their business and make decisions about the network or about their customers, or whatever it might be.

[08:56] And I think it's been one of the biggest challenges we've had with legacy in telco, is that a lot of that context is embedded in code and in customizations and in integrations, and I think AI now gives you an ability to expose it in specs, written descriptions of how we run our business. And so, we have to really think about how we're managing and maintaining that, because it's easy in a silo, and it's very difficult at scale, and I think that's where we are right now when it comes to AI.

Chris Wade: [09:25] Sure.

DR: [09:25] And so, you guys are touching live production networks. Right? Every time I go talk to a telco about their network, they're like, "It's an act of Congress to get a change in the network." And so, how are you guys getting people comfortable that your AI won't make a mistake?

Chris Wade: [09:39] Sure. So I think we have to start with the premise that AI is going to make a mistake, the context we give it, the MCP tools we give it. It's going to make the right decision with that context, but as business owners, it might not make the decision we want it to make. So we have to start with the same components and governance we had with automation, with logs and auditing, and all the stuff we've talked about. But I think the other thing we need to think about is, we've been encoding all of these things, and the good thing about determinism is it makes the same decision every time.

DR: [10:06] Yup.

Chris Wade: [10:06] The bad decision is that it makes the same decision every time. So we've had to encode all these, what if, else this. What if the device is offline? What if it's on Tuesday? What if we're having a weather event? And once we can remove that and do the combination of determinism with the reasoning, you actually have a smaller footprint of technical debt to deal with, and I think we're going to find ourselves in a much more robust situation.

[10:30] Even with automation, we want to have perfect data. We want a perfect infrastructure as code and perfect outcomes. But the reality is, especially at telco, I mean, what's inventory? 62% accurate? And we have that infrastructure as code with humans. Now, we have this intelligence from reasoning. The combination of the two can really overcome the traditional challenges we've had just to scale automation.

DR: [10:50] Yeah. In the Totogi ontology, which is my company's product, one thing we've done is really built in some ability to see the decisioning that the AI is doing, and then dynamically change it right there on the spot. You can see why it made the decision that it did. Was that a rule that we set? Was that pulling from some certain spec or document?

[11:11] But you build something small. Humans are overlooking it. They see that it works, and then they see the logs. They see the traces of the decisioning, and then we expand and grow. And so, it's been able to build trust pretty quickly.

Chris Wade: [11:24] Sure.

DR: [11:24] You guys have a great customer, Lumen, that's been able to really expand their footprint with your product. You took them from 16 automated workflows to over 350 on your platform, which is amazing. So how'd you guys do that?

Chris Wade: [11:37] Sure. So Greg Freeman, who's a VP of their network over there, I think he really focused on the human element. A lot of our customers get mired in the details of APIs and pipelines. The goal here is to get the most people to participate in automation. Most people build small tiger teams, software developer teams that take a bunch of the automation efforts.

[11:56] Greg took the opposite approach. He knew that the people in the field had the expertise of what he needed to automate. So he built a core team that understood the tooling, and he would invite these members from the field into a rotation in the automation team to say, "Hey, bring your idea. Bring your expertise," because the core team doesn't necessarily have the expertise of somebody who's doing transport equipment in a certain region.

[12:17] So they come in with their idea. They're given some training, the opportunity to automate their role, and then either stay or rotate back out, and maybe take those lessons back out to the field. So he's done a wonderful job with that. And then the benefit with AI on top is, now we have all these automations within the system. And now, when we start to put NLP interfaces on top, we start to hook it into maybe customer order systems or help desk systems or the NOC.

[12:41] Now, people, through agentic interfaces, NLP interfaces, can now access those automations that those experts from the field came in. So it's really a story of human transformation. A lot of times, especially folks like myself get stuck in the nuts and bolts of APIs and MCP, but getting the most people to participate in automation is really the goal here.

DR: [13:00] That's super cool. It sounds like you guys were able to build a little bit of a flywheel of success.

Chris Wade: [13:04] Yeah.

DR: [13:05] Well, I was just talking about how telcos and their culture, it's slow. They almost brag about how slow they are. And so, we both know that building a startup is incredibly difficult. Most startups fail. I think nine out of 10 don't make it. And so, the telco industry, it's a whole new level of brutality, long sales cycles, incumbents that have been around forever. Plus, you need a ton of money to go the distance.

[13:30] And so, my advice to any founder thinking about starting a telco-focused company is you pretty much have to 10x your timeline expectations, make sure you have great investors that understand the industry, because this is going to take way longer than you think. And so, you've been doing it for over 10 years. What's your advice on how to make it in our industry?

Episode 138 | From programmable networks to agentic AI (Chris Wade)
Released April 14, 2026

Chris Wade: [13:50] So I think my comments are more on what we need to do to help move the industry forward so we can create a startup ecosystem.

DR: [13:57] Sure.

Chris Wade: [13:57] Our industry is one where we believe the software is specific to telco. We have a lot of verticalized vendors in our space that basically support the telco industry. And when I look at what we're doing with networking, I'm seeing more similarities between hyperscalers, enterprises, and telcos than I've ever seen, maybe less the RAN.

DR: [14:16] Yeah.

Chris Wade: [14:16] So I think as an industry, we need to move towards thinking of broad-based networks and maybe less vertical for us, so we can get the benefit and the R&D and the scale. And my traditional issue with our space is that the vendors consider their APIs and their data intellectual property, and we've tried to do the TMF thing. But what is successful is we need vendors to publish their APIs. We need to be an open ecosystem. We don't ask the hyperscalers to create a standard interface for using compute or storage, but in our industry, that's been acceptable.

DR: [14:48] Yeah.

Chris Wade: [14:49] So we really need telcos to ask vendors to publish their APIs to operate in an open way, and that's the only way we're going to foster startups in our space. We really need to broaden, bring in some enterprise software space, and also ask maybe our telco friends in procurement to not buy software unless their vendors are open and publish APIs.

DR: [15:09] Well, I think it's going to be more than that on just that point. I think especially with AI, the old-school vendors are going to need to provide a command-line interface, a CLI, to their products, or else people are going to swap them out for products who do. And so, I think just the market's going to demand that, and they were able to build their walled gardens for 40 years. That day is ending, and they're going to be replaced.

[15:37] And again, it won't be overnight. It'll probably take a decade. But if they don't start moving and embracing AI and openness so they can be part of a broader ecosystem, and not just, again, building their walls around their product, I think their days are numbered. So I totally agree with you on that.

Chris Wade: [15:54] Sure.

DR: [15:55] Well, Chris, this was an awesome conversation learning more about Itential. So thanks so much for coming onto the podcast.

Chris Wade: [16:02] Thanks for having me.

DR: [16:03] Awesome. Stick around. We end each podcast with a "Telco in 20" takeaway. I've got two minutes to tell you something you need to know. Chris and I just talked about AI agents operating infrastructure. It made me think about something Aaron Levie, the CEO of Box, posted recently on X, that AI agents, not humans, are becoming software's primary users.

[16:33] Think about what that means. Software's most important user isn't a person staring at a screen anymore. It's an agent calling APIs and command-line interfaces. If a system can't be accessed this way, it doesn't matter how good it is. To AI, it does not exist. This is a massive problem for telco. Most of your vendor systems are locked behind proprietary interfaces, completely invisible to AI, not because the technology isn't there, but because their business model depends on keeping you trapped inside their walled garden where only their consultants can access it.

[17:08] But in an agentic world, every system you operate needs to be open so any AI can reach it, not just your dinosaur vendor's copilot. You need agents working across billing, provisioning, customer care, and the network. And when the agents hit a system they can't talk to, your whole AI strategy fails. So here's my challenge to you. Ask your vendor one question, "Can my AI agents access your system through an open API or CLI?" If the answer is no, you now know which vendor is going to be the problem.

Episode 138 | From programmable networks to agentic AI (Chris Wade)
Released April 14, 2026

[17:41] Stop evaluating vendors on what their AI agents can do, and start evaluating them on what your AI can reach, because in a world where agents are the primary users of software, a closed system isn't just inconvenient. It's irrelevant. And if your vendors won't give you that access, Totogi can add it for you, on top of any vendor system. No rip and replace required.

[18:07] DM me on LinkedIn or X, @TelcoDR, and I'll show you how. Until then, tune into more "Telco in 20" episodes. Like and follow, and leave us a five-star review. Don't forget to sign up for my must-read email newsletter on TelcoDR.com, and check out our awesome YouTube channel and hit that subscribe button. Later, nerds.