MY WEIRD PROMPTS

Podcast Transcript

EPISODE #201

The Open Door: How Fast Can Hackers Find Your Server?

Published January 08, 2026 • Runtime: 23:58

https://myweirdprompts.com/episode/unprotected-server-background-noise/

EPISODE SYNOPSIS

What happens if you leave a server online without a password? In this episode of My Weird Prompts, Herman and Corn dive into the "background radiation" of the internet—the constant, automated scanning by botnets looking for any open door. From Z Map scans to the monetization of compromised servers through crypto-mining and Initial Access Brokers, discover why your digital security is under threat the moment you go live. Learn how the ecosystem of WordPress plugins and the rise of AI-augmented scanning are changing the landscape of cyber defense in 2026.

DANIEL'S PROMPT

Daniel

Firewalls are essential for network security, protecting against malicious actors. Sometimes, networks are left exposed to create "honeypots" for cybersecurity research. Beyond sophisticated, state-sponsored campaigns, there is a constant background noise of automated botnet activity on the internet. WordPress is a major target due to its large attack surface. These large-scale automated attacks are often used for things like cryptocurrency mining or ad injection. Herman and Corin, if someone were to set up an unprotected server on the internet—like a CRM with no password—how long would it take for it to be discovered and compromised? How does this background level of hacking work?

TRANSCRIPT

Corn

Hey everyone, welcome back to My Weird Prompts. I am Corn, and I am sitting here in our living room in Jerusalem with my brother.

Herman

Herman Poppleberry, at your service. It is a beautiful day outside, but we are about to talk about some of the darker corners of the digital world.

Corn

Exactly. Our housemate Daniel sent us a really interesting audio prompt today. He was asking about what happens if you just... leave the door open on the internet. Specifically, if you set up a server, like a customer relationship management system or a database, and you do not put a password on it. He wants to know how long it takes for someone to find it and what that background noise of the internet actually looks like.

Herman

I love this question because it touches on something most people do not see. We think of the internet as this place we go to visit websites, but beneath the surface, there is this constant, pulsing rhythm of automated activity. It is like the background radiation of the digital universe.

Corn

It is funny you say that because when we were talking about home networks back in episode two hundred and seventy, we touched on how firewalls are our first line of defense. But Daniel’s prompt goes a step further. He is asking about the predators in that environment. So, Herman, let’s start with the big question. If I spin up a fresh server on a cloud provider like Amazon Web Services or DigitalOcean right now, and I leave it totally unprotected, how long do I have?

Herman

The short answer? Minutes. Sometimes less. It is actually a bit terrifying when you see it in real time. There have been many studies over the last decade where researchers set up honeypots, which are essentially decoy servers designed to be attacked, and they often see the first connection attempts within minutes, sometimes even within a minute of the IP address becoming reachable on the public internet.

Corn

Sixty seconds? That feels almost impossible. How is someone sitting at a keyboard finding a random new server that quickly?

Herman

Oh, they are not. That is the key. No human is doing this manually. It is all automated. There are these massive botnets and scanning tools that are constantly traversing the entire IPv4 address space. If you think about it, the total number of IPv4 addresses is about four point three billion. That sounds like a lot, but with modern computing power, you can scan the entire thing in a surprisingly short amount of time.

Corn

Right, I remember reading about a tool called Z Map. Doesn’t that allow a single machine with a good connection to scan the whole internet in under an hour?

Herman

Exactly. Z Map and another one called Masscan are the industry standards for this. Back in the day, scanning was slow. You had to wait for a response from each IP before moving to the next. But these tools are asynchronous. They just fire off packets as fast as the network card can handle them and then listen for the echoes. So, if you are a bad actor, or even just a researcher, you can have a fresh map of every open port on the entire internet every single day. And by twenty twenty-six, security researchers are increasingly discussing AI■augmented scanning, where bots may start using predictive models to guess which IP ranges are likely to have newly provisioned cloud servers, though this is still an emerging area.

Corn

So, when Daniel asks how long it takes for a server to be discovered, the answer is basically as long as it takes for one of these global scans to cycle back to your specific IP address.

Herman

Precisely. And because there are thousands of different groups running these scans simultaneously for different reasons, your discovery time is essentially very close to zero. The moment you are on, you are visible.

Corn

That brings up an interesting point about what they are looking for. Daniel mentioned a CRM with no password. Is a bot smart enough to recognize a specific piece of software, or is it just looking for any open door?

Herman

It is a multi-stage process. First, they do what we call port scanning. They check if common doors are open. Port eighty for web traffic, port four four three for encrypted web traffic, port twenty-two for secure shell access, and so on. If they find a door open, the next step is fingerprinting. The bot sends a specific request to that port and looks at the response. Every piece of software has a unique signature or banner that it sends back. It might say, Hey, I am an Apache web server version two point four, or I am a Microsoft SQL server.

Corn

And once they know what it is, they look up the menu of known vulnerabilities for that specific version.

Herman

You got it. It is like a burglar walking down a street. First, they check if any front doors are unlocked. That is the port scan. If they find one, they look at the brand of the lock or the type of security system. That is the fingerprinting. Then, they reach into their bag for the specific master key that fits that brand. In the digital world, that master key is an exploit.

Corn

We should talk about WordPress here because Daniel mentioned it specifically. It is such a massive target. I think something like forty-three percent of all websites run on WordPress as of this year, twenty twenty-six. Why is it such a magnet for this background noise?

Herman

It is the sheer scale. If you are writing a script to automate attacks, you want the biggest return on your investment. If you find a vulnerability in a niche blogging platform used by ten people, it is not worth your time. But if you find a flaw in a WordPress plugin that is installed on millions of sites, you just hit the jackpot. In past incidents, there have been massive waves of attacks targeting vulnerabilities in popular plugins, and within hours of some of those exploits being published, automated bots have been observed compromising tens of thousands of sites.

Corn

And it is not just the core WordPress software, right? It is the ecosystem.

Herman

Exactly. The core software is actually quite secure because it is scrutinized by so many people. The problem is the long tail of plugins and themes. Someone might install a cool photo gallery plugin they found five years ago, and the developer stopped updating it three years ago. If a vulnerability is found in that plugin today, every single site using it is suddenly a sitting duck.

Corn

I have seen this on some of the sites I help manage. You look at the logs and you see thousands of attempts to access files like wp-config dot p-h-p or xml-rpc dot p-h-p. It is relentless.

Herman

That xml-rpc one is a classic. It was originally designed for remote posting and pingbacks, but it became a favorite for brute-force attacks. Instead of trying one password at a time through the login page, an attacker could use a single xml-rpc request to try hundreds of password combinations at once. It is much more efficient for the bots.

Corn

So, let’s go back to Daniel’s hypothetical. I have a server, it has a CRM on it, and it has no password. The bot finds the port, fingerprints the CRM, realizes there is no authentication required, and then... what? What is the actual compromise that happens in those first few minutes?

Herman

Usually, the first thing that happens is a script kiddie or an automated worm will drop a web shell. This is a tiny piece of code that allows the attacker to run commands on your server through a web browser. Once they have that, they have a foothold. Now, here is where it gets interesting: what do they actually want? In twenty twenty-six, the most common use for a random, low-value server is not stealing your data—unless you are a high-value target. Most of the time, they just want your resources.

Corn

Like cryptocurrency mining?

Herman

That is a huge one. Even though the era of easy Bitcoin mining is long gone, there are plenty of privacy-focused coins like Monero that can still be mined on standard server CPUs. If an attacker can compromise ten thousand random servers and run a mining script on each one, they are making pure profit because you are the one paying the electricity and the cloud hosting bill.

Corn

That is such a parasitic way of operating. They aren't even looking at your files; they just want your processor cycles.

Herman

Exactly. Or they use your server as a proxy or a relay. If they want to attack a bank or a government agency, they don't want the attack to come from their own IP address. They want it to look like it is coming from your unprotected CRM server in Jerusalem. It provides them with a layer of anonymity.

Corn

So your server becomes a foot soldier in a larger botnet. We talked about this a bit in episode eighty-one when we were discussing how firewalls handle stateful inspection. But it is wild to think that your empty server could be contributing to a massive distributed denial of service attack on the other side of the world.

Herman

It happens every day. And there is also ad injection. If they get into your web server, they might not take it down. They might just inject a few lines of code that show extra ads to your visitors or redirect them to affiliate sites. You might not even notice it for weeks, but they are siphoning off a little bit of value from every person who visits your site. And increasingly, security researchers warn that data scraped from CRMs could be fed into AI models used for social engineering. Your customer list isn't just a list of names; it's a goldmine for training a phishing bot to sound exactly like your company.

Corn

You mentioned earlier that most of this is automated. But is there ever a point where a human steps in? Like, if the bot finds something that looks particularly juicy?

Herman

Definitely. This is where we get into the concept of Initial Access Brokers. These are specialized hackers who don't necessarily want to do the final crime. Their whole business model is just finding the open doors. They run the bots, they gather the list of compromised servers, and then they sell that list on dark web forums.

Corn

So it’s like a real estate agent for stolen servers.

Herman

Precisely. They might say, I have access to five hundred servers in the healthcare sector with these specific specs, and then a ransomware group will buy that access to launch a targeted attack. So, Daniel’s unprotected CRM might start as a target for a simple miner, but if the bot sees that the server is part of a corporate network or has access to sensitive databases, it gets flagged for a human to take a closer look.

Corn

That is the lateral movement piece. Once they are in one room of the house, they start looking for the vents and the crawlspaces to get into the rest of the building.

Herman

Right. And that is why the no password scenario is so dangerous. It is not just about that one server; it is about what that server can see. In a modern office, a CRM usually has an API key to the email system, or a connection to the billing database. If I am an attacker and I find your CRM is wide open, I am going to scrape every single API key and password I can find in the configuration files.

Corn

This reminds me of the Shodan search engine. For those who don't know, Shodan is often called the search engine for the Internet of Things. Instead of searching for web pages like Google, you search for devices.

Herman

Oh, Shodan is a goldmine for researchers and, unfortunately, for attackers too. You can literally search for authentication disabled or default password and find thousands of webcams, industrial control systems, and even medical devices that are just... sitting there. It is the public index of the background noise we are talking about.

Corn

I remember seeing a report where someone found a building’s entire heating and cooling system accessible via an unencrypted page with no login. You could literally change the temperature of an office building on the other side of the planet with a click.

Herman

It is a massive problem. And it brings us back to Daniel's question about background noise. A lot of people think that if they aren't a big company, nobody is looking for them. But these bots don't care who you are. They are just looking for the string of numbers that represents your IP address. It is entirely impersonal.

Corn

It's like the ocean. The tide doesn't care if you're a giant ship or a little sandcastle; it's going to wash over you regardless.

Herman

That's a great analogy. Now, there is one nuance I want to add here about twenty twenty-six. We are seeing a shift because of IPv6. In the old IPv4 world, there are only four billion addresses, which is small enough to scan. But IPv6 has... well, it's a number so large it's hard to wrap your head around. It is three hundred and forty undecillion addresses.

Corn

Wait, spell that out for me. How many zeros?

Herman

It is a thirty-nine-digit number. If you tried to brute-force scan the entire IPv6 space at the same speed we scan IPv4, it would take astronomically longer than the current age of the universe—effectively impossible.

Corn

So, does that mean we are safer in an IPv6 world?

Herman

Yes and no. It means blind scanning is much harder. An attacker can't just start at address zero and go to the end. They have to be more strategic. They look at DNS records, they look at leaked data, or they wait for your device to call home to a server they control. So the background noise is changing. It is becoming less about where is every open door and more about let me follow the people I see walking around.

Corn

That is a fascinating shift. It moves from a broad-spectrum broadcast to a more targeted tracking approach. But let’s talk about the honeypots Daniel mentioned. This is the good guy version of this activity, right?

Herman

Exactly. Cybersecurity researchers set up these intentionally vulnerable servers to see what the attackers are doing. It is like a biological research lab that studies viruses to create vaccines. By watching how a bot tries to break into an unprotected CRM, researchers can identify new zero-day vulnerabilities or track the evolution of botnet code.

Corn

I’ve heard of Canary Tokens too. They are like digital tripwires. You leave a file named passwords dot t-x-t on your server, but it’s actually a special file that sends you an alert the moment someone opens it.

Herman

I love Canary Tokens. They are a great way to get a heads-up that someone is poking around where they shouldn't be. It changes the game from passive defense to active monitoring.

Corn

So, if we were to give Daniel a concrete timeline... if he sets up that CRM today, at two p-m in Jerusalem, with no password. By two-zero-five, he has been scanned. By two-fifteen, a bot has probably identified the software. By two-thirty, there is a good chance an automated exploit has been attempted. And by dinner time, his server is likely mining Monero for someone in another country.

Herman

That is a very realistic timeline. In some cases, it might even be faster. If his IP address happens to be in a range that is currently being targeted by a specific campaign, it could be a matter of minutes.

Corn

It really puts the security by obscurity myth to bed. You are never too small to be noticed because the things doing the noticing aren't human.

Herman

Exactly. And that leads us to some practical takeaways. Because it’s not all doom and gloom. We know how to stop this.

Corn

Right. The first one is obvious: never, ever put a service on the public internet without authentication. Even if it’s just for a minute while you test something. That minute is all a bot needs.

Herman

And if you do need to expose something, use a non-standard port. It won't stop a dedicated attacker, but it will filter out a lot of the basic background noise bots that only look at the common ports.

Corn

Also, allow-listing. If only you and I need to access that CRM, we should set the firewall to only allow traffic from our specific home IP addresses. To the rest of the world, the server will look like it doesn't even exist.

Herman

That is the stealth approach, and it is incredibly effective. Another big one for twenty twenty-six is the move toward Zero Trust architecture. The idea is that you never trust a connection just because it’s inside your network. Every single request, even from one of your own servers to another, should be authenticated and encrypted.

Corn

It’s like having a lock on every internal door in the house, not just the front door.

Herman

Exactly. So if a bot does get into that unprotected CRM, it can't just walk over to the database. It hits another locked door.

Corn

I also think it’s worth mentioning that cloud providers have gotten a lot better at this. If you use something like AWS, they often have Security Groups that are closed by default. You have to intentionally go in and open a port to the world.

Herman

True, but as we always say, the most dangerous part of any system is the human configuration. People get frustrated when things don't work, so they just hit the allow all button to see if that fixes the connection issue. And then they forget to turn it back off.

Corn

The I’ll fix it later trap. That is where the bots live. They live in that gap between I’ll just test this and I’ll secure it tomorrow.

Herman

You know, I was looking at some logs from our own router the other day—remember when we talked about the home network in episode two hundred and seventy? Even our little home IP address gets hit thousands of times a day. Most of it is just knocking on port twenty-two, looking for a Raspberry Pi with a default password of raspberry.

Corn

It's wild to think that while we are sitting here drinking coffee, there are digital entities all over the world trying the handle on our front door.

Herman

It really is a different way of looking at the world. It's not a static place; it's a constant, dynamic struggle.

Corn

So, to wrap up Daniel's question, the background level of hacking is essentially a global, automated, twenty-four-seven audit of every connected device. It is driven by the desire for resources—computing power, bandwidth, and anonymity. And for popular platforms like WordPress, the pressure is even higher because the exploit-to-reward ratio is so good for the attackers.

Herman

Well said. It's not personal; it's just physics. Digital physics. If there is a void—an unprotected server—the pressure of the internet will rush in to fill it.

Corn

I think that is a perfect place to leave it. Daniel, thanks for the prompt. It definitely made me want to go back and double-check my own server settings.

Herman

Same here. I might go rotate some API keys just for the peace of mind.

Corn

Before we go, I want to say a huge thank you to everyone who has been listening. We've reached nearly three hundred episodes now, and the community feedback has been incredible. If you are enjoying these deep dives, please take a moment to leave us a review on your podcast app or on Spotify. It really does help other curious people find the show.

Herman

It really does. And if you have your own weird prompt you want us to tackle, head over to our website at myweirdprompts dot com. There is a contact form there where you can send us a message or an audio clip just like Daniel did.

Corn

You can also find our full RSS feed there if you want to make sure you never miss an episode. We are on Spotify as well, of course.

Herman

This has been My Weird Prompts. I’m Herman Poppleberry.

Corn

And I’m Corn. Thanks for diving down the rabbit hole with us. We will see you next week.

Herman

Stay secure out there!

Corn

Bye everyone.