JULY 2026 EDITION

SkillWeed

AGENTIC AI & THE NEW RISK FRONTIER THE GRC PULSE

GOVERNING THE NEXT WAVE OF AUTONOMOUS AI SYSTEMS

TABLE OF CONTENTS

EDITOR'S NOTE

Every few months, the ground shifts under this industry just enough that the people who noticed early end up with a real advantage. July is one of those moments.

Two things happened in the last few weeks that will shape how Cyber GRC professionals work for the rest of the year. First, AI stopped being something organizations merely use, and became something that acts on its own — agentic systems that plan, call tools, and make multi-step decisions with very little human supervision. Second, regulators blinked. The European Parliament approved a package of delays to the EU AI Act's toughest obligations, buying organizations more runway, but not an exemption.

Both of these stories point to the same lesson: the rules are not getting simpler, they are getting more dynamic. The professionals who win in this environment are not the ones who memorize a framework once and stop. They are the ones who treat governance as a living practice — checking what changed, asking what it means for their organization, and adjusting before the next audit forces them to.

This edition is built around that idea. We look at why agentic AI is a different category of risk than the chatbots most policies were written for. We break down exactly what did and didn't change in the EU AI Act's newest delay. And because we never want this publication to speak only to people who are already deep in the field, we've included a beginner's playbook for writing your very first AI governance policy — no jargon, no assumed experience.

Whether you are five years into a GRC career or five weeks into deciding this is the field for you, welcome to the July edition. Let's stay ahead of where the industry is going, not where it has been.

Akingbade Akinfenwa

CHAPTER 1:
AGENTIC AI — GOVERNANCE'S NEXT FRONTIER

FROM CHATBOTS TO AUTONOMOUS ACTION

Most AI governance policies written in the last two years were built around a simple mental model: a person asks a question, a model answers, a human decides what to do with that answer. Agentic AI breaks that model.

An agentic system doesn't just respond — it plans a sequence of steps, calls tools and APIs, reads and writes data, and moves on to the next step without waiting for a person to approve each one. That can mean drafting an email and sending it, querying a database and updating a record, or chaining several other software services together to complete a task end to end.

WHY THIS CHANGES THE RISK CALCULUS

- Fewer checkpoints. A single bad assumption early in a chain of actions can compound across every step that follows, with no human in the loop to catch it.
- Real credentials, real access. To act, an agent typically needs standing access to systems — email, code repositories, cloud infrastructure, customer data — which makes it a genuine identity to govern, not just a tool to monitor.
- Harder to audit. Reviewing one AI-generated output is straightforward. Reviewing a chain of twelve autonomous decisions, each influencing the next, is a different kind of audit problem entirely.
- New attack surface. Every tool an agent is authorized to call is a potential path for an attacker who compromises that agent's credentials.

WHAT'S ALREADY GOING WRONG

The risk isn't theoretical. In June 2026, security researchers reported that an employee at Vercel had signed up for an AI productivity tool, Context.ai, using their corporate Google Workspace account and granted it broad, "allow all" permissions. When attackers later compromised Context.ai itself, they inherited that employee's access and used it to reach into Vercel's own environment — a company that wasn't even a customer of the breached AI tool. Security researchers covering the incident pointed to OAuth tokens, not any novel AI exploit, as the actual attack surface.

That detail matters for governance teams: the danger of agentic and AI-connected tools rarely comes from the AI model misbehaving. It comes from how much standing access gets handed out the moment someone clicks "allow," and how few organizations track which apps hold which permissions across their identity providers.

GOVERNING AGENTIC SYSTEMS

- Scope every credential. Give each agent the minimum access it needs for its specific job, never a standing "allow all" grant.
- Add approval checkpoints for high-impact actions. Sending money, deleting data, or messaging customers externally should pause for human sign-off, even in an otherwise autonomous workflow.

- Log every action, not just every output. An audit trail that only shows the final answer is not enough; you need the full chain of steps the agent took to get there.
- Treat new AI tools as vendor risk, not personal productivity choices. Move OAuth consent to admin review before any new app can touch corporate accounts.
- Build a kill switch. Know how to immediately revoke an agent's access if it starts behaving unexpectedly, before damage compounds.

THE BOTTOM LINE

Agentic AI isn't dangerous because it's smarter than the tools before it. It's dangerous because it acts, and action without a checkpoint is exactly the kind of risk traditional GRC frameworks were built to catch. The skill that matters now is translating governance principles you already know — least privilege, segregation of duties, audit trails — onto a category of system that didn't exist when most of those principles were written.

CHAPTER 2:
INSIDE THE EU AI ACT'S DIGITAL OMNIBUS

WHAT JUST CHANGED, AND WHAT DIDN'T

If you've been tracking the EU AI Act's August 2026 deadline with some anxiety, there's real news to share — but it is relief, not an exit ramp.

On 16 June 2026, the European Parliament approved a package of amendments known as the Digital Omnibus on AI, following a political agreement reached by EU negotiators in May. The Council of the EU still needs to formally adopt the package, but the direction of travel is now clear.

WHAT'S BEING DELAYED

- High-risk Annex III systems (use-based high-risk AI, such as tools used in hiring, credit decisions, or law enforcement) move from an August 2026 deadline to 2 December 2027 — a 16-month postponement.
- High-risk Annex I systems embedded in already-regulated products (medical devices, lifts, machinery) shift from August 2027 to August 2028.
- The watermarking requirement for AI-generated content under Article 50(2) moves from August 2026 to 2 December 2026 — a shorter, four-month grace period.
- National AI regulatory sandboxes, which member states were required to stand up, now have until August 2027 instead of August 2026.

WHAT'S NOT CHANGING

- Most Article 50 transparency duties still apply from 2 August 2026 — including the requirement to tell people when they're interacting with an AI system such as a chatbot.
- Prohibited practices remain banned, and a new prohibition was added: AI systems that generate child sexual abuse material or non-consensual intimate imagery are now explicitly outlawed, with compliance required by 2 December 2026.
- GPAI model obligations, already in force since August 2025, are untouched by this package.

WHAT THIS MEANS FOR YOU

The temptation after good news like this is to relax the timeline internally. Resist it. The delay gives organizations more time to build a high-risk AI compliance program properly — with real risk assessments, documentation, and testing — instead of rushing a thin version of it to hit a deadline that, until recently, looked unmissable. Transparency obligations are still landing on schedule in August, and the package itself isn't final until the Council formally adopts it and it's published in the Official Journal.

THE BOTTOM LINE

Treat the extra runway as a gift to your program, not your priority list. Organizations that use this delay to build durable AI governance structures will be in a completely different position by December 2027 than those that simply pushed the deadline to the back of their calendar.

CHAPTER 3:
WHAT JUST CHANGED, AND WHAT DIDN'T

If chapters like the two before this one feel like a lot, that's normal. Most people entering Cyber GRC right now are being asked to understand AI governance before they've even written their first policy of any kind. So this section skips the theory and gives you a starting point you can actually use this week.

STEP 1 — INVENTORY WHAT'S ALREADY HAPPENING

Before writing any policy, find out what AI tools your organization (or team, if you're practicing on a smaller scale) is already using — officially and unofficially. Ask colleagues directly. Most “Shadow AI” isn't malicious; people just never got asked.

STEP 2 — SORT TOOLS BY WHAT THEY TOUCH

- Low risk: tools that only touch public information (e.g. drafting a generic blog outline).
- Medium risk: tools that touch internal but non-sensitive information (e.g. summarizing a meeting transcript).
- High risk: tools that touch customer data, financial data, source code, or anything regulated (e.g. an AI tool with access to a CRM or codebase).

STEP 3 — WRITE A ONE-PAGE POLICY, NOT A FORTY-PAGE ONE

Your first AI usage policy should answer four questions in plain language: what tools are approved, what data may never be entered into an AI tool, who approves a new tool request, and what happens if someone breaks the policy. That's it. You can expand it later — a policy nobody reads because it's forty pages long protects no one.

STEP 4 — ADD ONE HUMAN CHECKPOINT

Pick the single highest-stakes use of AI in your organization right now — a hiring decision, a customer-facing message, a financial calculation — and require a named person to review the AI's output before it's acted on. One real checkpoint beats ten checkpoints that exist only on paper.

STEP 5 — REVIEW IT EVERY QUARTER

AI tools and the risks around them move faster than most policy review cycles. Put a recurring reminder on the calendar now, before launch-week excitement fades.

A STARTER TEMPLATE TO ADAPT

"[Organization] employees may use the following approved AI tools for [approved use cases]. Customer data, employee personal information, and unreleased financial information may never be entered into an AI tool that has not been reviewed by [approving role]. Any new AI tool must be requested through [process] before use. Outputs used in [high-stakes use case] must be reviewed by [named role] before being acted on. This policy will be reviewed every [quarter/month]."

You will never feel fully ready before writing your first policy. Write it anyway, get feedback, and improve it. That's the job.

CHAPTER 4:
VENDOR RISK PULSE — WHEN AI TOOLS BECOME THE ATTACK SURFACE

REAL INCIDENTS, REAL LESSONS

Third-party risk didn't go away when AI entered the picture — it just found a new door in. That door is increasingly OAuth: the small “continue with Google” or “continue with Microsoft” prompts employees click dozens of times a year without a second thought.

VERCEL VIA CONTEXT.AI

Disclosed June 2026 · Source: DarkReading

An employee signed up for AI productivity tool Context.ai using their corporate Google Workspace login and granted it broad, "allow all" permissions. When attackers compromised Context.ai, they inherited that access and used it to reach environment variables inside Vercel — a company that was not even a Context customer.

NYC HEALTH + HOSPITALS

Disclosed May 2026 · Source: TechCrunch

A months-long intrusion traced back to an undisclosed third-party vendor exposed medical records, government identification, and biometric data for at least 1.8 million people, making it one of the largest healthcare breaches reported in 2026.

THE ONCOLOGY INSTITUTE

Disclosed May 2026

A multi-state cancer care provider confirmed patient data was compromised through a third-party software vendor. Kroll, acting as the vendor's administrator, notified the Institute of unauthorized access to systems holding patient data.

THE PATTERN ACROSS ALL THREE

None of these organizations were breached because their own systems had a flaw. In each case, attackers reached their target by compromising someone the target had chosen to trust. Security researchers tracking this trend describe it simply: attackers increasingly log in, not break in.

WHAT TO DO ABOUT IT

- Move OAuth consent to admin review. Most Google Workspace and Microsoft 365 tenants still let any employee grant a third-party app access to the company account. Change that default.
- Inventory every app with standing access to your identity provider, not just the ones procurement formally approved.
- Treat OAuth tokens as high-value credentials, with the same rotation and monitoring discipline you'd apply to a database password.
- Ask AI vendors directly what scopes their tool requests and why — and decline tools that ask for more access than their function requires.

THE BOTTOM LINE

Your vendor risk program doesn't need a separate "AI risk" track bolted on top of it. It needs the same rigor you already apply to any vendor — extended to a category of tool that employees can now adopt in minutes, with a single click of "continue with Google."

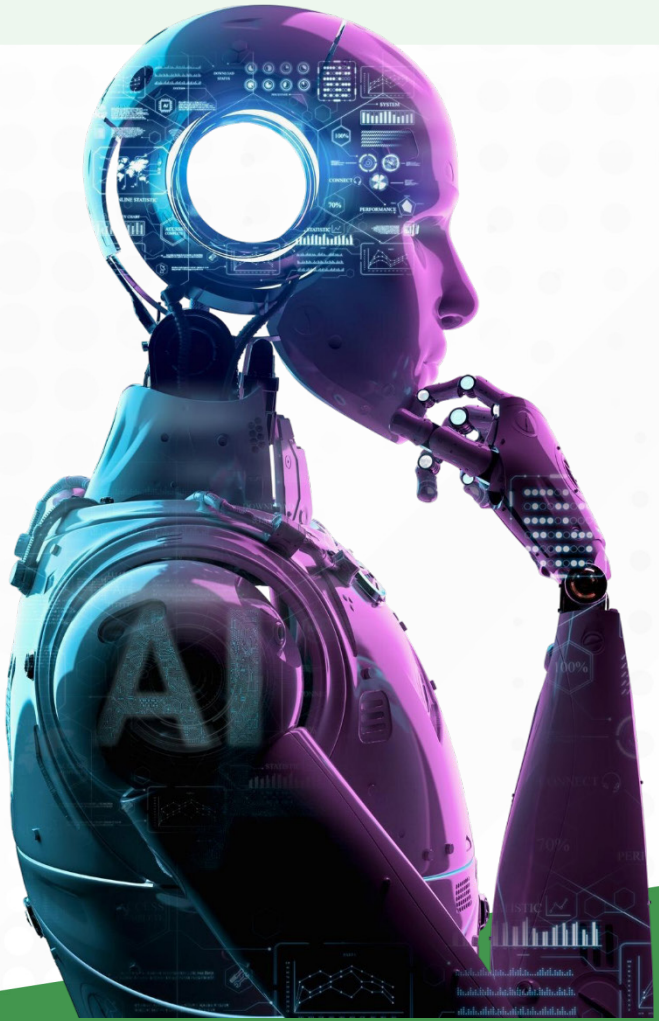

CHAPTER 5:
TOOL OF THE MONTH — DRATA

CONTINUOUS COMPLIANCE, NOW WITH AI BUILT IN

Last edition we featured Vanta. This month we're looking at Drata, another major player in the compliance automation space, and one that's leaned hard into agentic AI features over the past year.

WHAT IS DRATA?

Drata is a security and compliance automation platform that connects to an organization's tech stack — cloud providers, identity systems, code repositories — and continuously monitors whether the underlying controls behind frameworks like SOC 2, ISO 27001, HIPAA, GDPR, and PCI DSS are actually operating as intended. It supports more than twenty frameworks and over two hundred integrations.

WHAT STANDS OUT THIS YEAR

- AI agents for third-party risk. Drata's AI features can draft criteria from existing vendor questionnaires, pull evidence automatically from a vendor's published Trust Center, and complete routine follow-ups — shrinking work that used to take a reviewer hours into minutes.
- Continuous control monitoring. Rather than a point-in-time check, Drata pulls live configuration data from systems like AWS, Azure, and Okta to flag drift the moment a control stops behaving as expected.
- Built-in internal audit workflows. Teams can run internal audits directly inside the platform, with a shared evidence viewer for both internal audit and compliance owners.
- Vendor Trust Center visibility. Reviewers can immediately see whether a vendor maintains its own Trust Center and what assurance documentation is already available, speeding up vendor assessments.

WHAT THIS MEANS FOR GRC PROFESSIONALS

Tools like Drata are quietly redrawing the job description for compliance analysts. The valuable skill is shifting away from manually chasing screenshots and toward reviewing what an AI agent already gathered, catching what it missed, and making the judgment calls a tool can't make on its own.

THE CAVEAT

No automation platform replaces an understanding of the controls it's monitoring. Misconfigured integrations can create false confidence just as easily as they create real assurance. The professionals who get the most value from tools like Drata are the ones who still know what “good” looks like underneath the dashboard.

CHAPTER 6:
CYBER GRC ROLE SPOTLIGHT

ROLE SPOTLIGHT: CYBER GRC ANALYST

If you're trying to figure out whether Cyber GRC is actually the right path for you, here's everything you need to know in one place.

WHAT THE ROLE ACTUALLY IS

A Cyber GRC (Governance, Risk & Compliance) Analyst is the bridge between technical security teams and business leadership. Instead of writing code or configuring firewalls, you assess risk, build policies, track compliance against regulatory frameworks, and translate technical findings into language executives can act on.

A TYPICAL DAY LOOKS LIKE:

- Reviewing and updating a risk register
- Mapping security controls to a framework (NIST, ISO 27001, SOC 2)
- Coordinating with a vendor on a third-party risk questionnaire
- Preparing evidence for an upcoming audit
- Drafting or revising a security policy
- Sitting in on a meeting to explain a risk finding in plain English

CORE SKILLS YOU'LL NEED

- Risk assessment and risk register management
- Familiarity with at least one major framework (NIST CSF, ISO 27001, or SOC 2 — pick one to start)
- Policy writing
- Vendor/third-party risk evaluation
- Communicating technical risk to non-technical stakeholders
- Increasingly: a working understanding of AI governance and AI risk concepts

CERTIFICATIONS WORTH PURSUING (NOT REQUIRED TO START, BUT VALUABLE FOR GROWTH)

- CompTIA Security+ — solid foundational entry point
- CRISC (Certified in Risk and Information Systems Control) — strongly correlated with higher pay
- CISA — valuable once you have a few years in
- IAPP's AIGP — increasingly recognized as the standard credential for AI governance specifically

SALARY EXPECTATIONS (U.S., 2026)

- Entry-level: roughly $63,000–$85,000
- Mid-level: roughly $90,000–$130,000
- Senior/specialized (e.g., Cybersecurity GRC, IT GRC): often $130,000–$170,000+, with top earners surpassing $200,000 depending on industry and location

These numbers vary by city, industry, and certification — insurance, finance, and tech tend to pay above average, and remote roles widen the range further.

JOB TITLES TO SEARCH FOR

GRC Analyst, Cybersecurity GRC Analyst, IT Risk Analyst, Compliance Analyst, Third-Party Risk Analyst, Vendor Risk Analyst, Privacy Officer, Data Protection Officer, Regulatory Compliance Specialist, AI Governance Analyst, AI Risk Analyst.

WHY IT'S A STRONG ENTRY POINT INTO CYBERSECURITY

You don't need a computer science degree, a coding background, or years of IT experience. What you need is the ability to think critically about risk, understand a handful of frameworks, and communicate clearly — skills that transfer directly from careers in audit, compliance, finance, operations, or project management.

THE BOTTOM LINE

Cyber GRC is one of the most accessible entry points into cybersecurity — no coding, no IT background, no computer science degree required. And with AI governance now a core part of the job, the timing has never been better. Skillweed's upcoming AI-infused Cyber GRC class is built specifically around this shift, giving you both the foundational GRC skills employers expect and the AI governance literacy that's quickly becoming a differentiator in this field.

Visit www.skillweed.com for more information and to reserve your spot.

CHAPTER 7:
VENDOR COMPARISON: AI GOVERNANCE TOOLS

CHOOSING THE RIGHT TOOL, WITHOUT THE GUESSWORK

AI is moving into nearly every business function faster than governance programs can keep up — which means GRC professionals are increasingly expected to know not just *what* AI governance is, but *which tools* organizations actually use to operationalize it.

There's no single "best" platform; each one is built around a different starting point, whether that's compliance documentation, bias auditing, or extending a privacy program you already run. Here's a side-by-side look at four of the leading platforms in 2026 to help you understand the landscape before anyone asks you to pick one.

Tool	Best For	Strengths	Limitations
Credo AI	Regulation-mapped compliance documentation	Purpose-built for AI governance (not bolted onto another product); pre-built policy packs for EU AI Act, NIST AI RMF, ISO 42001; strong audit-evidence automation	Built for GRC teams to operate, not developer self-serve; no real-time enforcement at the model layer
Holistic AI	Bias and fairness auditing	Deep technical rigor on bias/fairness measurement; strong for regulated, high-stakes use cases (hiring, lending, healthcare)	Assessment-focused, not enforcement-focused — flags gaps but doesn't control live AI systems
OneTrust AI Governance	Extending an existing GRC/privacy program	Natural extension if you already use OneTrust for GDPR/CCPA; broad integrations; added real-time monitoring and AI agent detection in 2026	An add-on to a broader platform, not purpose-built for AI; less specialized than dedicated AI tools
IBM watsonx.governance	Model lifecycle governance at enterprise scale	Strong fit for organizations already in the IBM ecosystem; solid for hybrid/multi-cloud model documentation and risk monitoring	Most valuable when paired with existing IBM infrastructure; less natural fit outside that stack

HOW TO CHOOSE, AS A BEGINNER:

- Already using a GRC platform for privacy/compliance? OneTrust is usually the path of least friction.
- Need deep regulatory documentation for EU AI Act or NIST AI RMF specifically? Credo AI.
- Working in a bias-sensitive industry (HR, lending, healthcare)? Holistic AI.
- Already inside the IBM ecosystem? watsonx.governance.

A quick rule of thumb the industry repeats often: match the tool to where your organization actually is today, not where you hope to be in two years. Most teams start with one framework and one tool, then expand.

THE BOTTOM LINE:

The right AI governance tool isn't the one with the longest feature list — it's the one that matches where your organization already is. Already running a GRC platform? Extend it. Need EU AI Act-ready documentation fast? Go purpose-built. Operating in a bias-sensitive industry? Prioritize auditing depth. This kind of practical, tool-agnostic judgment is exactly what our upcoming AI-infused Cyber GRC class is designed to build — visit www.skillweed.com to learn more.

CHAPTER 8:
COMMUNITY HIGHLIGHT

INSIDE THE SKILLWEED COMMUNITY

Every edition, we like to step away from frameworks and headlines for a moment and check in on what's actually happening inside the Skillweed community — the part of this publication that's less about the industry and more about the people building a career inside it.

WHAT'S BEEN HAPPENING THIS MONTH

- Planning for the October 3rd Cyber GRC class is officially done and you can now enjoy our updated Cyber GRC and AI class by visiting skillweed.com

- A member of our community who took the last prep class just aced the CRISC certification. This can be you when you sign up for the CRISC On-demand classes. Send an email to info@skillweed.com or visit the website to start preparing.
- The AI Governance Internship continues to give participants hands-on exposure to real frameworks and scenarios, bridging the gap between "I've learned it" and "I've done it" — the kind of proof of application employers are increasingly asking for.
  Send an email to info@skillweed.com to get the link for the previous class and start learning.
- More members are showing up in the community chat to answer each other's questions before an instructor even sees them — a sign the culture is doing what it's meant to do.

AVAILABLE RESOURCES

- Free e-books: Our resource library has grown to 60+ free downloadable guides covering everything from AI governance and LLM security to NIST/ISO frameworks, GDPR, CCPA, IT general controls, third-party risk, and more — including beginner-friendly picks like *50 AI Prompts for Cyber Security GRC* and the *FREE Cyber GRC Toolkit*. Browse and download the full library at skillweed.com/assets/resources.
- Paid books: For members who want a deeper, structured read, our Amazon bestselling titles — *Third-Party Risk Management*, *Risk Assessment*, *Guide to IT General Controls*, and *Cyber Security Operational Technology Best Practice* — are available now. Visit academy.skillweed.com/pages/resource-library to browse and shop.
- YouTube: Prefer to learn by watching? Our YouTube channel breaks down GRC and AI governance concepts in short, digestible videos. Subscribe at youtube.com/@skillweedAcademy.

WHAT PEOPLE ARE SAYING

> “I stopped feeling like I was learning alone the moment I asked my first question in the group chat and got three replies in ten minutes.”

> “I always learn something new from the community events.”

> “The live sessions are what keep me consistent. On-demand alone, I would have quit by week two.”

> “I didn't expect to actually enjoy studying for a certification, but the way it's broken down made it click.”

READY TO BUILD THE SKILLS THIS MAGAZINE JUST WALKED YOU THROUGH?

Join the October 3rd Cyber GRC & AI Governance Class

No coding. No IT background needed. Six weeks, live instructor-led classes, real-world labs, and a curriculum built around exactly what this issue covered — frameworks, risk, compliance, and AI governance.
Reserve your spot: **skillweed.com**

Skillweed — Where this entire magazine, and the community behind it, comes from. From live classes to mentorship to career support, Skillweed exists to take people with zero background and turn them into confident, employable GRC and AI governance professionals.
skillweed.com

MORE FROM SKILLWEED

Geotela — Every place has a story. Geotela is Skillweed's geointelligence navigation tool, built for more than getting from A to B — it helps investors and explorers uncover real opportunity in a region, place by place. Already a community of 500+ and growing, with a waitlist open now.

Join the waitlist by visiting geotela.com

Landzille — Founded by Skillweed's own Akingbade Akinfenwa, Landzille helps families and investors acquire land with confidence across North Texas. With 567+ acres sold across Roxton, Leonard, Gainesville, Bonham, and surrounding communities, it's proof that the same governance and risk thinking that protects companies can help you build wealth too.

Visit Landzille.com to find available projects.

CHAPTER 9:
FINAL MESSAGE
TURNING KNOWLEDGE INTO ACTION

Understanding the trends in this edition is only useful if it leads somewhere. Here's where Skillweed can take you next.

Until Next Month, that's a wrap on the July edition of the Skillweed Cyber Magazine.

If there's one thing this issue should leave you with, it's this: the professionals who win in the next five years of cybersecurity won't be the ones who know the most frameworks. They'll be the ones who can connect risk, compliance, and AI governance into one coherent story — and tell it clearly to the people making the decisions.

That's exactly the skill set we build at Skillweed.

We'll be back next month with more frameworks, more real-world incidents, and more of the practical, no-fluff guidance that's helped thousands of people break into this industry from zero.

Until then — stay curious, stay compliant, and stay ahead.

The Skillweed Team


SkillWeed