

► BUILDING CAPACITY



Darryl Booth, MBA

Building Capacity by Understanding AI Agents

Editor's Note: A need exists within environmental health agencies to increase their capacity to perform in an environment of diminishing resources. With limited resources and increasing demands, we need to seek new approaches to the practice of environmental health. Acutely aware of these challenges, the *Journal* publishes the Building Capacity column to educate, reinforce, and build on successes within the profession using technology to improve efficiency and extend the impact of environmental health agencies.

Column contributors are guest authors. The conclusions of this column are those of the author(s) and do not necessarily represent the views or policies of NEHA, nor does it imply endorsement of any products, services, or resources mentioned.

Darryl Booth has been monitoring regulatory and data tracking needs of environmental public health agencies for over 20 years. He is the manager of U.S. operations for Hedgerow Software.

I had previously imposed an AI cooling-off period for my column. I felt that the bombardment of news, updates, releases, and hype did not need amplification. There is, however, something new. And it is something that I want to introduce: AI agents (aka agentic AI).

The word “agent” signals a new expectation—the ability for some AI products to act on their own on behalf of somebody else. For better or worse, this buzzy concept will wash over us in 2026. AI agents shift from advice to action, which is why leaders should pay attention.

Revisiting AI Basics

My experience when speaking to audiences is that every room is made up of equal parts of people who are either AI skeptics, AI curious, or AI enthusiasts. It is good fun to get all three perspectives in one room, especially

in the locked-down world of environmental health regulators.

For the skeptics and the curious, the least you need to recall is that AI is the ability of AI services to take a prompt (i.e., a written request, a question, an instruction) and generate a response that seems insightful, useful, and relevant.

For our enthusiasts, we have invested a lot of time aligning this tool to real-world recipes suitable for environmental health professionals in an office setting. The examples below are truly useful.

Recipe 1: Summarize a Complex Document for a Certain Reader

Take a complex regulation, enter it into an AI chat window, and ask for a summary, outline, or translation. This example is very useful and appropriate.

Recipe 2: Create a Repository of Knowledge

Here we use Google's NotebookLM, preload it with PDFs (e.g., the Food and Drug Administration *Food Code* and appendices, standard comments, guidance documents, local ordinances, and forms), and then ask it questions about those resources.

Recipe 3: Generate a PowerPoint on Any Topic

Using PowerPoint's built-in Copilot, click Create with Copilot and give it a topic such as, “Give me a 20-minute presentation explaining the function of the environmental health department.” Give it a little time, and you will have a credible presentation.

Recipe 4: Start a Letter, Email, or Flyer

Many writers struggle to just get started. Ask any AI tool to generate a letter, paragraph, or outline to get you started.

If the recipes above are valid, think of three more, and you are on your way to helping lead your organization to a new way of working. I acknowledge important overriding matters such as agency policy, attribution, hallucinations, and other AI issues. Conversations around these matters are crucial. Do not skip steps.

Agentic AI

As mentioned, the words “agent” and “agentic” are working their way into the hype cycle. Why? Because the previous examples are all strictly interactive. They save time and ideally improve quality, but they still require an orchestrator (you) to make the

request, assess and tune the result, and then bring it to its final form (e.g., send an email, make a presentation, meet with a graphic designer). It is a copy/paste exercise before you are really done.

AI pundits are seeking new, increasingly autonomous actions. Why? To increase their value, visibility, and utility. Let us look at some examples to make this point clear.

Agentic AI Example

Consider the following example:

Read each unread email in my inbox, based on what you know about me and my job, move junk into the Archive folder, and accept appointments that A) are consistent with my priorities and B) fit my calendar. Send confirming responses where needed, and flag emails that require my action as high priority.

In this example, we are not just asking for a text response. We are asking for it to act—and that is the important difference. It promises a “set it and forget it” utility.

This autonomous agent requires a new trusted relationship on a few levels:

1. Can we trust the underlying AI model to understand our instructions?
2. Can we trust the underlying AI model to understand the content (and intent) of emails?
3. Can we trust the underlying AI agent to act on our behalf (where an error—even an honest error—impacts our reputation and credibility)?

The third question is concerning. To my thinking, what happens out of our view, but in our name, deserves oversight akin to having an intern manage your inbox. A competent intern would gain our trust over time, but probably not initially.

Another Agentic Experiment

I have another demonstration where I open one of several new agentic browsers, such as Comet or Alias. Agentic browsers can see

AND click what appears in your browser window to meet a task. I launched our data management system and asked in the sidebar: “Using this software that is used to manage facilities, permits, and inspections for regulated facilities, go get a list of active facilities permitted in the Air Quality Program and export to CSV.”

The experience caught my attention. I watched the AI model read the screen, clicking links, and even backtracking until the task was done. In the browser sidebar, I could monitor the step-by-step reasoning as it typed out things like, “I clicked the facilities icon and I can see a filterable grid,” and “Let’s put ‘Air Quality’ into the filter.”

Several observations are worth noting. First, it sometimes clicked the wrong link and had to backtrack and recover. It took longer than it would have taken me. As it reasoned, it took maybe five times longer than a trained user. Note also that it sometimes took different paths to accomplish the task.

In my talk, I aligned the example to those morning activities, such as reviewing the complaints that came in overnight, checking that timesheets were entered by the end of the week, running manager reports, and posting inspection results to the internet.

What if 10 independent agents started their tasks at 6 a.m., such that when the humans come in at 8 a.m., the complaints are already sorted and assigned, along with other tasks?

An Agentic AI for Environmental Health Professionals?

Not today and maybe not in 2026, but sooner than many of us expect. Agents and agent-like features will start showing up in the tools we use every day, arriving in small increments and building toward a greater number of AI-driven tasks and services.

For environmental health programs, we can ignore the hype and should not be fearful. A crawl-walk-run approach might look like this:

- **Crawl:** Use AI to explain and summarize.
- **Walk:** Allow AI to prepare actions (e.g., build a report, assemble a spreadsheet, draft replies), but require a human to edit and approve. Give attribution to AI, if appropriate.
- **Run:** Only after trust is earned, allow limited “hands-on-keyboard” automation within a sandbox and only with narrow permissions. Require audit logs and easy rollback.

In other words, treat agents the way you would treat a new intern with system access. Start small. Supervise. Expand responsibilities when proven reliable.

If the past 3 years were about learning to prompt, the next few years will be about learning to delegate safely. And, as regulators, educators, and leaders, delegation cannot undermine the trust you have built. 🌸

Corresponding Author: Darryl Booth, U.S. Operations Manager, Hedgerow Software.
Email: darrylbooth@hedgerowsoftware.com

CALL FOR SUBMISSIONS

The *Journal* seeks guest authors for the Building Capacity column. Our goal is to provide a platform to share capacity building successes occurring across the country and within different sectors of the environmental health profession, including academia, private industry, and state, local, tribal, and territorial health agencies. Submissions will be reviewed by the NEHA technical advisors for data and technology and *Journal* staff for appropriate content, relevance, and adherence to submission guidelines. To learn more about the submission process and guidelines, please visit www.neha.org/contributors.

Did You Know?

Our Government Affairs program advocates for support of environmental health programs and professionals at federal, state, and local levels of government. We function as a liaison between environmental health professionals and government officials to inform decisions that support and fund our workforce. You can stay up to date on our work at www.neha.org/advocacy.