HFS

HFS Spring Summit 2025

THE RISE OF SERVICES-AS-SOFTWARE

THE AGENTIFICATION OF EVERYTHING

Dana Daher
Practice Leader, HFS Research

When was the last time you used a real map?

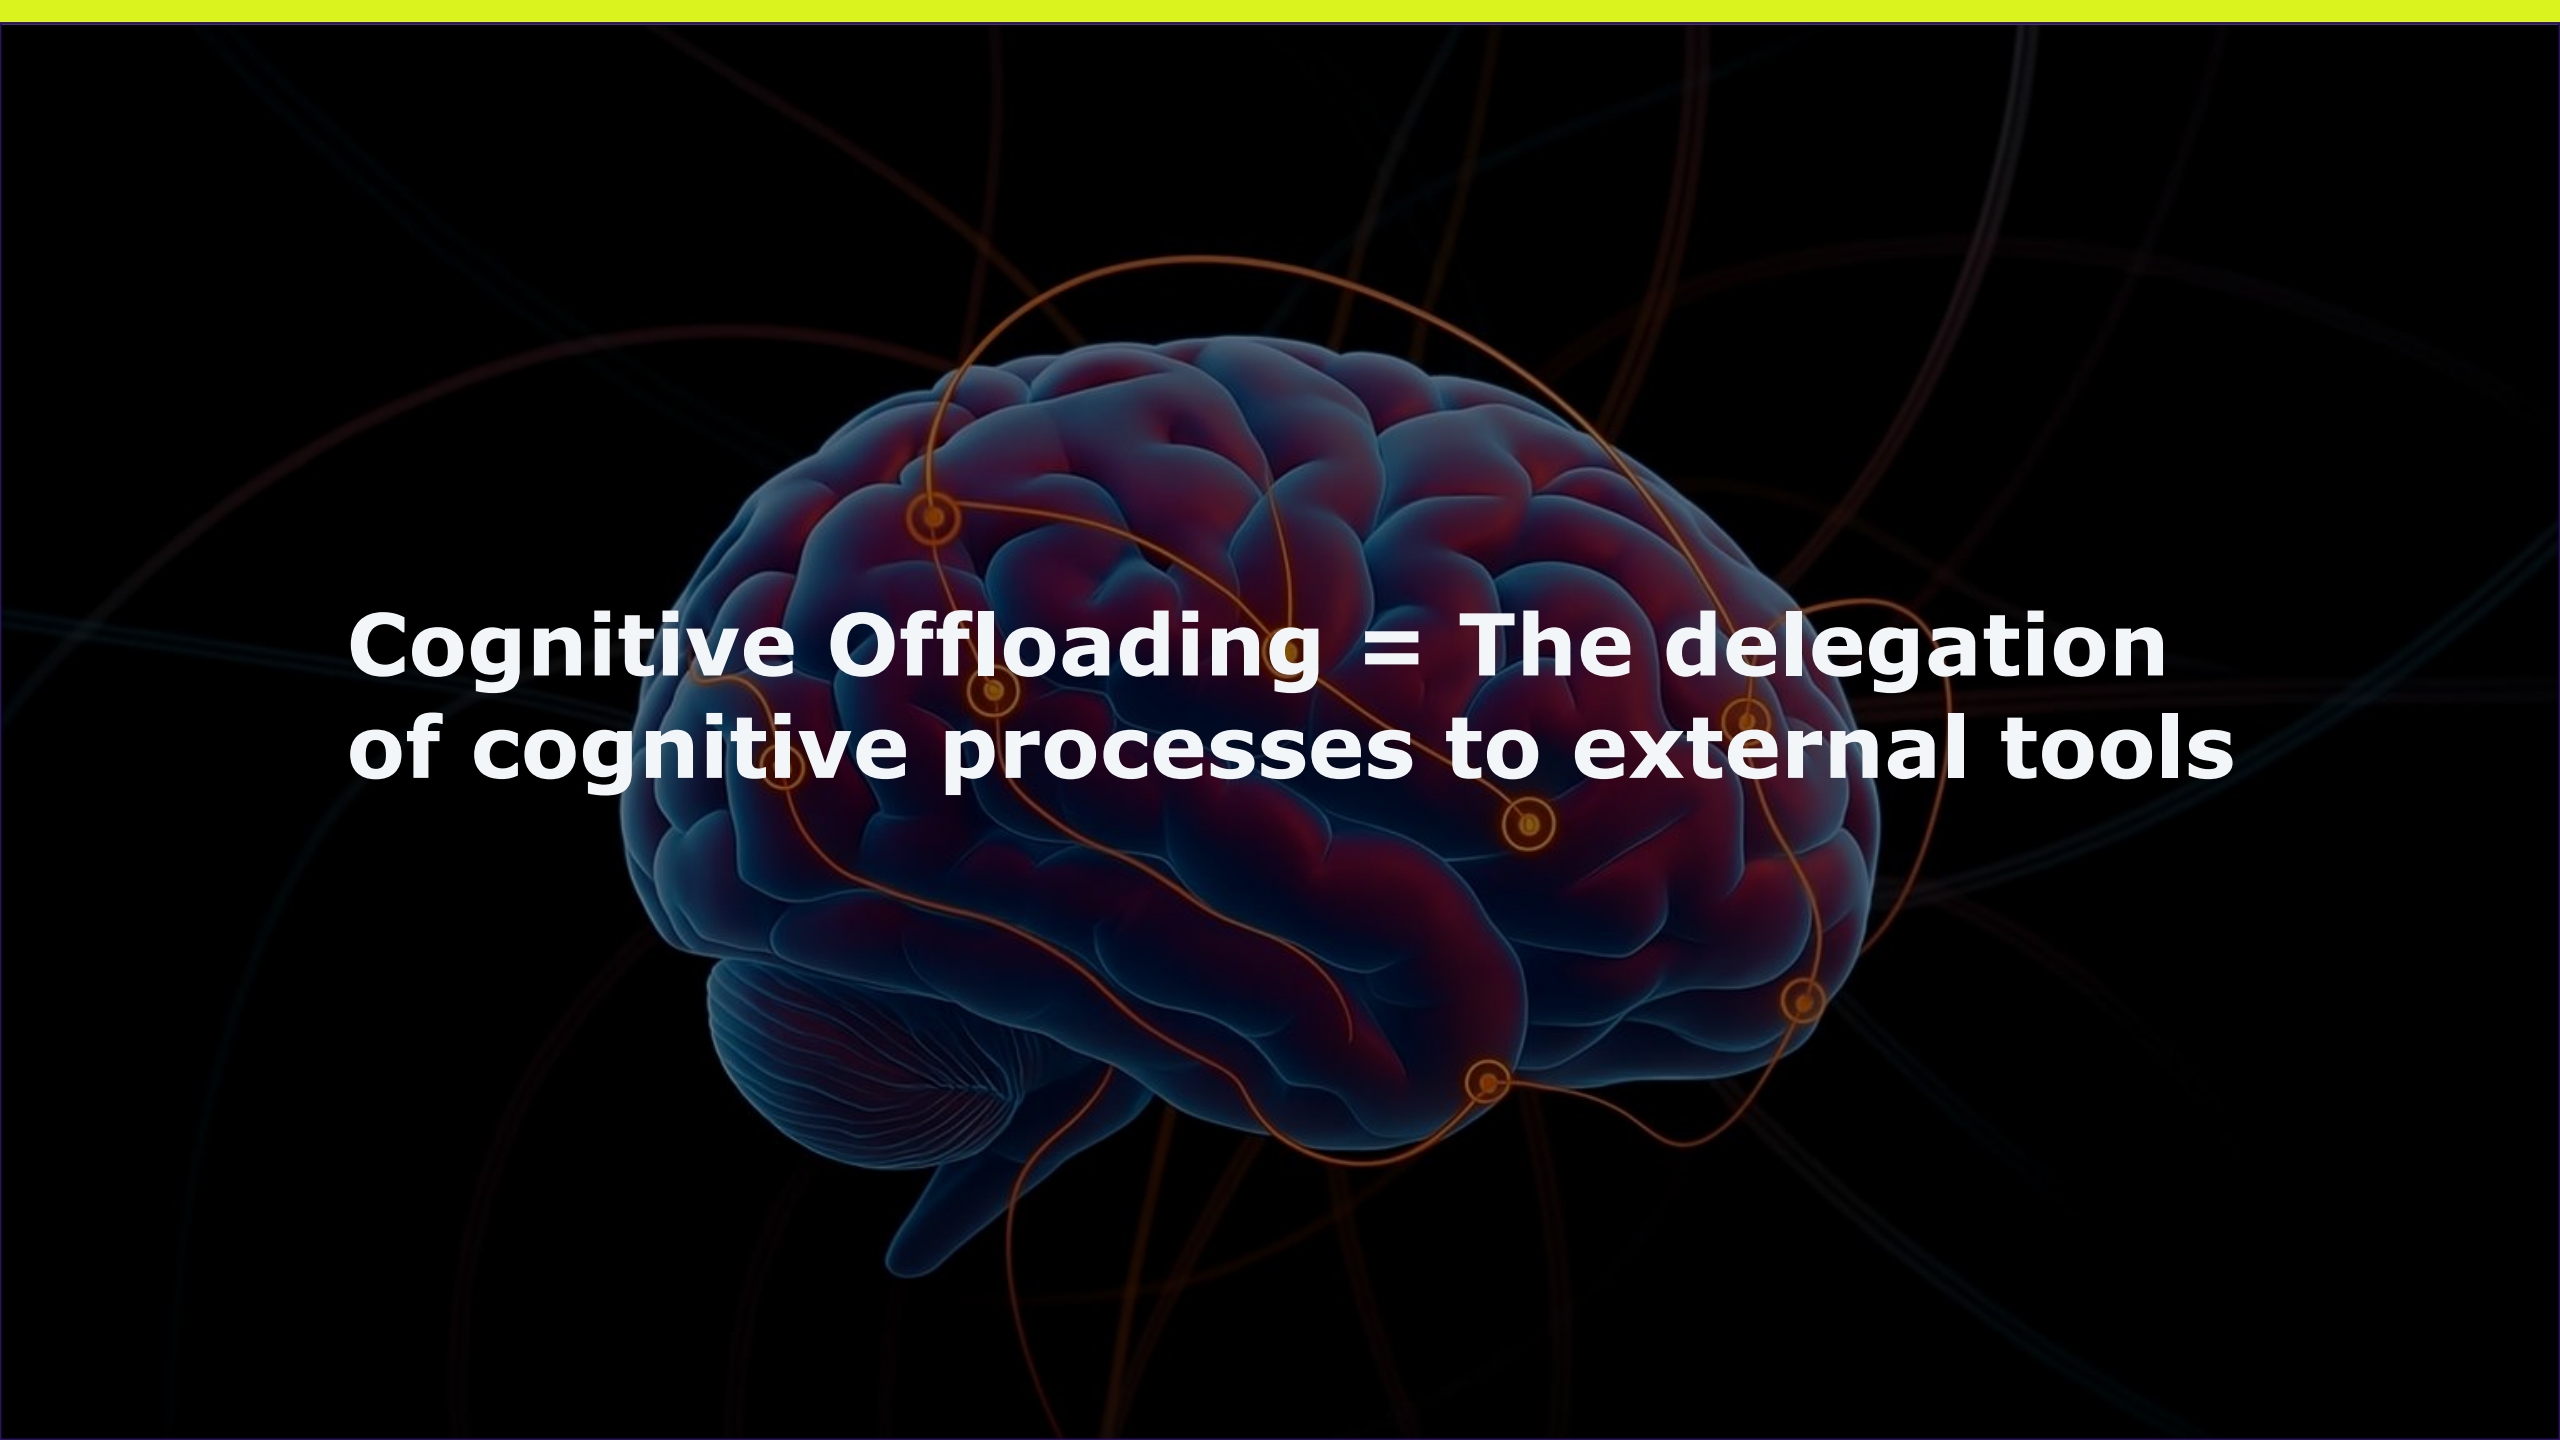
Cognitive Offloading = The delegation of cognitive processes to external tools

We’ve always offloaded — memory to paper, stories to tablets, math to machines, navigation to apps.

That’s not new.
That’s human progress.

The Three Waves of Cognitive Change

Wave 1: The Age of Access

(2007-2012)

Technologies
Google, iPhone, Wikipedia

What Changed:
Knowledge retrieval replaced memorization.

Impact

Rise of "Why remember when you can Google it?" thinking.

Type of offloading

Memory offloading

Wave 2: Age of Conformity

(2012-2023)

Technologies
Facebook, YouTube, TikTok, Twitter

What Changed:
Reality curated by algorithms; emotional over truth

Impact

- Echo chambers & polarization
- Reduced independent critical thinking
- Unseen algorithms shape opinions

Type of offloading

Opinion offloading

Wave 3: Age of Complacency

(2023+)

Technologies
ChatGPT, Claude, Copilot

What Changed:
Entire workflows offloaded — from research to writing

Impact

- Rapid output, shallow engagement
- Diminished confidence in judgment
- Flattened expertise development

Type of offloading

Decision offloading

Source: HFS Research, 2025

Enter GenAI – A new kind of offload

This isn't just a new tool. GenAI is the first technology that *mirrors how we think*.
It doesn't calculate or recall—it *predicts*, *composes*, and *decides*.
And that changes everything.

- GenAI isn't just a helper… It *simulates cognition*
- It drafts your ideas *before* you've fully formed them
- It speaks with fluency and confidence, even when it's wrong
- It *feels* like thinking, which makes it *easy to stop thinking*
- It makes offloading effortless—and invisible

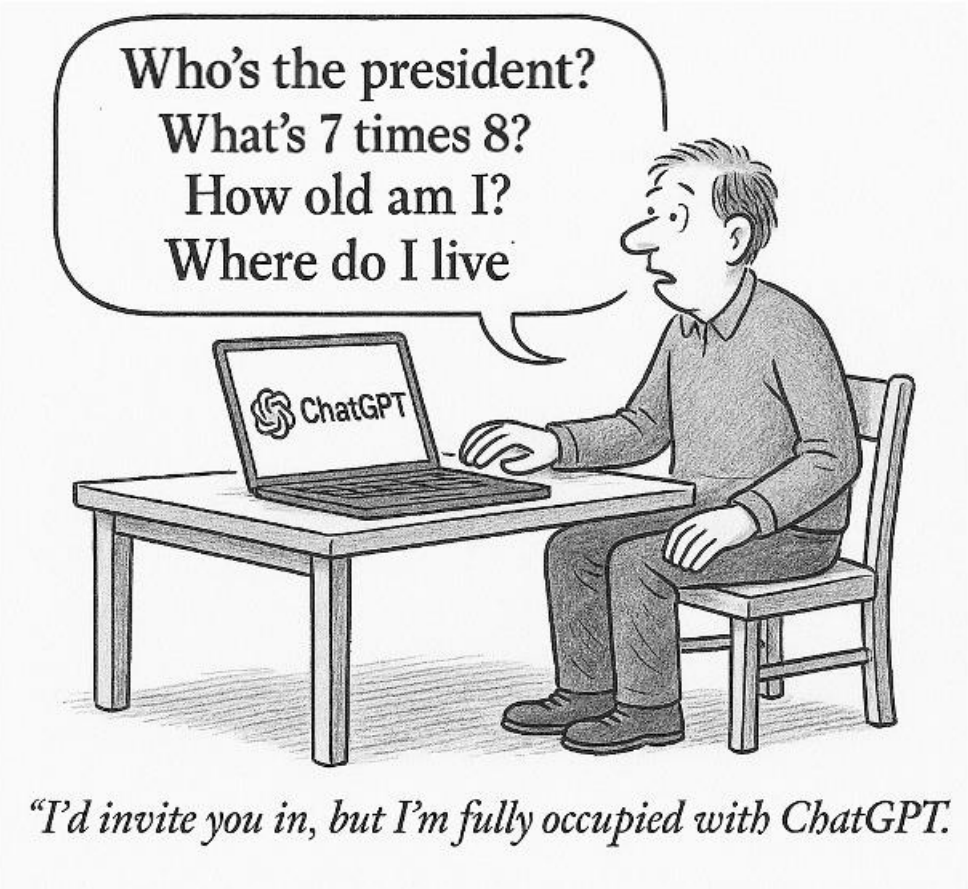

Why it’s easier to surrender than to think

- GenAI feels fluent, even when it’s wrong
- We’re rewarded for output, not insight
- Most workflows don’t value reflection
- It’s easier to copy than to challenge
- Pressure to go fast, scale, and appear smart

What early research is telling us about GenAI and thinking

- **GenAI shifts thinking from doing to supervising.**
- **Perceived "good enough" output reduces critical engagement.**
- **Trust in GenAI leads to overreliance.**
- **Frequent AI use can erode confidence and motivation.**
- **Students and professionals bypass the learning process.**

Generative AI Can Harm Learning

Hamsa Bastani,[1*] Osbert Bastani,[2*] Alp Sungu,[1*†]
Haosen Ge,[3] Özge Kabakcı,[4] Rei Mariman

[1]Operations, Information and Decisions, University of Pennsylvania
[2]Computer and Information Science, University of Pennsylvania
[3]Wharton AI & Analytics, University of Pennsylvania
[4]Budapest British International School

buted equally.
alpsungu@wharton.upenn.edu.

revolutionize how humans
ignificantly improving hu-
on is how generative AI af-
kills as they perform tasks.
oductivity gains, especially
han experts must check its
pecifically OpenAI's GPT-
ses at a high school. In a
nts, we have deployed and
standard ChatGPT inter-
gned to safeguard learning
15% of the curriculum in
our results show that ac-

The Impact of Generative AI on Critical Thinking: Self-Reported Reductions in Cognitive Effort and Confidence Effects From a Survey of Knowledge Workers

Hao-Ping (Hank) Lee
Carnegie Mellon University
Pittsburgh, Pennsylvania, USA
haopingl@cs.cmu.edu

Advait Sarkar
Microsoft Research
Cambridge, United Kingdom
advait@microsoft.com

Lev Tankelevitch
Microsoft Research
Cambridge, United Kingdom
levt@microsoft.com

Ian Drosos
Microsoft Research
Cambridge, United Kingdom
t-iandrosos@microsoft.com

Sean Rintel
Microsoft Research
Cambridge, United Kingdom
serintel@microsoft.com

Richard Banks
Microsoft Research Cambridge
Cambridge, United Kingdom
rbanks@microsoft.com

Nicholas Wilson
Microsoft Research
Cambridge, United Kingdom
niwilson@microsoft.com

Abstract

The rise of Generative AI (GenAI) in knowledge workflows raises questions about its impact on critical thinking skills and practices. We survey 319 knowledge workers to investigate 1) when and how they perceive the enaction of critical thinking when using GenAI, and 2) when and why GenAI affects their effort to do so. Participants shared 936 first-hand examples of using GenAI in work tasks. Quantitatively, when considering both task- and user-specific factors, a user's task-specific self-confidence and confidence in GenAI are predictive of whether critical thinking is enacted and the effort of doing so in GenAI-assisted tasks. Specifically, higher confidence in GenAI is associated with less critical thinking, while higher self-confidence is associated with more critical thinking. Qualitatively, GenAI shifts the nature of critical thinking toward information verification, response integration, and task stewardship. Our insights reveal new design challenges and opportunities for developing GenAI tools for knowledge work.

CCS Concepts

• **Human-centered computing** → **Empirical studies in HCI**.

Keywords

Critical thinking, Generative AI tools, Knowledge worker, Bloom's taxonomy, Survey

ACM Reference Format:

Hao-Ping (Hank) Lee, Advait Sarkar, Lev Tankelevitch, Ian Drosos, Sean Rintel, Richard Banks, and Nicholas Wilson. 2025. The Impact of Generative AI on Critical Thinking: Self-Reported Reductions in Cognitive Effort and Confidence Effects From a Survey of Knowledge Workers. In *CHI Conference on Human Factors in Computing Systems (CHI '25), April 26–May 01, 2025, Yokohama, Japan.* ACM, New York, NY, USA, 23 pages. https://doi.org/10.1145/3706598.3713778

ACM ISBN 979-8-4007-1394-1/25/04
https://doi.org/10.1145/3706598.3713778

1 Introduction

Generative AI (**GenAI**) tools, defined as any *"end user tool [...] whose technical implementation includes a generative model based on deep learning"*,[1] are the latest in a long line of technologies that raise questions about their impact on the quality of human thought, a line that includes writing (objected to by Socrates), printing (objected to by Trithemius), calculators (objected to by teachers of arithmetic), and the Internet.

Such consternation is not unfounded. Used improperly, technologies can and do result in the deterioration of cognitive faculties that ought to be preserved. As Bainbridge [7] noted, a key irony of automation is that by mechanising routine tasks and leaving exception-handling to the human user, you deprive the user of the routine opportunities to practice their judgement and strengthen their cognitive musculature, leaving them atrophied and unprepared when the exceptions do arise.

In response, research has begun looking closely at how different activities are impacted by GenAI and the extent to which cognitive offloading [8] occurs, and whether this may be an undesirable thing. Some work has focused, for instance, on studying the effects of GenAI use on memory (e.g., [1, 106]) and on creativity (e.g., [28, 100]). Moreover, design research has also been developing interventions that *improve* the ability of people to think in certain ways (e.g., [24]). We review these lines of work in Section 2.

In this paper, we focus on a higher-level concept that captures another aspect of thought considered desirable and worthy of preservation: *critical thinking* (defined in Section 2). The effect of the use

[1]While there is no broad consensus on how to define this now-common term, for clarity we adopt this definition, a rationale for which is given in [115].

What we risk losing when we stop thinking

Erosion of Judgment

When we stop evaluating, we stop owning decisions

Erosion of Quality

When "fast enough" replaces "thoughtful enough," standards collapse quietly.

Erosion of Confidence

When we delegate too much, we lose belief in our ability to reason.

Erosion of Originality

When GenAI writes everything, everything starts to sound the same.

Erosion of Expertise

When we skip the struggle, we weaken our depth.

Erosion of Diversity

When we optimize for speed, we lose voices, perspectives, and nuance.

The power of GenAI when used well

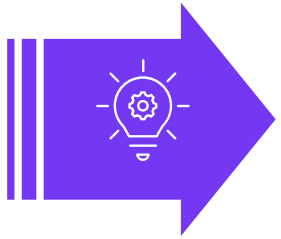

Accelerates Insight

Surface patterns summarize noise and synthesize fast.

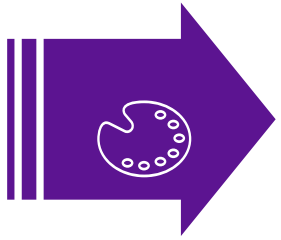

Unlocks Creativity

Unsticks ideation, generates options, and reframes problems.

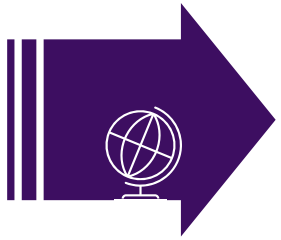

Improves Access

Bridges language, supports neurodivergence, and democratizes content

Q: What are the primary drivers behind your organization's investment in GenAI initiatives?

Rank	Drivers of GenAI investments
1	**Enhancing productivity and efficiency**
2	**Increasing organizational agility and responsiveness**
3	**Improving customer or user experience**
4	Gaining a competitive advantage
5	Opportunity to reinvent or enhance services and products
6	Fear of missing out on emerging technologies
7	Strategic mandate from board members
8	Improving data-driven decision-making and insights
9	Enhancing data security and privacy
10	Increasing revenue
11	Expanding into new markets or demographics
12	Pressure from employees to adopt

Sample: 553 survey participants
Source: HFS Research, 2024

Are you offloading — or surrendering?

❌ Surrender to AI	OR	✅ Partner with AI
Let AI make decisions *for* you	⟷	Use AI to enhance judgment
Accept outputs as “good enough”	⟷	Verify, question, and revise
Skip the thinking process	⟷	Stay in the cognitive loop
Automate expertise you haven’t developed	⟷	Offload grunt work, not insight
Treat AI as a thinking substitute	⟷	Treat AI as a creative partner

Should I use GenAI here? Ask yourself this first.

1. Am I avoiding effort or amplifying insight?
2. Do I know what “good” looks like?
3. What’s the consequence if this is wrong?
4. Is the output final or foundational?
5. Can I explain or defend the result without GenAI?
6. Am I building understanding or just moving faster?

Once you use GenAI, use it with intent.

01 **Interrogate the output.**
Don’t accept answers at face value; ask, “Does this actually make sense? Is this true?”

02 **Do the first draft yourself.**
Force your brain to think before you prompt. Even five minutes builds muscle.

03 **Switch from “what should I do?” to “what am I missing?”**
Use AI to expand your thinking, not replace it.

04 **Test yourself without the tool.**
Rebuild confidence by solving things manually once in a while.

05 **Reflect on the thinking you didn’t do.**
After you use GenAI, ask: What part of this did I actually think through? What did I skip?


Don’t offload your mind.

Thank you.

THE RISE OF
SERVICES-AS-SOFTWARE
THE AGENTIFICATION OF EVERYTHING

www.hfsresearch.com

hfsresearch

www.horsesforsources.com

www.horsesmouthpodcast.com