



■ LUNCH & LEARN • BRINGING AI ABOARD

Towing the Fleet Forward

Bringing AI and LLMs aboard, under governance, starting now.





Welcome aboard

Who’s talking, why this course exists, and where the next thirty minutes go.

YOUR PRESENTER


Zach Geller

Executive IT Support • N. Texas Principal Lead — Business Systems & Transformation, RTX

THE FOUNDATION

ITILv4 (Axelos) • Train-the-Trainer (ATD) • Architecture & Design (CMU) • Strategy & Innovation (Stanford) • Agentic AI (MIT — in progress)

HER MAJESTY • OFF THE CLOCK



Ms. Lil Buddy — gargoyle gecko. I am staff; the second print is the kingdom I maintain. Also: *top-10 Magic: The Gathering* worldwide, *Go Birds*, and a gardener.

THE COURSE, IN FOUR LEGS



- 1 Who I am — and why we’re here**
Two minutes of introductions, then we sail.
- 2 The what & why behind AI**
The talk: the paradox, the landscape, the play, execution — about twenty minutes.
- 3 The how — a live demo**
Ten minutes in our governed tool, worksheet in hand. Today I drive; this week, you do.
- 4 Q&A — your questions, your start**
Reactions, pushback, and where you’d run your first tow.

WHY THIS COURSE


The MIT coursework is where “towing” was born — and it’s the dawn of this technology. **Pioneering happens now; it should happen here.**

■ THE HOOK


The generation-ship paradox

“There’s always a faster ship — and a faster model.

Perfect timing doesn’t exist.”



The voyage — A crossing of centuries — generations live and die aboard.



The overtake — Launched decades later, engines invented after you left.



The trap — There’s always a faster ship. You never launch.

Every departure date is wrong — and so is waiting. That’s the paradox.

Now swap ships for models. Every few months a faster, cheaper, smarter model launches. Wait for the “definitive” one and you never leave port.

That’s why we’re in this room: not to pick the perfect model — to know the landscape well enough to sail in spite of the paradox.

■ RESOLVING THE PARADOX



We can't out-wait it, so we tow

The trap holds only if hull and engine are welded together. **Towing splits them.**



- 1 **Keep the hull** — our people, operations, knowledge — the vessel we've built.
- 2 **Swap the engine** — hitch to the newest; never bet the fleet on one.
- 3 **The flip** — the faster ship becomes our next engine — waiting loses.

■ WE TOW WHEN — A TRIGGER, NOT A DATE

$$\text{value}_{\text{now}} > \text{friction}_{\text{switch later}}$$

where **value** = what starting today earns us · **friction** = the cost of re-hitching when a better engine arrives

Capability clears the bar · governance is ready · cost beats manual.

(A framework built around the uncertainty itself — the way a qubit harnesses it.)

The paradox never closes — towing means it doesn't have to. **We clear all three triggers today, so we start. That's the case for this class.**



AI didn't arrive, it accelerated

So what is AI? Just a new program — but it matured so far past everyone's expectations that a generational shift came with it.



WE ARE HERE — governed & compliant LLMs. Options in this class — analogous to our use case, swappable as things change:

Copilot (gov cloud)

ChatGPT (enterprise)

IBM watsonx

THE FRONTIER. AI agents and deep self-reasoning — Claude, GPT-5, Gemini. Available now, just not in our environment yet. That's the next tow.

**internal estimate of typical compliance-cloud feature lag — the published open-vs-frontier gap is ~4 months (Epoch AI), and the trail keeps shortening.*

The lesson: every leap was underestimated — early (not perfect) movers kept a lead the laggards never closed. **And we're already on the road:** RTX putting governed AI in our hands. The compounding has started.



Two AIs on the table, and a gap that’s closing

Governed enterprise AI

OUR LANE — EXAMPLES OF THE CLASS, NOT COMMITMENTS

- ✓ Compliant by default — runs inside a governed compliance boundary
- ✓ Our data never leaves the boundary
- \$ Seat-licensed — predictable cost, already in our stack
- *Constrained, for now*

Copilot (gov cloud)

ChatGPT wrappers

IBM watsonx

Frontier models

AVAILABLE NOW — NOT IN OUR ENVIRONMENT

- ✓ Fast and highly capable
- ✓ State-of-the-art reasoning
- \$ Metered tokens — pay-per-use, easy to start, easy to overrun
- *External and unbounded*

Claude

GPT-5

Gemini




The gap is closing. Private cloud can bring whole frontier models inside governed walls — or split them down (distillation, SLMs) when it fits — on whichever compliant platform we ride. The “safe vs. powerful” trade-off is dissolving.

Where this goes

Today is AI 101. Agentic AI — models that act, not just answer — is the 201, and it's coming (even cost-effectively).

101 — today



today

the next 12-24 months

then: everywhere

201 — agents

01 Frontier power, in-house

Run the whole model in-house for integration and the heaviest loads — or split it down (SLMs, distillation) when the job fits.

02 From assistants to agents

AI that executes multi-step work — not just answers questions. The 201 arrives mid-course.

03 The default collaborator

AI embedded in every tool and workflow, for everyone — as ordinary as a browser tab.

The lesson: the moat won't be access to models — everyone gets that. It'll be readiness: clean data, mature governance, fluent people. That's what we start building now.



Tow the fleet

Elevate what already works, no rip-and-replace. Three lines, one cable — **and today we pull hard on exactly one of them.**

01 Software


Wrap legacy systems in AI layers. Modernize the intelligence, keep the proven core.

02 Operations

Accelerate proven workflows inside policy-compliant guardrails. Speed within governance.

03 People

Arm our people rather than replace them. This is where the real change happens — and it's the only line we pull today.




Governance is the cable's tension, not a brake on it. *Software and Ops are already rigged — they're the next pull, not the forgotten ones.*



A shared language is on us

Our AI-native recent grads and their managers can't be speaking two different languages. Closing that gap isn't theirs to solve alone — we own the outcomes, so we own the fluency.



- To the business — the outcomes are ours to deliver
- To new hires — we arm them, we don't strand them
- To each other — and to ourselves, as peers

The twist: in the space that actually ships our product — compliance, safety, customer trust — **we're the newer, faster ship.** A fancy AI assistant on the old vessel doesn't know these waters.

"The language gap is a leadership responsibility, not a generational excuse."




The tow line runs both ways

Two tows, one crew — and AI translating between the decks.

RECENT GRADS TOW ESTABLISHED LEADERS the HOW — tool fluency • execution speed • AI instincts



ESTABLISHED LEADERS TOW RECENT GRADS the WHAT & WHY — compliance • safety • customer trust • judgment



And it isn't clean by generation. Plenty of established leaders are the fluent ones here — and plenty of recent grads arrive skeptical of AI itself: the data centers, the energy, the jobs. That's a different conversation, and it's a fair one.


Either direction, you tow the slower ship into the new speed — nobody gets left behind, and nobody gets cut.



Small cost now. Compounding return.

Drawn to scale — and the advantage it buys is the part you can't purchase later.

THE RETURN — IT COMPOUNDS, QUARTER OVER QUARTER



The model was never the moat.

- Your data — organized so AI can actually use it
- Your workflows — rebuilt around the new speed
- Your people — fluent from thousands of reps

value_{now} > **friction** of switching later

The trigger from earlier — and it's already cleared.

HOW IT COMPOUNDS — THE FLYWHEEL, AND EVERY LAP ADDS TO THE STACK




You can buy the model. You can't buy two years of your team having used it. So the expensive choice isn't starting — it's waiting. **That's why I'm here.**



Governance is what lets us go fast

The part you asked me to focus on — and the part that actually lets us say yes.



1 What goes in

Know your data class before you paste. Approved data, approved tools — all inside our governed environment. If it couldn't leave the boundary in an email, it doesn't leave in a prompt.

2 Where it runs

Governed tools only. Sensitive information never touches open, consumer models — the boundary is what keeps our clearances and contracts intact.

3 What comes out

A person owns every output. Verify before it ships — the work carries your name, not the AI's. Rigor on the workflow, always.

WHEN IN DOUBT:

Ask before you paste

Default to the governed tool

Verify, then send


Guardrails don't slow the tow — they make the speed safe, and safe speed is what lets us commit.



Start with a pilot you can measure

Not a transformation program — one first launch, with a course you can chart.

launch pad — today



under way — week 1

the first measured win — day 30



1 Pick one routine task

WEEK 1

Low-risk, easy to double-check, done every week — a task where a win is obvious.

status emails • meeting notes • report digests

2 Set the measure up front

BEFORE LAUNCH

Before you launch, write down the one number that will prove the win:

hours saved

cycle time

error rate

3 Prove it, then share it

DAY 30

Run it weekly, verify every output, and say the number out loud at day 30.

the win funds tow #2 →

It's not only about ROI. The first pilot is how a team gets comfortable, breaks the uncertainty wall, and starts the compounding now.



Choosing a model: capability, cost, and fit

It's like buying a car — same decisions, same trade-offs. **Analogy only, and security-agnostic:** the car says nothing about which roads you're cleared to drive.

BUYING A CAR		PICKING A MODEL
Make	Toyota · Ford · Honda	Anthropic · OpenAI · Google
Model	Camry · F-150 · Civic	Claude · ChatGPT · Gemini
Trim	XL → XLT → Platinum	Haiku → Sonnet → Opus → Fable Luna → Terra → Sol
Drive mode	Eco → Comfort → Sport	quick answer → default → extended thinking
Fuel	gas, burned per mile	tokens, metered per use
Octane	87 regular → 91 → 93 premium	training & weights: fast → balanced → pro
Fit	how it drives	how it thinks & explains

TWO ROADS, ONE GARAGE

guardrails, speed limits — and everyone gets home

blistering laps, no guardrails — not where we commute

“Claude explains what it’s doing and why, in real time — so I learn while it works.”

Capability · Cost · Fit — pick the model that fits the job and the person using it. Governance decides the boundary; the analogy never does.



How I use AI

*“My mind races faster than my skills — **until now.**”* Three things in this room were made that way. Here’s the honest split.

■ THIS DECK

27 slides and a worksheet

AI DID

Every layout, every draft, twelve rounds of revision — in a few evenings.

I DID

The argument, the running order, and twelve rounds of “no, do it again.”

■ THE ELEVATOR AD

a 17-second animation

AI DID

Drew and rendered every frame from a description. No stock, no template.

I DID

The joke, the timing, and the call to cut it before it got cute.

■ THE TOKEN CALCULATOR

a small working tool

AI DID

Wrote it, debugged it, and explained it back to me in one sitting.

I DID

Knowing what question it had to answer — and checking that it did.

Rigor: iterate ~5×, meta-prompt to sharpen, verify every output — speed without slop.

Habits: keep a thought-board — jot ideas while big tasks run, hit enter when usage is back; and don’t sweat typos, models read right through them.

Start here: the email you already send every week. That is the entire on-ramp — nothing on this slide started bigger than that.

■ THE QUALITY GUARDRAIL



Slop is real.

AI is a probabilistic yes-man: it predicts the next likely word, so it hands you confident, wrong, unverified output. It looks real, and it won't be flawless for years.

■ OUR JOB NOW



Hallucinations look exactly like good work:

- A source or citation that was never written
- A number that's confidently, subtly wrong
- A summary that quietly invents a detail

Without governance and your judgment, these slip through, **and “AI-assisted” becomes slop with your name on it.**

“We'll adopt once it's perfect” is the perfect-ship trap again.

There's always a cleaner version coming. Govern the flaws. Tow anyway.



Your part starts this week

Three days. Three moves. Every prompt below is printed on your worksheet, word for word — type it exactly, then swap one noun for your own work.

MONDAY

Pick it, then draft it

One routine thing you already do every week. Hand it over and read the draft — don't touch it.

■ TYPE THIS

Draft the weekly status email I send my director: what shipped, what slipped, what I need, what's next.

WEDNESDAY

Critique — don't edit

Say what's wrong in your own words and make it go again. Three rounds. No hand-polishing.

■ TYPE THIS

Too corporate and too long. Cut it by half. You buried the ask — put it in the first two lines.

FRIDAY

Meta-prompt — the one nobody tries

Make it improve your prompt, then run its version. This is the move that changes how you work.

■ TYPE THIS


Rewrite my original prompt so it gets a better first draft on the first try - then run it.

Keep the rigor. Verify, think critically, own the result. That's how we get the speed without the slop.



Governance keeps it safe. Practice makes it ours.

We don't wait for the perfect ship. We tow the one we've built: our people, our operations, our knowledge, into a future that keeps compounding. **The rest is just doing the work.**





The elevator slide-deck ad · *the 2026 director's cut*



2016 — one deck, one elevator ride... almost.

THE SAME TASK, TEN YEARS APART

2016

One deck. One elevator ride. Still at 16% when the doors open.

2026

Seventeen seconds of finished animation, drawn from a sentence.

What didn't change: someone still had to decide it was funny, and decide where to cut.

17 seconds. Every frame drawn by AI from a description — the joke, the timing, and the cut are mine.

Roll it — then we go hands-on. The next ten minutes are live, on my screen, with the worksheet in front of you.

And now — onto the demo.

Ten minutes, live in Copilot. Today I drive — this week, you do.


THREE THINGS TO CATCH ME DOING:

1 • critique, don't edit

2 • the meta-prompt moment

3 • verify before it ships

Grab the worksheet. Its demo cues match these three — check them off as you spot them, then keep it for your first tow this week.



the demo is the tow line — watch the hitch happen live



This deck towed itself

Everything we just asked you to do was done to build what you watched.

■ THE PROMPT • VERBATIM, EVERY TYPO

“page 10 rockets are clipping and not high quality. do a pass to make sure all space ships are high quality for 1200dpi printing — recreating them better so you can resize with vectoring on your backend?”

why it worked: it named the defect and set the bar (1200 DPI, vector) — but left the how to the engine.




THE FIRST PASS

Flame cropped mid-burn, 125 DPI, stray pixels in the corners.



three rounds of critique after that prompt — zero hand-edits



SHIPPED

Vector-rebuilt from the original hand art, ≈1,900 DPI. Recreated, never regenerated.

1

Critique, don't edit

No slide survived its first draft — “How to start” went through three full redesigns before it made sense.

2

Verify before ship

All eight sources in the appendix were re-checked against the live originals the week of this talk.

3

Own the result

A person signed off on every slide, every claim, every ship. The AI pulled; a human steered.

The worksheet in front of you runs this exact loop. Follow along during the demo — then hand it to your team.



“I don’t understand at all.”

■ THE PROMPT • VERBATIM, EVERY TYPO

“how to start page 13 needs to be better UI and unique, i dont understand at all”

why it worked: no fix, no spec — just the honest symptom. That’s the whole move.

■ HOW TO START

Start with a pilot you can measure

Not a transformation program — a first crossing, with a course you can chart.

1 Pick 1–2 workflows
Low-risk, high-verification tasks where a win is obvious.

2 Set a measurable outcome
Decide the metric up front — before you sail.

3 Prove it, then scale
Near-term wins fund the next leg and build the muscle.

It's not only about ROI. The first pilot is how a team gets comfortable, breaks the uncertainty wall, and starts the compounding now.

THE DRAFT — a route map with diagonal waypoints. Dense, clever, unreadable.



■ HOW TO START

Start with a pilot you can measure

Not a transformation program — a first crossing, with a course you can chart.

1 Pick one routine task
Low-risk, easy to double-check, done every week — a task where a win is obvious.
status emails • meeting notes • report digests

2 Set the measure up front
Before you sail, write down the one number that will prove the win:
hours saved • cycle time • error rate

3 Prove it, then share it
Run it weekly, verify every output, and say the number out loud at day 30.
the win funds tow #2 →

It's not only about ROI. The first pilot is how a team gets comfortable, breaks the uncertainty wall, and starts the compounding now.

Paradox • What AI Is • The Play • **Execution** • The Human Part • Close

13 / 23

SHIPPED — same content, one straight line: pad → three steps → first win.

The critique wasn’t a design spec — it was one honest sentence. Nobody hand-polished the broken version. It was told what was wrong, and it went again.



The live tow — eight moves

Run it with me on your screen. These are the real prompts — the same words are printed on your worksheet. Type them exactly, then swap one noun.

THE LOOP — every task, every time

0 Pick the task

— no prompt yet: the weekly update you already send your leader.

1 Draft — one honest prompt

Draft the weekly status email I send my director: what shipped, what slipped, what I need, next.

2 Critique, don't edit — ×3

Too corporate and too long. Cut it by half. You buried the ask — put it in the first two lines.

3 Meta-prompt — the jaw-drop

Rewrite my original prompt so it gets a better first draft on the first try - then run it.

4 Verify before it ships

List every factual claim as bullets so I can check each one, and flag anything you inferred.

GEAR SHIFTS — same draft, three more gears

5 Flip the audience

Rewrite that update for my team instead of my director - same facts, their language.

6 Flip the format

Now turn it into a five-bullet summary I can paste at the top of a slide.

7 Red-team it

Argue against this update. What's weak, risky, or missing? Rank the top three.

≈ 9 minutes, start to finish. Eight moves, one draft, zero hand-edits — and the same loop runs on your task at your desk this week.

Swap one noun. "Status email" becomes your thing — the loop doesn't change.

While we fly: governed tool only • nothing sensitive goes in • nothing ships until a human verifies it.



Q & A

Reactions, questions, and where you'd start.

There is no wrong question here — including the skeptical ones. If it's on your mind, it's on someone else's too.

STILL OPEN AFTER TODAY

Grab me in the hall, or send the prompt you're stuck on. *I'd rather rewrite one real prompt with you than answer ten hypotheticals.*

IF NOBODY GOES FIRST, I'LL ASK THESE:

?

What do I do when it's confidently wrong?

verify before it ships — that's move four.

?

What can I put into it — and what absolutely can't I?

governed tools only; nothing sensitive goes in.

?

How do I start my team without making it a project?

one routine task, one honest prompt, this week.

?

Isn't this just going to write worse versions of what I already write?

not if you critique instead of edit.

Then Monday: one routine task · one honest prompt · critique instead of edit · and the meta-prompt on Friday.



Five habits that do most of the work

Not tricks. These five show up in almost every prompt I write — and they're what built this deck.

1 Batch it. Number it.

One message, twelve numbered notes — not twelve messages. It answers as one coherent pass instead of twelve patches that fight each other.

2 Say what's wrong. Don't fix it.

"I don't understand at all" beat any design spec I could have written. Name the symptom; let the engine find the cure.

3 Give the constraint, not the solution.

"Make it look like Google made it." One constraint carries further than a paragraph of instructions — and leaves room to be surprised.

4 Bring your own material.

My hand-drawn ships. My role, my voice, my gecko. The output is only as original as what you feed it — AI recreates, it shouldn't invent you.

5 Make it show its work.

Sources re-checked, files hashed, every claim listed so I can verify it. Speed is worthless if you can't put your name on the result.

■ A REAL PROMPT FROM BUILDING THIS DECK • VERBATIM

`"1) i dont like the left vertical story block bar accent, make it look like google made it • 2) "interrupt me" take out that whole call to action, its obnoxious • 3) youre right its not 12-18 months, close to 6-12 months • ...and nine more"`
why it worked: twelve numbered notes in one message — blunt, batched, and not one of them told it how to fix anything.

None of this is prompt engineering. It's the same thing you already do with a sharp new hire: be clear, be honest, and check the work.



Glossary: the working vocabulary

Every term this deck leans on — one breath each, for whenever you come back to it.

LLM — large language model — AI trained on vast text; the engine behind Copilot, Claude, ChatGPT.

SLM — small language model — compact enough to run privately, at near-frontier quality.

Frontier model — the most capable current models from the big labs. Powerful — and external.

Agent / agentic AI — AI that plans and executes multi-step work with tools, not just answers.

GCC High — Microsoft’s US-government compliance cloud — one example of where this class of governed AI can run.

Copilot — Microsoft’s AI assistant across M365 — in gov-cloud form, one example of governed AI.

Prompt — the instruction you give the model. Quality in, quality out.

Meta-prompting — asking AI to improve your prompt before you run it.

Token — the unit AI usage is metered in — the fuel, or a lease’s mileage cap.

Context window — how much the model can hold in mind at once — bigger fits more of your docs.

RAG — retrieval-augmented generation — the model looks things up in your documents first.

Fine-tuning — extra training on your own data to specialize a model.

Model weights — the learned numbers that are the model. “Open weights” = you can host it yourself.

Distillation — teaching a small model to mimic a big one — one way frontier power comes inside.

Inference — running a trained model to get an answer (vs. training it).

Guardrails — the technical and policy limits on what a tool can see and do.

Hallucination — confident, plausible output that is simply wrong. Looks exactly like good work.

Slop — unverified AI output shipped as finished work — what rigor exists to stop.



Where the claims come from

Kennedy, “The Wait Calculation” (JBIS, 2006)

The real math behind the generation-ship paradox — when waiting beats launching, and when it stops.

OpenAI — ChatGPT launch (Nov 2022)

GPT-3.5 goes public and AI goes mainstream — the 2022 node.

Epoch AI — open- vs. closed-model gap (2025)

Open models trail the frontier by roughly four months — the gap is closing.

Microsoft 365 Roadmap — Copilot in GCC High

One example of tracking governed availability inside a compliance boundary.

Honest labels: “≈6–12 months behind the frontier” is our internal working estimate of typical compliance-cloud feature lag, and the compounding chart is illustrative — directional, not measured.

Mollick — Co-Intelligence (2024)

The best single on-ramp for leaders: working with AI as a teammate, not a tool.

Microsoft — Work Trend Index 2026

Annual data on AI at work — what “frontier firms” are doing differently.

One Useful Thing — Mollick’s newsletter

Short, practical, hype-free — a working manual that updates weekly.

■ FURTHER READING — IF YOU ONLY PICK THREE