


From Elevators to Agents

The Three Checkpoints for AI Acceptance

Pankaj Tyagi · Head of Technology · Sycamore Informatics Inc.
PHUSE SDE Chennai · 19 September 2026



Sycamore Clinical Agentic Framework™

SPD IS PROTOCOL	MDR METADATA	DTM TRANSFORM	DSW OPEN-SOURCE	SCE ANALYSIS	CDR REPOSITORY
--------------------	-----------------	------------------	--------------------	-----------------	-------------------

Background: Elisha Otis's elevator patent drawing, US 31,128 (1861), US National Archives · Wikimedia Commons, Public domain

THREE CHECKPOINTS

Every transformative technology has to clear all three.



CAPABILITY

Is it good enough?

REGULATORY

Can it be governed?

ACCEPTANCE

Will anyone use it?

1850s: the city runs out of ground

Industrial revolution → people to cities

Cities → out of land

The only direction left was up

Buildings stopped at ~7 storeys.


Not because of steel.

Because of stairs.



Broadway looking north from between Grand and Broome Streets, c. 1853-55 · attributed to Silas A. Holmes · public domain

Elisha Otis had solved it. Nobody was buying.



Elisha Graves Otis, 1811-1861 · Wikimedia Commons, Public domain

8

elevators sold in 1854

15

elevators sold in 1855

Otis Elevator's own archive

Sales figures: Otis Worldwide, Our history. "Orders were few until May 1854" — Britannica.

Because everybody knew what a hoist was

A plank on a rope, outside
the building, carrying bricks
and the men who trusted it.

Workers wouldn't ride it.
It crashed to the ground
when the cable broke.

WHAT PEOPLE PICTURED

An open platform.

A hemp rope.

A load of bricks.

No Manhattan customer is getting in a machine that looks like that.

Otis, 1859: ten men killed in four years in New York — on other people's ropes.



✓ CHECKPOINT 1 — CAPABILITY

The brake

Steel

cable that held ✓


Steam

motors that were reliable ✓

The safety brake

the thing that mattered ✓✓✓

If the cable breaks, a spring snaps open and locks into toothed rails.
The load's own weight tightens the grip.



✓ CHECKPOINT 2 — REGULATION


Patents and Local Approvals

Intellectual Property

key safety elevator patents ✓

Formal Approvals

building codes and mechanical certifications ✓



E. G. Otis, "Improvement in Hoisting Apparatus", US Patent 31,128, 15 January 1861 · US National Archives · public domain

✓ CHECKPOINT 2 — REGULATION

The 68-year gap

Safety brake invented: 1852

First unified American elevator safety code (ASME A17.1): 1921


1852

1921

68 years

Governed in the meantime by patents, insurers, and whatever the city happened to say.

"First published by ASME in 1921, A17.1 has been a primary source for all subsequent standards governing elevators..." — ASME



X CHECKPOINT 3 — ACCEPTANCE

New York Crystal Palace, May 1854

He stood on the platform.
He had them cut the rope.

“All safe, gentlemen, all safe.”

The platform dropped a few inches and stopped.



Artist's depiction of the 1854 demonstration · Wikimedia Commons, public domain. Quote attributed via Otis's corporate history, not a primary source.

And then he engineered the feeling

- **The operator**
a trained human, in the box, every trip
- **1948 · Otis Elevoice™**
the first talking elevator
- **1948 · Palmolive Building**
Otis pipes in music
- **1950 · Autotronic™**
the human finally comes out
↑ the two years of parallel run

**The reassurance shipped two years
before the human was removed.**



Elevator operator, Martha Washington Hotel, New York, 1917 · Bain News Service · Library of Congress · public domain

New York City subway, today

Five deep-bore stations.
Elevators automated in the 1970s.
The operators are still there.

MTA tried to remove them in
2003. 2007. 2018.
Pushed back every time.

**Fifty years of proven automation
did not retire the need
for a person.**

The parallel

ELEVATOR

Open freight hoist

Safety brake

Patent + public diagram

Operator in the car

Mirror, light, music

Cutting the rope in public

AGENT

Unfenced chatbot

Guardrails

Validation pack + audit trail

Human review that is real

Review UX that makes work faster

Evals you will show Quality

Same three gates. New machine.


	1854	2026
CAPABILITY	steel · steam · the brake	frontier models · guardrails
REGULATION	ASME A17.1, 68 years late	FDA at the table, now
ACCEPTANCE	operator · demo · Elevoice	?

Checkpoint 1 – The Capability



✓ CHECKPOINT 1 — CAPABILITY

The capability curve



This is not one company's bet

FRONTIER

Anthropic Claude

OpenAI GPT

Google Gemini

xAI Grok

Amazon Nova

OPEN WEIGHTS

Meta Llama

Mistral

Alibaba Qwen

DeepSeek

Microsoft Phi


Google Gemma

IBM Granite

WHERE WE RUN

AWS Bedrock — data residency, no training on your data

Frontier models require strict containment and insulation:



Secure Containment Architecture

- **Agent Container Security:** Secure sandbox execution environments.
- **Network Isolation:** Ensuring zero data leaks outside designated environments.
- **Validation Guardrails:** Strict input and output validation protocols.
- **The Otis Parallel:** Just as Elisha Otis's safety brake gave society the trust to build upward into skyscrapers, strict containment systems serve as the essential safety mechanism enabling the safe scaling of frontier AI.


Gate 1 is open.

Steel didn't make the elevator acceptable.

The brake did.

A capable model is a hoist.

Ship the brake.



1861



2026

What “good enough” now means

Keep the model safe

Aligned refusal · grounded generation · tool allow-lists

Contain the agent

OS + network isolation · least privilege · no shell by default

Do not leak the study

Zero data retention where required · no cross-study memory

Do not trust the prompt

Injection defense · untrusted-document handling

Prove what happened

Signed traces · replayable runs · model + prompt pin

Fail closed

Human escalation · deterministic hand-off

Eleven deterministic rails · three phases

PRE · 2

- **prompt_injection**
role-hijack · homoglyph / zero-width normalised
- **input_pii_detection**
PII/PHI in the inbound prompt



TOOL · 2

- **tool_param_validation**
path traversal · SQL/shell · study scope
- **tool_output_size**
truncates output over 32K



POST · 7

- **phi_detection**
SSN · email · phone · DOB · MRN
- **cross_study_leakage**
- **data_leakage_enhanced**
- **prompt_leak**
canary planted in the system prompt
- **output_grounding**
- **topical_confinement**
- **toxicity_claims**

■ block □ warn (tunable per agent)

Not an LLM judging an LLM. Deterministic checks.

KILL SWITCH

one admin-only, audited, reversible flag stops all LLM use platform-wide

Checkpoint 2 – Regulation

US & Global

- FDA AI for Drugs & Biologics Guidance
- FDA PCCP for AI-enabled medical devices
- FDA AI-enabled SaMD lifecycle guidance
- GMLP principles (FDA / Health Canada / MHRA)
- 21 CFR Part 11 • HIPAA • Section 1557
- ICH E6(R3) GCP computerized systems
- WHO AI for Health ethics & regulatory guidance

EU, UK & Asia

- EU AI Act (high-risk healthcare AI)
- EMA AI Reflection Paper (2024)
- FDA-EMA 10 Good AI Practice Principles (2026)
- EU MDR / IVDR Rule 11 (SaMD)
- GDPR • Clinical Trials Regulation 536/2014
- MHRA AI Airlock • Health Canada MLMD guidance
- China NMPA AI + Drug Regulation • Order 818

✓X CHECKPOINT 2 — REGULATION

The regulator is not asking
“is your model good?”

It is asking:

Who decided?

Can you show me?

**Can you show me again in seven
years?**

YOUR AI STRATEGY MUST COVER

The regulatory gate, assembled

Human review

Required, skilled, recorded — not a rubber stamp

Taint / watermark

Every AI-touched artifact stays marked

Effective review

Why, not only what — CoT on demand

Injection defense

Untrusted inputs cannot steer the agent

RBAC

The agent inherits the user's rights — never more

ZDR / tenancy

No cross-study contamination

Reproducibility

Pin model, prompt, tools — re-run or explain

Audit trail

Who asked, what ran, what changed, who signed

1 • Human review, enforced by classification

```
if readOnly, ok := tool.(ReadOnlyTool); ok && readOnly.ReadOnly() {  
    return false    // no human review needed  
}  
return true        // everything else: side-effecting
```

A tool is read-only only if it says so.

**New tool, unknown tool, broken tool →
human review.**



2 • Taint the output. Two ways.

The tool path

every write_file, with a SHA-256

The syscall path

strace -ff on every child process,
catching any script writing anywhere

Captures write-intent, not proof-of-change —
biased to over-flag. The safe direction.

user_files

adae_t14_2_3.sas

ai_gen_status

none

ai_generated

human_reviewed

ai_trace_url

→ [view the full generation trace](#)

Re-generate the content and it
drops back to ai_generated.

CDR-SCE user_files.ai_gen_status + ai_trace_url, shipped 26.2 (#11280).
Illustrative record.

3 · A signoff you can't justify is theatre

Reasoning

every ReAct step, every tool call, every result

Injection defence

Unicode-normalised; canary planted in the prompt

Authorisation

only a signing role may sign

Retention

prompt/response text is a toggle — the hash is always kept

Isolation

cross-study writes denied by the OS, per-job uid drop

4 • The model will be gone before the trial is



frontier model deprecation window 12–24 months



your retention obligation 7+ years

5x – 20x

Bit-identical re-execution is not the regulatory bar.

Proving what the model did, and why a human approved it — is.

Your AI strategy must answer all four

1 • REVIEW

enforced by default-unsafe classification

2 • TAINT

every AI byte flagged, two capture paths

3 • TRACEABILITY

reasoning, scope, roles, retention

4 • REPRODUCIBILITY

(skill × model) pair, never auto-migrated

Miss one and the other three don't count.

Checkpoint 3 – Acceptance



User Experience

Effective Reviews

Parallel Operations

Night-shift

"Chief-of-staff" Bots


X CHECKPOINT 3 — ACCEPTANCE

1 • Make the validated path the easy path

SDV — Skill Development & Validation



ADV — Agent Development & Validation



The machine does the labour. The human does the deciding.

2 · An agent is a lifecycle, not a launch

- EOL notice arrives
- new draft version minted automatically
- eval suite runs
- diff + score delta presented
- **a human signs** ← the only gate
- pair promoted to validated

Provider substituted outside the skill's validated set?

blocked_reason:
provider_not_validated

A skill may declare
`allow_provider_substitution: false`
when GxP validation was performed against one provider only.

3 • Review has to be possible, at scale

sdtm-review — Phase 5, human gate

Per-domain disposition, computed **deterministically in Python** from finding severity and domain criticality against the validation plan's rubric.

The language model does not decide it. Same findings, same disposition, every time.

Resumable • 72-hour window
Ends in a human e-signature

SDTM delivery review — STUDY-1042 — 8 domains

DOMAIN	ERRORS	WARNINGS	DISPOSITION	
DM	0	2	ACCEPT	view trace
AE	14	31	REJECT	view trace
LB	3	128	REVIEW REQUIRED	view trace
VS	0	7	ACCEPT	view trace
CM	1	22	REVIEW REQUIRED	view trace
EX	0	0	ACCEPT	view trace
MH	0	4	ACCEPT	view trace
DS	6	11	REJECT	view trace

[Approve All](#) [Approve by Domain](#) [Reject All](#) [Add Note](#)

Illustrative mock-up. Dispositions are severity-derived, not model confidence scores.

The three gates, and our job at each



CAPABILITY

stop relitigating it. Ship the brake.

REGULATORY

HITL · taint · transparency · reproducibility

HUMAN ACCEPTANCE

make the validated path the easy path —
and make the reviewer's job possible

Otis needed a brake and a showman and an operator. So do we.



The elevator didn't make buildings taller.
It made the top floor the place you wanted to be.

Somebody in this room will see a cancer beaten.


The work is getting the evidence there faster — and doing it in a way a regulator, a reviewer, and a patient can all trust.

**Build the brake.
Then open the gate.**

Thank you

Pankaj Tyagi · Head of Technology · Sycamore Informatics Inc.

PHUSE SDE Chennai · 19 September 2026



Sycamore Informatics

Sycamore Clinical Agentic Framework™

SPD
PROTOCOL

MDR
METADATA

DTM
TRANSFORM

DSW
OPEN-SOURCE

SCE
ANALYSIS

CDR
REPOSITORY

Inventor,

E. G. Cat

BACKUP — DO NOT PRESENT · HOLD FOR Q&A

B1 — “What's your accuracy?”

There are NO measured accuracy figures in our documentation. Say so.

tlg-programmer generates SAS, runs it inside CDR-SCE on real study data via sycamore-cli, parses the log, fixes, re-runs. Proven on a sponsor study / POC-04. sdtm-review came out of a named sponsor Proof-of-Concept. 75% cost reduction on SDTM Qualification, measured ($\sim \$10 \rightarrow \sim \$2-4$ per run). Prompt-cache savings now measured and displayed, not assumed. $\geq 70\%$ eval quality gate for promotion to validated. Six production skills declare RunScript today.

BACKUP — DO NOT PRESENT · HOLD FOR Q&A

B2 — Full guardrail table

Eleven rails × phase × severity. Full table in FACT-CHECK.md §B3.

The honest caveats — lead with these:

toxicity_claims downgraded block→warn (false-positived on legitimate clinical reporting)

the tool phase had zero production callers until 26.2

semantic-paraphrase injection still open (tier-2 grader designed, not built)

PHI patterns are US/English-centric; post-execution scan adds ~200ms

per-agent severity tuning means an audit must inspect agent CONFIG, not just the registry

the kill switch fails OPEN on a settings-read error, so a DB blip can't wedge the platform

BACKUP — DO NOT PRESENT · HOLD FOR Q&A

B3 — Three-tier audit schema

audit_log

one row per significant business action

ai_audit_log

one row per LLM call: feature, provider, model, skill version, tokens, latency, stop reason, always-present prompt hash, temperature, max_tokens, stop_sequences

events

event-sourced record with correlation ID and sequence

Immutability lives at the SQL boundary, not in application code.

BEFORE UPDATE / BEFORE DELETE triggers raise an exception — so accidental ORM writes, buggy code and malicious access all fail identically. Nothing hard-deletes.

BACKUP — DO NOT PRESENT · HOLD FOR Q&A

B4 — Cost controls

\$2.00 default per-execution ceiling, enforced by CostTracker halting the ReAct loop

ReAct loop bounded at 20 iterations with 7 stop conditions

Prompt caching, 3 breakpoints — ~10% of input rate on a cache hit

Admin-only, audited, reversible kill switch

The war story:

a runaway knowledge-consolidation loop cost ~\$35/day until we found it —
after which that whole subsystem became opt-in, default off.

BACKUP — DO NOT PRESENT · HOLD FOR Q&A

B5 — Otis sources

US Patent 31,128 — patents.google.com/patent/US31128A/en

Otis Worldwide, Our history — otis.com/en/uk/our-company/history

ASME, Elisha Graves Otis — asme.org

ASME A17.1 catalogue, first published 1921 — files.asme.org

Britannica, Elisha Otis — britannica.com/biography/Elisha-Otis

Paumgarten, “Up and Then Down”, The New Yorker, 21 Apr 2008

SEIU 32BJ, Our History (the September 1945 strike) — seiu32bj.org/history

Bethune, Maclean's, 24 Mar 2014 — the penthouse inversion

BACKUP — DO NOT PRESENT · HOLD FOR Q&A

B6 — Image credits

All imagery is public domain or Creative Commons.

Slides 1, 6 — Otis elevator patent drawing, US 31,128 (1861). US National Archives. PD.

Slide 3 — Broadway looking north from between Grand and Broome Sts, c.1853–55.

Attributed to Silas A. Holmes. Wikimedia Commons, PD.

Slide 4 — Elisha Graves Otis. Wikimedia Commons, PD.

Slide 5 — Lewis Hine, ironworker on the Empire State Building, 1930–31. NARA, PD.

Slide 8 — artist's depiction of the 1854 Crystal Palace demonstration. Commons, PD.

Slide 9 — Elevator girl, Martha Washington Hotel, New York, 1917.

Bain News Service, Library of Congress. Wikimedia Commons, PD.

Slide 10 — 191st Street station elevators. MTA of the State of New York, CC BY 2.0.

Slide 12 — METR, “Measuring AI Ability to Complete Long Tasks”. Reproduced with attribution.

Slide 26 — Atrium, Hyatt Regency Embarcadero. Dale Cruse, CC BY 2.0.