

How powerful is your AI operation?

Agentic Power measures how much skilled human-equivalent work an AI operation produces for every hour of human direction.

How much skilled work can one hour of human attention command?

Developed by **Dr. Mark Allen** at **HeroForge.AI**

AP = 30x

One hour of human direction produces approximately thirty hours of comparable skilled work.

Measure. Increase. Direct.

AI is becoming a workforce you can direct.

People and organizations can delegate useful work to AI while supplying goals, context, judgment and accountability.

Measure the power of your AI operation.

Agentic Power measures skilled human-equivalent work produced per hour of human direction.

$$\text{AP} = \text{Skilled Human-Equivalent Hours} \div \text{Human Direction Hours}$$

AP = 30× means approximately thirty hours of comparable skilled work for one hour of human direction. Count accepted work and all attributable direction, including setup, review and correction.

Understand its economic efficiency.

Agentic Power Efficiency (AP\$) compares equivalent human labor cost with the attributable cost of producing that work.

$$\text{AP\$} = \text{Equivalent Labor Cost} \div \text{Total Attributable Agentic Cost}$$

AP\$ = 10× means an estimated \$10 of comparable human labor cost for every \$1 of attributable agentic cost. This does not establish actual savings or profit.

Understand what shapes your power—and where to improve it.

Delegation examines the responsibility AI carries. Autonomy examines how independently work proceeds. Self-improvement examines how the operation gets better. Together, they help you identify constraints and choose changes that could increase individual or organizational AP.

Adopt capability instead of building everything yourself.

Adopt stronger models, agent working environments, tools, reusable skills, workflows and complete AI staff. Assemble and adapt existing capability to your work.

Use Agentic Engineering to turn capability into greater power.

Improve how the components work together: responsibilities, context, permissions, coordination, evaluation and human oversight. Measure whether the changes produce more accepted work per hour of direction.

Build improvement into the operation.

Use results to improve instructions, workflows and system design. People can make those improvements, use AI assistance, or authorize evaluated changes within defined boundaries.

Direct greater power toward worthwhile outcomes.

More capacity and better economics create options. Human judgment determines which goals deserve that capacity. Assess the outcomes, not just the power available.

How much skilled work can one hour of human attention command?

Busy AI is not productive AI.

A hundred agents producing a thousand drafts nobody uses is not productive power.

Tokens, prompts, agent counts, runtime and model benchmarks do not directly answer the practical question:

How much useful skilled work did the system produce per hour of human direction?

A familiar idea: horsepower.

When steam engines began replacing animal power, people needed a practical way to understand and compare what unfamiliar machines could do. James Watt helped popularize horsepower as a common language for machine power.

Horsepower gave people an intuitive measuring stick for a new source of productive capability. AI now presents a similar problem for knowledge work.

Agentic Power provides a common measure for that work: useful skilled output relative to the human attention required to direct it.

Agentic Power.

Agentic Power measures how much skilled human-equivalent work an AI operation produces for every hour of human direction.

$$AP = \frac{\text{Skilled Human-Equivalent Hours}}{\text{Human Direction Hours}}$$

SAME SCOPE. REPRESENTATIVE WINDOW. SAME QUALITY BAR.

Skilled Human-Equivalent Hours (HEH)

Human-Equivalent Hours estimate how much time a competent skilled person or team would reasonably need to deliver comparable-quality accepted work.

A useful analysis that would take a skilled analyst ten hours represents about **10 HEH**—even if you could not do that work yourself.

SAME OUTCOME · SAME QUALITY BAR · REASONABLE HUMAN BASELINE

What counts as human direction?

Attributable human attention: briefing, setup, training, supplying context, connecting tools, review, correction, coordination, exception handling, maintenance, evaluation and improving the system.

Include everyone who directs the operation. Do not count unattended AI runtime.

Agentic Power is not intelligence, runtime, automation percentage, or financial ROI. It measures one thing: **skilled human-equivalent output per hour of human direction.**

Elapsed time is how long the work took on the calendar. Human direction is the combined active time people contributed; only that time goes in AP's denominator.

Economic translation

Not all skilled hours cost the same.

Agentic Power deliberately treats one Human-Equivalent Hour as one hour of comparable skilled work.

That keeps AP simple.

But different skills carry very different labor costs.

Ten hours of bookkeeping and ten hours of specialized legal work both count as 10 HEH, while the market cost of obtaining that human capacity may differ dramatically.

Equivalent Labor Cost

$$= \Sigma (\text{HEH} \times \text{benchmark hourly labor cost})$$

For each kind of work, multiply its HEH by the appropriate hourly labor cost, then add the amounts.

10 HEH bookkeeping × USD 50/hour =

USD 500

10 HEH specialized legal × USD 400/hour =

USD 4,000

Illustrative rates, not observed market benchmarks.

Agentic Power Efficiency (AP\$)

Economic efficiency

Equivalent Labor Cost translates HEH into benchmark human labor cost. Total Attributable Agentic Cost includes direct AI spending, the cost of human direction, and other costs attributable to the same work and period.

AP\$

$$= \text{Equivalent Labor Cost} \div \text{Total Attributable Agentic Cost}$$

AP\$ = 10× means the work represents an estimated \$10 of comparable human labor cost for every \$1 of total attributable agentic cost. This does not establish actual savings or profit.

Agentic Power tells us how much skilled work human attention can command. AP\$ tells us how economically that work was produced compared with obtaining comparable human labor.

It does not tell you what the work was worth to the business. This is not ROI, realized savings, profit, or business value.

ILLUSTRATIVE EXAMPLE

AGENTIC POWER · AP = 17.5×

35 HEH ÷ 2 hours human direction

EQUIVALENT LABOR COST

USD 4,500

20 HEH bookkeeping × USD 50

10 HEH software × USD 150

5 HEH specialized legal × USD 400

TOTAL ATTRIBUTABLE AGENTIC COST

USD 420

Direct AI Cost USD 120

Human Direction Cost USD 300

Other attributable costs USD 0

SECONDARY ECONOMIC METRIC

AP\$ = 10.7×

4,500 ÷ 420 = 10.714...

Every USD 1 of attributable agentic cost produced accepted work equivalent to approximately USD 10.70 of benchmark human labor cost.

This is not ROI, realized savings, profit, or business value.

A current operating view covers the selected period. A fully loaded view also accounts for attributable setup and earlier investment across the lifecycle.

Current operating view · Labor market: United States national (fictional assumption) · Rate basis: custom internal rate · Source: fictional teaching assumptions, not observed benchmarks · Source date: September 2026 · Currency: USD. Human direction: 2 hours × USD 150/hour. Setup investment is not estimated; this is not a lifecycle claim.

Agentic Power measures productive amplification.

AP\$ measures economic efficiency.

Business outcomes tell us whether the work was worth doing.

Two Ratios. Two Different Questions.

	AP	AP\$
	Agentic Power	Agentic Power Efficiency
Question	How much skilled work can one hour of human attention command?	How much comparable human labor cost can one dollar of agentic cost command?
Formula	$\text{HEH} \div \text{Human Direction Hours}$	$\text{Equivalent Labor Cost} \div \text{Total Attributable Agentic Cost}$
Example	17.5x	10.7x
Measures	Productive amplification	Economic efficiency

Neither metric tells you whether the work was worth doing. Business outcomes do.

Economic methodology and reporting rules

Three operations. Three estimates of Agentic Power.

An AI employee handles a recurring specialist role; an AI department handles a whole business function. Persistent operation means software keeps working within defined permissions. These describe how the work is organized—not how much Agentic Power it produces.

AI EMPLOYEE Ajay Jani, MD AP ≈ 16.7× ≈200 HEH ≈12 human-direction hours ≈2 elapsed days The work and the approach After learning Claude Cowork and reusable Skills—saved instructions and procedures for recurring work—Ajay tackled unfinished administration. This built on earlier instruction in writing prompts and using AI to improve them. The work included a credentialing application, tax-document preparation, license tracking and document organization. Ajay described saving “hundreds of hours.” 200 estimated HEH ÷ 12 self-reported direction hours. Dr. Mark Allen confirms the work met a high standard.	AI DEPARTMENT Rowena Chua, MD + Emmanuel Cecilio AP = 55× ≈22,000 HEH ≈400 combined human-direction hours ≈110 elapsed days The work and the approach Rowena had no technical background; Emmanuel had some IT experience but was not a software engineer. During a three-day Mastermind retreat, Dr. Mark Allen taught them to install, configure and use Ruflo as a pre-built software engineering department. Their website, marketing and communications, AI clinical reporting, education platform and staff clinical workflows are now in production. Rowena describes the system as integral to her business. A continued-learning and Mastermind journey over approximately 110 days; other systems remain in development. Retrospective HeroForge.AI / AI-assisted work estimate.	PERSISTENT AI OPERATING SYSTEM Reuven “rUv” Cohen AP ≈ 33,000× FRONTIER CASE AI-DERIVED PRODUCTION-HISTORY ESTIMATE ≈80M HEH estimated ≈2,404 human-direction hours modeled ≈253-day measurement window The work and the approach Public software including RuView, Ruflo, RuVector and MetaHarness. rUv ran a detailed audit prompt in his own Codex environment and verbally reported an estimate of 10 million working days for his year-to-date code output. Human direction estimated using an assumed daily schedule.
--	--	---

Operating models describe the work. AP measures the result.

Different work and estimation methods; these cases are not a competitive leaderboard. Broader delegation or greater autonomy does not by itself establish higher AP, better quality or better outcomes.

Behind the large number: a visible body of work.

rUv builds high-visibility open-source software. These public projects make part of that output inspectable.

- [RuView](#) – WiFi-based presence and movement sensing.
- [Ruflo](#) – AI-agent coordination, shared memory and workflows.
- [RuVector](#) – vector and graph memory infrastructure.
- [MetaHarness](#) – tools for building focused working environments for AI agents.

“As an advisor to Cognitum.One, I have an unusual view into rUv’s software output beyond what is visible on GitHub.

I have personally watched his autonomous systems produce more high-quality software during a week when he was on vacation than I have seen many conventional organizations produce in a month.”

Dr. Mark Allen

Dr. Mark Allen is an advisor to Cognitum.One. This is first-person observation, not an independent benchmark.

A different operating model can create a different scale of output.

Proven work keeps moving when human attention moves elsewhere. That is why autonomy can make Agentic Power so large. Elapsed time still matters: the same AP delivered in two days or two months can serve very different needs.

What is your Agentic Power?

The Agentic Power Audit examines accepted work, estimates comparable skilled effort and reconciles human direction. It produces a Profile that leads with AP, then connects accepted work, comparable labor cost, actual costs and outcomes.

Example Agentic Power Profile

FICTIONAL EXAMPLE – ILLUSTRATIVE RATES

Your Agentic Power

AP ≈ 17.5x

184 HEH ÷ 10.5 direction hours

Seven-day window

AGENTIC POWER EFFICIENCY

AP\$ = 10.8x

Current Operating AP\$ · USD 17,100 ÷ USD 1,585

Accepted work

184 HEH

Equivalent Labor Cost: **USD 17,100**

Total Attributable Agentic Cost: **USD 1,585**

Direct AI Cost USD 220

Human Direction Cost USD 1,312.50

Other attributable costs USD 52.50

Labor market: United States national (fictional assumption) · Rate basis: custom internal rate · Source: fictional Profile assumptions, not observed benchmarks · Source date: September 2026 · Currency: USD. Human direction: 10.5 hours × USD 125/hour. Historical setup investment is not estimated; this is not a fully loaded result. These are not Dr. Mark Allen’s or HeroForge.AI’s economics.

The Agentic Power Framework · Dr. Mark Allen | HeroForge.AI · v1.5.0

11 / 42

What did it produce?

Accepted customer research, a released reporting application, an approved sales campaign and three completed administrative processes.

Work records, HEH and illustrative rates

What was it directed toward?

Return routine operator time while maintaining quality.

Routine direction: 9 → 6 hours

Across comparable weekly workloads.

Routine direction is part of total human direction. The 10.5-hour total also includes other attributable effort, including improvement work.

Partial coverage: three of four declared sources. Figures describe covered work only; the example does not establish causation.

Diagnostics

These observations help identify constraints and choose an improvement to test. They cover different parts of the work, not one combined operating level. The next sections explain how to use delegation, autonomy, system design and improvement to increase AP.

Responsibility delegated

Recurring specialist roles account for the largest share of accepted work.

System design

Good task context; tools, permissions and the working environment need improvement.

Main bottleneck

Repeated human handoffs

How work runs

Interactive work accounts for the largest share of accepted work.

Improving the system

AI suggests changes; people approve and apply them.

Next improvement

Test one recurring workflow with clear acceptance checks.

Choose how deeply to look.

The Basic Audit gives a quick estimate from available conversation context and information you provide. The Full Evidence Audit follows a guide to examine work records, files and completed outputs.

For the Full Audit, use a **chatbot with file access** through uploads, connected services or folders you share. **No command line required.** Start in one conversation: your chatbot reviews relevant records across the sessions and projects it can access, tells you what it searched, and helps you supply missing evidence.

You don't need to read the entire guide before starting. Your chatbot uses it to lead you through the audit. [Read the Full Evidence Audit Guide](#), or use either copyable starting prompt on the framework website.

Using this in a workshop? Run the Basic Audit before you begin. Return to the same chat afterward and use the [workshop follow-up prompt](#) to measure just the new interval. If other chats are inaccessible, one short recap supplies the missing work.

[See how the Audit works](#) · [Try the AP calculator](#)

The audit runs in your chosen AI environment. Downloading a guide does not send HeroForge.AI your work records.

Your AP tells you how much productive power the operation currently delivers. To increase it, examine how responsibility is delegated, how independently work proceeds, and how the operation improves. These reveal possible constraints and opportunities; measuring again tells you whether a change helped.

Adopt capability. Build your power.

You can increase your own or your organization's AP by selecting, combining and adapting capability that already exists.

Start with the constraint in your work. You might need a model better suited to the task, an agent working environment, a tool, a reusable skill, a proven workflow, or a complete AI employee, team or department. Adoption can happen at any of these levels.

A **model** provides the AI's underlying capabilities. A **harness** is the working environment that connects the model to instructions, tools, memory and controls. Adopting a better combination can change what work you can delegate and how much attention it needs.

Capability creates possibilities. Agentic Power is realized in use.

You still supply goals, context, permissions, review and judgment. The same model or system can produce different AP in different hands. Count setup, adaptation and onboarding effort when measuring the result.

Use delegation, autonomy and self-improvement to find opportunities. Agentic Engineering helps you turn those opportunities into a working operation.

The higher you climb, the more you direct goals.

As AI becomes more capable, humans can delegate broader responsibility: from individual tasks to roles, outcomes and entire operations. Use the ladder to ask what responsibility your operation could reliably carry next. Humans remain accountable.

AI staff is an organizational metaphor: software systems can take on roles and responsibilities without being human employees.

R labels the responsibility level, from R0 (human work) to R6 (an AI portfolio).

↑ Broader responsibility • Read upward from R0

R6	AI Portfolio Govern multiple AI-native operations.	"Allocate resources across these operations within our goals and risk limits."
R5	AI Company Delegate organizational goals + constraints.	"/goal Grow revenue while maintaining these quality, margin and customer constraints."
R4	AI Department Delegate a business function.	"/goal Increase qualified marketing pipeline by 25% within this budget."
R3	AI Team Delegate an outcome.	"/goal Launch the customer research report by Friday."
R2	AI Employee Delegate a recurring role.	"Own our weekly competitive intelligence."
R1	AI Assistant Delegate tasks.	"Draft this email."
R0	Human You perform the work.	No responsibility delegated to AI.

Instructions become more abstract as responsibility rises. A higher rung needs the capability, coordination and controls to deliver what was delegated.

Task	→	Role	→	Outcome	→	Goal	→	Goals + constraints
------	---	------	---	---------	---	------	---	---------------------

/goal is an illustrative interface metaphor, not a universal command. Goals still need quality standards, constraints and escalation rules.

The higher you climb, the less you direct tasks and the more you direct goals.

If repeated task-by-task instructions consume your attention, test delegating a complete recurring responsibility with clear standards and escalation rules. Broader delegation helps only if the accepted work justifies the human effort it requires.

Skills + Tools / Connectors


A useful AI Employee needs more than a good prompt. It needs to know its job and have the tools required to do it.

Skills make the role repeatable. Tools / Connectors provide access to approved information and allow permitted actions, with appropriate context and checks.

Work can begin without waiting for you.

Autonomy Levels help you identify where proven work could proceed with less avoidable intervention, including when you are not there to initiate each action.

A1–A5 label how a job starts and continues, from interactive work to proactively pursuing an ongoing goal.



Autonomy changes what initiates the work.

A4 responds to a defined event. A5 identifies what needs doing and initiates the next action. Running continuously alone does not establish A5.


Look for unnecessary starts, handoffs and check-ins. A defined assignment, schedule or event trigger may let proven work proceed with less attention. Keep review and escalation where they protect quality, and measure the resulting AP.

When accepted work keeps growing without the same growth in human direction, AP can rise dramatically.

An AI Employee can proactively identify and address needs within its role. An AI Department can still wait for a human assignment. Responsibility and autonomy are different choices.

Turn capability into a productive operation.

Agentic Engineering is the human discipline of selecting, combining, designing, directing, evaluating and improving AI systems that perform useful work. It turns adopted or newly built capability into an operation you can rely on and improve.



Each step zooms out from one instruction toward the entire operation. Improve the layer that constrains useful work, then measure whether the change increased AP.

Put AI staff into practice.

When capability is organized as AI staff, use this lifecycle to put it to work and improve it. Models, harnesses, tools and workflows support the staff at every stage.

Hire

Choose or build a specialist for a real job.

Onboard

Connect context, tools, permissions and standards.

Train

Teach Skills and evaluate representative work.

Deploy

Run within boundaries and collect evidence.

Promote

Expand responsibility when performance earns trust.

Deployment → evidence → learning → improvements → earned trust

Evidence feeds back into Skills, context, tools and operating design at every stage. Use what you learn to improve the next cycle.

Improve the system doing the work.

Proactive work advances the goal. Self-directed improvement strengthens the system that pursues it.

Autonomy asks how independently it works. Improvement asks how it gets better.

Use results from completed work to improve the instructions, tools or workflow used for future jobs. People can lead this feedback loop or give AI increasing ownership of it.

Proactive work can use an unchanged process.

A research agent may notice an emerging issue and prepare a briefing without being asked, using the same research Skill.

Improvement can happen between assignments.

A research agent may identify a weakness in its research process, test a revised Skill, and retain the revision after demonstrating better results. People set the goals and boundaries; higher-risk changes can require approval.

Goal → **Act** → **Measure** → **Evaluate** → **Learn** → **Improve** → **Act again**

↻ **Learn from the outcome. Improve the next cycle.**

Some Agentic Power can be reinvested into creating more future Agentic Power.

Choose a recurring error or source of rework. Use its evidence to improve the instructions or workflow, test comparable work again, and measure AP with the improvement effort included. Measure whether the change improves accepted work, quality or human attention.

Goal-Directed Self-Improvement is an Agentic Engineering method that can support every stage of the AI staff lifecycle. It can operate within an AI Employee, Team, Department or Company.

Improve toward the actual goal: revenue, better software, or less owner / operator attention with equal or better quality. Human judgment sets the constraints, quality standards and safety boundaries.

For a detailed assessment of who owns the improvement process, see the [Audit's improvement statuses](#).


Understand your operation. Choose your next improvement.

Ask three questions: **Responsibility: What does it own?** **Autonomy: How independently does it work?** **Improvement: How does it get better?** The map shows responsibility and autonomy. The improvement loop helps you change the operation over time; its detailed status belongs in the Audit.

These numbers label categories; they are not scores to add or multiply.

Use your position to identify constraints and choose a change to test. Measure AP before and after comparable work to find out whether the change helped.

Rowena and Emmanuel's software operation illustrates the notation: **R4 · A2 – AI Department · Delegated**. They delegate departmental work through assignments and lead improvements themselves. A question to test might be whether a proven job could start on a schedule. The labels do not establish that a change is needed—or determine the case's AP = 55×



[Download the framework graphic](#) · HEH = Skilled Human-Equivalent Hours; Human Direction is measured in hours.

Not every workflow should maximize autonomy. Risk, quality, economics, regulation, and human preference matter.

Four levers. One productive system.

AP rises when accepted HEH grows faster than the human direction required to produce it. Use what you learned about delegation, autonomy and improvement to choose among four practical levers. You can adopt existing capability or engineer a change at any of these points.

Capability	System design	Scale + autonomy	Learning
Give AI the ability to do better work.	Make the operation reliable and coordinated.	Let proven work run more broadly and independently.	Turn performance evidence into better future execution.

The levers can reinforce one another. Better system design can reduce correction; reusable learning can improve the next assignment. **Measure the result, including the human effort invested in improvement.**

Once you have measured an improvement in AP, you can consider what it means for the capacity available to your organization.

Productive capacity becomes an organizational resource.

Productive Output

$$= \text{Agentic Power} \times \text{Human Direction Capacity}$$

Productive output counts work that meets the agreed standard. Output is measured in skilled human-equivalent hours; direction capacity is measured in human hours.

Human Direction Capacity means the human hours available to direct the operation during the chosen period.

Higher Agentic Power. Less human attention for the same output.

SAME OUTPUT: **1,000 HEH**

AP = 10×

100 human-direction hours

AP = 100×

10 human-direction hours

Same work. Same quality. Illustrative arithmetic.

For a given quantity and quality of output, higher AP means fewer human-direction hours are required. For a fixed amount of human-direction time, higher AP means more productive capacity is available. Apply this relationship at a comparable AP and quality level; a past ratio does not guarantee unlimited scale.

Higher Agentic Power can make dramatically leaner organizations possible, while also allowing the same workforce to produce far more.

From one workflow to a company.


Measure AP for a workflow, project, person, team, department or company.

Aggregate the accepted work. Reconcile attributable human-direction hours. Then divide. Do not average individual AP ratios.

Comparable AI capability is getting cheaper.

Agentic Power can rise as models become more capable, systems are better engineered, delegation and autonomy increase, and Goal-Directed Self-Improvement improves the operation. At the same time, recent costs of running AI at a fixed capability level have shown extremely rapid, often multiplicative declines.

This chart compares the cost of running models at a similar score on a broad knowledge benchmark (MMLU). Tokens are the small units of text used to price that processing.



Epoch AI measured fixed-performance price declines of roughly 9× to 900× per year across benchmarks, including approximately 40× for one GPT-4-level science threshold. Rates vary substantially; the fastest declines are recent and may not persist. [Sources and limits](#).

A cheaper unit of capability does not guarantee a smaller AI bill. Organizations can use more AI, run more agents, process more context, reason longer and automate more workflows. Stanford's 2026 AI Index reports record compute costs and infrastructure investment. [Source](#).

Agentic Power can rise while the unit cost of the intelligence underneath it falls.

More productive capacity per human hour, at a falling unit cost of machine intelligence, changes the economics of skilled work.

Cheaper intelligence does not only reduce the cost of existing work. It expands the universe of work worth doing.

When the cost of producing intelligence falls, work that was previously too expensive becomes economical. Falling AI capability costs do not only threaten existing labor economics; they can create demand for new forms of cognition and work. Whether recent rates of decline persist is unknown. This can make existing work cheaper or make previously impractical work worth attempting. [The jobs section explores both possibilities](#).

What does this mean for the people we hire?

In an AI-native economy, we will increasingly hire people not only for what they know, but for how much useful capacity they can command.

If two people perform broadly comparable work at the same quality bar, with similar scope and under comparable conditions, a consistent, substantial difference in AP is economically meaningful.

That matters for hiring, compensation, bonuses, promotion, equity, team structure and workforce planning.

AI-native companies should increasingly recruit, develop, and reward people who can create and command high Agentic Power.

Share the value people create.

When a small number of people create dramatically more productive capacity, there is a strong economic argument for sharing more of the value they create with them.

Higher compensation, bonuses, equity, profit sharing or broader responsibility may follow. This is an economic implication, not a salary formula.

Some people raise everyone else's AP.

Agentic Engineers build and improve the instructions, tools, workflows, coordination and evaluations that help an organization use AI effectively. Their value can extend far beyond their personal output.

Capacity and economics tell you what is possible. Human judgment determines which outcomes deserve that power.

Power is not the destination.

Agentic Power is not the destination. It is the power available to reach the destination.

Horsepower does not tell you where to drive.

Agentic Power does not tell you what to build.

Agentic Power multiplies execution.

Human judgment determines where that power goes.

Higher Agentic Power, well directed with good judgment, should increase an individual or organization’s ability to achieve the outcomes it cares about.

Those outcomes may be revenue, profit, retention, customer satisfaction, quality, software released, research completed, innovation, cycle time, cost reduction, lives improved, risk reduction or human time returned.

Track the power alongside the outcome.

AP ↑ Profit ?	AP ↑ Revenue ?	AP ↑ Customer satisfaction ?
AP ↑ Release quality ?	AP ↑ Cycle time ?	AP ↑ Owner / operator time returned ?

If AP rises but the intended outcome does not improve, investigate goal selection, judgment, prioritization, quality problems, organizational bottlenecks, inability to absorb the new capacity or effort directed toward low-value work.

Track comparable windows and changed conditions. A correlation between AP and an outcome does not establish causation.

An early outcome worth pursuing

First, get your time back.

For many people, the first important outcome of Agentic Power is simply getting their time back.

A physician

stops spending nights on credentialing.

A founder

stops manually assembling reports.

A small-business owner

stops doing repetitive administration.

A marketer

stops manually researching every campaign.

A developer

stops supervising routine implementation.

An operator

stops personally moving information between systems.

The initial goal can be simple: **free the owner's or operator's time while maintaining or improving work quality.**

Reclaimed time can become family, health, rest, creativity, sales, strategy, learning, relationships, higher-value work or simply fewer hours worked. You do not have to reinvest every saved hour into more production.


AI can give you time back.

You choose whether—and where—to reinvest it.

What happens to jobs?

What happens when everyone becomes dramatically more productive? Higher Agentic Power means fewer human hours may be required for many categories of work people are paid to perform today. Some jobs will shrink, disappear, be redesigned, or require fewer people. That disruption is real.

Existing paid work is not the boundary of useful human activity. Nathaniel Whittemore, host of *The AI Daily Brief*, describes a useful distinction between “efficiency AI” and “opportunity AI.” Efficiency AI applies AI to work we already do; opportunity AI pursues work previously too expensive, slow, specialized, labor-intensive, small in market size, or difficult to staff. [Source and attribution.](#)



Efficiency AI reduces the cost of work we already do. Opportunity AI expands the set of work worth doing. The work humans are paid to do today is a bounded set. The opportunity space for new products, services, discoveries, experiences, businesses, and human ambitions is effectively unbounded—not a claim of mathematical infinity or unlimited demand.

The Industrial Revolution replaced vast amounts of human physical effort without eliminating human economic activity: it changed what humans did. Machines lowered task costs and made new industries and forms of work possible. AI may do something similar for cognitive work, but its pace, breadth, and social consequences may differ substantially. [Historical context.](#)

Agentic Power fuels both sides of the transition. Applied to existing work, it creates efficiency; applied to previously uneconomic possibilities, it creates opportunity. As [comparable AI capability costs fall](#), a small company might test hundreds of product concepts instead of three, or build software for a niche internal problem that never justified an engineering team.

Illustrative thought experiment; fictional numbers: an expert analysis costing \$50,000 is uneconomic if the decision it informs is worth only \$5,000. If comparable-quality analysis costs \$50, thousands of previously unjustifiable analyses could become viable. That is opportunity AI—provided the value, quality, and total delivery cost hold up.

AI will likely eliminate some jobs, change many more, and create work that does not exist today. An identity built on “I perform these tasks” is more exposed when AI can do them better, faster, and cheaper. A stronger position is learning [Agentic Engineering](#): identify valuable problems, define goals, exercise judgment, build AI staff, direct teams, evaluate outcomes, discover opportunities, and improve systems. These capabilities can strengthen a person’s position without guaranteeing employment.

As Agentic Power increases, the human role can move up the [AI Staff Delegation Ladder](#): from doing tasks to directing roles, managing outcomes, setting goals and constraints, and deciding what deserves to be done at all. The practical aim is to become someone who can direct AI toward useful outcomes. [Jobs, displacement, and new work: the deeper discussion.](#)

Our work moves upward.

The future of work is not humans becoming faster chatbot users. It is humans learning to direct increasingly capable digital organizations.

Doing → Directing → Designing → Deciding

Agentic Power gives us a way to measure how much productive capacity that transition creates. Leadership determines what we do with it.

“My mission is to increase my own Agentic Power, help others increase theirs as quickly as possible, and direct that power toward good.”

– Dr. Mark Allen

At HeroForge.AI, that means helping people and organizations become better Agentic Engineers, increase their Agentic Power, and direct that capacity toward outcomes that matter.

Agentic Power multiplies execution.

Human judgment determines where that power goes.

Methodology and endnotes.

Measurement rules and limitations

Accepted work and a credible reference

An AP number is only as credible as its Human-Equivalent Hours estimate. Good AP measurement makes the baseline inspectable.

Estimate the effort a competent skilled person or team would reasonably need for the same accepted outcome at the same quality. Record scope, skill assumptions, evidence, uncertainty and deduplication. A human-only deliverable does not enter an AI operation's numerator. Do not count raw output, unused drafts, duplicates, failed work, token volume, agent runtime, internal machine actions or output that humans must recreate.

For mixed work, attribute only the accepted AI-produced portion. Count accepted intermediate outputs only when independently useful and not also counted inside a final deliverable.

All attributable direction

Include setup, training, context, tools, review, failed-attempt correction, coordination, exception handling, evaluation, maintenance and system improvement. Use active person-hours across all human directors, with shared effort reconciled.

Preserve both a representative current window and earlier investment in the full-project view. Do not hide setup costs by selecting a convenient interval. Zero or unknown direction never yields infinity.

Individual and organizational AP

Use the same rule for a workflow, project, individual, team, department or company: total accepted HEH divided by reconciled attributable direction hours. Never average individual AP ratios. Avoid counting overlapping work products or shared human time twice.

AP comparisons are useful when scope, quality, baseline method, available tools and working conditions are sufficiently comparable. They need not be limited to one person over time.

A hiring use case

A standardized evaluation could hold constant the desired outcome, quality threshold, available tools, time window and security / policy constraints. Then record accepted work products, HEH, human direction, AP, elapsed time, judgment and final quality.

This can reveal productive capability and inform rewards. Disclose assistance and inherited assets; examine repeatability and contribution to others' performance. This is a methodology direction, not a recruiting product or a salary formula.

AP's limits

Agentic Power measures human-attention productivity, not compute efficiency; AI and infrastructure cost should be tracked separately.

AP does not measure every dimension of intelligence, organizational health, resilience, strategic judgment, economic value or beyond-human capability. Agentic Power measures how much skilled human-equivalent work an AI operation produces for every hour of human direction.

That is a consequential management measure. Track it alongside intended outcomes, costs, quality and risk. AP alone does not establish ROI or causation.

Evidence, terminology and case limits

HEH estimates what comparable human work would have required; it is not observed replacement labor. The cases use retrospective or production-history estimates; they are not independently audited benchmarks. Report elapsed time and uncertainty with every serious result.

The framework is proposed management terminology, not an externally validated industry standard. Agentic Engineering terminology is evolving rapidly; Graph Engineering in particular remains an emerging term.

Measured:

Direct observation

Estimated:

Stated basis

Derived:

Arithmetic from inputs

Illustrative:

Explicit assumptions

Self-reported:

Subject's claim

Unknown:

Insufficient evidence

[Read the full standalone Audit methodology.](#)

See how the Audit works

Choose a representative window and a consistent quality bar. Inventory accessible evidence, identify accepted outputs, count each once, estimate competent-human effort, and reconcile all attributable human attention. Missing evidence stays visible.

The first snapshot gives a provisional AP with its assumptions and one optional question. The full Evidence Profile adds AP\$ when evidenced, accepted HEH, Human Direction Hours, Equivalent Labor Cost, Direct AI Cost, Human Direction Cost, Total Attributable Agentic Cost, Elapsed Window, Work Products and Intended Outcomes. Human Direction Cost is required for AP\$; when unknown, show "AP\$ pending Human Direction Cost." Diagnostics cover delegation, autonomy, Agentic Engineering, self-improvement, the bottleneck and the Next Power Move.

Use the Quick Audit for a short progressive estimate and the Full Evidence Audit for detailed accounting, evidence collection, diagnostic classifications and public-safe export rules.

Improvement Audit: who owns the improvement process?

S1–S5 describe increasing AI ownership of improvement. S1 is human-led. Any level can use a systematic, continuous process.

S1 Manual	People identify, test and apply improvements.
S2 Assisted	AI suggests or helps with changes; people drive the improvement process.
S3 Managed	AI runs an assigned improvement cycle and proposes tested changes; people approve adoption.
S4 Adaptive	People define the improvement targets or categories; AI tests and adopts changes within them. Higher-risk changes need approval.
S5 Self-Directed	AI identifies and prioritizes what to improve in its skills or system, tests changes, and retains demonstrated improvements within defined permissions.

S4 improves within human-defined targets or categories; S5 also identifies and prioritizes what to improve, guided by human-set goals and constraints. A5 proactively advances the work; S5 directs improvement of the skills or system doing it. Neither level establishes the other. Require evidence of evaluation, retained changes, rollback and escalation. Persistent runtime alone does not establish an improvement status. Keep unknowns unknown; greater ownership alone does not establish higher AP.

Jobs, displacement, and new work

Tasks are not whole jobs

AI can reduce the labor required for existing tasks. Jobs also bundle relationships, judgment, accountability, coordination, context, and decision-making. Automation may remove parts of a job before an entire role disappears; employers may also reorganize the remaining work across fewer people. Some occupations will likely shrink or disappear, while many others change substantially. Task exposure is not an observed job-loss count. The ILO's task-based research helps distinguish these questions. [Sources and limits](#).

How falling costs open new work

The economics layer supplies a mechanism, not an employment forecast. As Agentic Power rises and unit AI capability cost falls, the cost of producing comparable-quality skilled outputs can fall. Human direction, verification, integration, and other attributable costs still matter. Higher AP alone does not establish a lower total cost or a valuable outcome.

When the cost of producing intelligence falls, work that was previously too expensive becomes economical. A service affordable enough to use once a year may become worth using weekly or continuously. Falling AI capability costs do not only threaten existing labor economics. They can create demand for entirely new forms of cognition and work. That demand depends on real usefulness, budgets, trust, and adoption; it is not automatic.

The fictional \$50,000 / \$5,000 / \$50 example in the main text illustrates a cost-benefit threshold, not an observed price, realized saving, or claim that all expert analyses can be automated. Comparable quality and the full cost of delivery must be established. Individualized tutoring, niche internal software, broader scientific hypothesis testing, and a solo entrepreneur serving a previously unstaffable market are possibilities, not measured outcomes.

From efficiency to opportunity to beyond-human scale

This is a conceptual progression, not a new metric, axis, or branded framework:

1. **Efficiency AI:** humans already perform the work; AI lowers its time or cost. For example, drafting a contract faster, subject to appropriate review.
2. **Opportunity AI:** humans could perform the work, but it was previously too expensive or impractical. For example, reviewing every historical contract in a small company once the marginal cost is low enough.
3. **Beyond-human-scale capability:** the same scale of work would never realistically be staffed with humans. For example, continuously evaluating millions of legal, scientific, security, software, or economic scenarios.

The existing [HEH boundary](#) still applies. Estimate reasonable human effort for the accepted, comparable outcome—not every internal machine search. When no defensible human-equivalent baseline exists, describe the capability separately and exclude it from AP. Opportunity does not license invented HEH or hypothetical labor-cost totals.

Opportunity does not erase transition pain

Industrial automation both displaced physical tasks and enabled new economic activity. That history offers a mechanism to examine, not a promise that AI will repeat the same employment trajectory. Cognitive automation may affect a broad range of occupations simultaneously. Timing, wages, bargaining power, geography, access to training, and the distribution of gains remain uncertain.

Aggregate future opportunity does not mean each displaced person can transition easily, quickly, or at comparable pay. New activity can benefit different people and places from those bearing the losses. Building judgment, goal selection, evaluation, and opportunity discovery can strengthen a person's position; it is not employment insurance, and individual adaptation cannot carry the whole social burden.

The Agentic Power Framework makes no macroeconomic employment forecast. The same productive-capacity change can reduce labor demand for existing work and make new categories of work viable. Which force dominates aggregate employment is outside the framework's empirical claim. The purpose here is to understand that underlying transformation and direct increasingly capable AI toward useful outcomes.

When Human-Equivalent Hours stop making sense

Agentic Power works especially well when AI accelerates work competent humans already know how to perform: software, research, analysis, administration, marketing, legal work, design and operations.

AI also enables work humans would never economically perform at the same scale. A system might simulate ten million software architectures overnight, correlate billions of security events, search enormous molecular or materials spaces, connect new papers across thousands of journals or test millions of pricing and supply-chain combinations.

It would be misleading to equate ten million machine experiments with ten million human experiments. No real business would have hired people to perform that internal process.

Measure the accepted human-comparable outcome, not the machine's internal search.

Internal machine activity

10 million architecture searches

Do not convert each search into hypothetical human labor.

Accepted outcome

One expert-quality production architecture

Estimate the reasonable human effort for a comparable accepted architecture.

A security system may inspect billions of events. Where possible, estimate the comparable accepted investigation, diagnosis, remediation plan or incident resolution / prevention work. Do not price the impossible manual reading of every log, or count speculative incidents prevented without evidence.

When even the final outcome has no human analogue

If no credible human-equivalent baseline exists, do not manufacture HEH. Report:

"Beyond-human-scale capability: no defensible HEH baseline."

Describe the outcome separately. Do not convert it into AP or silently include it in an aggregate numerator. A mixed operation can report AP for the defensible human-comparable scope and list the excluded capabilities alongside it.

Agentic Power measures AI-amplified human-equivalent production. Some of AI's most important future value may eventually come from work that was never economically possible for humans to perform at all.

If comparable AI capability continues to become cheaper, this boundary may become increasingly important. It is an open area for future work; this framework does not invent another metric to resolve it.

Origins of the Agentic Power Framework

The phrase 'agentic power' has appeared previously in unrelated academic and AI contexts. The Agentic Power Framework defined here uses the term specifically for skilled human-equivalent output per hour of human direction.

September 6, 2026

The OpenAI spark

OpenAI published "[Research acceleration: The view inside OpenAI.](#)" It reported approximately 3.1 agent-workdays of effort for every human workday in OpenAI Research.

This raised a question for Dr. Mark Allen: what does an agent-workday actually tell us about useful human-equivalent output? Runtime is not necessarily value. Agent count is not necessarily useful work.

September 10, 2026 · Vibecast

The rUv conversation

Dr. Mark Allen discussed the measurement problem with Reuven Cohen, rUv. rUv independently presented experimental "agentic horsepower," called APX, using an exponent-style formulation to compare human and agent output.

Both agreed that a unit for human-versus-agent productive output was needed. Dr. Mark Allen advocated a direct multiplier businesspeople could understand without reversing an exponent. He raised the importance of skilled work the operator could not personally perform and discussed client examples.

Autonomy emerged as a critical multiplier, and rUv demonstrated his persistent autonomous operating model. The Rowena + Emmanuel case uses 400 combined hours and $AP = 55\times$.

The current framework

A direct management measure

Dr. Mark Allen subsequently developed the Agentic Power Framework around $AP = \text{Skilled Human-Equivalent Hours} \div \text{Human Direction Hours}$, together with the AI Staff Delegation Ladder, the Autonomy Levels, Agentic Engineering, Goal-Directed Self-Improvement and the Agentic Power Audit and Profile.

OpenAI did not create or endorse Agentic Power. rUv is not represented as endorsing Dr. Mark Allen's final framework. Dr. Mark Allen did not invent APX. APX and Agentic Power are distinct metrics. The supplied automatic transcript was inspected; it supports the spoken estimate but is not independently verified raw audit evidence. Dr. Mark Allen supplies the Codex-run detail.

A reference you can share

How to cite this framework.

Dr. Mark Allen. “The Agentic Power Framework.” Version 1.5.0, October 2026. HeroForge.AI.
<https://heroforge.ai/agentic-power>.

The Agentic Power Framework

Version 1.5.0 · October 2026

Developed by Dr. Mark Allen at HeroForge.AI

[Explore the shareable framework cards](#)

Permanent publication: [HeroForge.AI/agentic-power](https://heroforge.ai/agentic-power). No DOI or external archival record is claimed. [Version record](#).

Sources, evidence & limits

Agentic Power measures how much skilled human-equivalent work an AI operation produces for every hour of human direction.

OpenAI: the initial measurement question

“Research acceleration: The view inside OpenAI,” September 6, 2026 [<https://openai.com/index/research-acceleration-view-inside-openai/>]. The reported 3.1 agent-workdays per human workday refers to agent effort, not this framework’s accepted human-equivalent output ratio. The source also discusses interpretive limits and research bottlenecks. No OpenAI endorsement is implied.

Horsepower: historical source and analogy

The U.S. Bureau of Standards’ [The Relation of the Horsepower to the Kilowatt, Circular 34, third edition \(1915\)](https://www.govinfo.gov/content/pkg/GOVPUB-C13-11c7b154c53443df43adc93a3bcceda9/pdf/GOVPUB-C13-11c7b154c53443df43adc93a3bcceda9.pdf) [<https://www.govinfo.gov/content/pkg/GOVPUB-C13-11c7b154c53443df43adc93a3bcceda9/pdf/GOVPUB-C13-11c7b154c53443df43adc93a3bcceda9.pdf>], section III, describes Watt’s use of horsepower for steam engines. AP borrows the function of a familiar comparative measure. It is not a physical unit analogous to watts.

Reuven “rUv” Cohen: September 10, 2026 Vibecast discussion

Conversation facts and excerpts were supplied by Dr. Mark Allen. The supplied extracted automatic transcript was inspected. It is not independently verified raw audit evidence; no public recording is claimed. Reuven Cohen, rUv, independently developed experimental agentic horsepower / APX. The discussion helped sharpen Dr. Mark Allen’s subsequent framework. Neither Dr. Mark Allen’s authorship of APX nor rUv’s endorsement of the final framework is implied.

Ajay Jani, MD: retrospective HeroForge.AI estimate

Ajay described saving “hundreds of hours” after approximately 12 hours of attention. This case uses an estimate of 200 Human-Equivalent Hours. Dr. Mark Allen confirms that the work was completed to a high standard. $200 \div 12$ gives $AP \approx 16.7\times$, over approximately two days. These are retrospective estimates, not observed competing skilled labor or a controlled benchmark.

Rowena Chua, MD + Emmanuel Cecilio: retrospective HeroForge.AI / AI-assisted estimate

Dr. Mark Allen supplies the continued-learning and three-day Mastermind retreat context, production status and $R4 \cdot A2 \cdot S1$ classification of their software operation. The retreat was instruction within a longer build, not its entire duration. The supplied central estimate is 22,000 HEH over approximately 110 calendar days, with approximately 400 combined human-direction hours. $22,000 \div 400$ gives $AP = 55\times$.

The original exercise suggested a conventional project could require 18 months, 6–9 developers, and approximately \$1 million. Those were AI-assisted estimates of what comparable human work would have required, not an observed competing project, verified savings, or established ROI. The HEH estimate is neither runtime nor token or code-line count. Creation of clinic software does not establish clinical validation.

Reuven “rUv” Cohen: self-reported output, assumed direction, and self-improvement

The supplied September 10 automatic transcript records rUv verbally reporting approximately 10 million human working days to reproduce 2026 year-to-date code output. Dr. Mark Allen confirms that rUv ran a detailed audit prompt in his own Codex environment. The transcript is not independently verified raw audit evidence. At 8 hours/day, that normalizes to approximately 80 million HEH. January 1 through September 10 is treated as 253 inclusive calendar days.

The central direction assumption is $253 \times 9.5 = 2,403.5$ hours. $80,000,000 \div 2,403.5 \approx 33,285$, displayed as AP $\approx 33,000\times$. At 8–12 active direction hours/day, the estimate is approximately $26,350\times$ – $39,526\times$. This is sensitivity to an assumed denominator, not a statistical confidence interval or a direction timesheet. Neither the output baseline nor its acceptance is independently audited here.

rUv’s own rough approximately 30,000-times-per-day APX comparison, described by him as loose math, is separate. He did not calculate this framework’s AP $\approx 33,000\times$ estimate. Similar scale does not establish equivalence.

The description of a persistent operating system—“Always running, always building, always learning”—comes from the supplied case material. Its mechanisms illustrate system improvement; no component-level causal contribution has been independently measured. [RuFlo Explained](https://github.com/ruvnet/ruflo/blob/main/docs/ruflo-explained.md) [https://github.com/ruvnet/ruflo/blob/main/docs/ruflo-explained.md] and the [MetaHarness guide](https://github.com/ruvnet/ruflo/blob/main/docs/metaharness-user-guide.md) [https://github.com/ruvnet/ruflo/blob/main/docs/metaharness-user-guide.md] provide technical context, not validation of the case estimates.

Agentic Engineering: practice and evolving terminology

OpenAI, “[Harness engineering](https://openai.com/index/harness-engineering/)” [https://openai.com/index/harness-engineering/] (February 11, 2026), provides a practical reference for the environment around coding agents. IBM, “[What is loop engineering?](https://www.ibm.com/think/topics/loop-engineering)” [https://www.ibm.com/think/topics/loop-engineering], discusses iterative agent processes. The stack here is a teaching synthesis. Its boundaries overlap. Loop Engineering is an emerging practice; Graph Engineering is a newer, less standardized descriptive term.

Stanford AI Index 2025 • fixed-capability cost

[2025 AI Index, Research and Development](https://hai.stanford.edu/ai-index/2025-ai-index-report/research-and-development) [https://hai.stanford.edu/ai-index/2025-ai-index-report/research-and-development]: GPT-3.5-equivalent MMLU performance (64.8) cost \$20.00 per million tokens in November 2022 and \$0.07 in October 2024, a reduction greater than 280-fold. The report also describes hardware price-performance costs falling 30% annually and energy efficiency improving 40% annually. These are different measures, not reductions to apply to an organization’s whole AI bill. Report year: 2025; accessed September 18, 2026.

Epoch AI • differing rates and uncertain persistence

[LLM inference prices have fallen rapidly but unequally across tasks](https://epoch.ai/data-insights/llm-inference-price-trends) [https://epoch.ai/data-insights/llm-inference-price-trends], March 12, 2025: fixed-performance declines ranged from $9\times$ to $900\times$ per year across benchmarks; one GPT-4-level PhD science threshold declined about $40\times$ annually. The fastest declines were recent; the rate differs by threshold and might not persist.

[How persistent is the inference cost burden?](https://epoch.ai/gradient-updates/how-persistent-is-the-inference-cost-burden) [https://epoch.ai/gradient-updates/how-persistent-is-the-inference-cost-burden], Jean-Stanislas Denain, February 16, 2026, describes very roughly 5– $10\times$ cheaper access to given capability annually on some tasks. This is an informal author analysis, not necessarily Epoch’s institutional view or a universal law. Both accessed September 18, 2026.

Stanford AI Index 2026 • total spending can rise

[2026 AI Index, Economy](https://hai.stanford.edu/ai-index/2026-ai-index-report/economy) [https://hai.stanford.edu/ai-index/2026-ai-index-report/economy] reports record AI-company compute costs and infrastructure spending. Falling unit costs and increasing total spending can coexist. More agents, longer reasoning, larger context and more workflows are ways usage can grow; this is an economic explanation, not a quantified causal attribution from the report. Accessed September 18, 2026.

A historical provider example • OpenAI, July 2024

[GPT-4o mini: advancing cost-efficient intelligence](https://openai.com/index/gpt-4o-mini-advancing-cost-efficient-intelligence/) [https://openai.com/index/gpt-4o-mini-advancing-cost-efficient-intelligence/], July 18, 2024: launch pricing was \$0.15 per million input tokens and \$0.60 per million output tokens. OpenAI described it as more than 60% cheaper than GPT-3.5 Turbo, with 82% MMLU and stronger academic benchmark performance. These are historical launch prices, not current quotes; this provider example is secondary to Stanford and Epoch evidence. Accessed September 18, 2026.

Economic methodology: costs, views and evidence

Agentic Power Efficiency (AP\$) measures how much benchmark human labor cost is represented by the accepted work produced for every dollar of attributable agentic cost.

The dollar sign in AP\$ represents economic cost, not U.S. currency. AP\$ can be calculated in any currency as long as numerator and denominator use the same currency and consistent cost basis.

For example: USD, EUR, GBP, CAD, AUD or JPY. Equivalent Labor Cost and Total Attributable Agentic Cost must use the same currency. If currencies differ, convert to one common currency before calculating AP\$. Record the currency, conversion basis and conversion date where applicable; never silently combine currencies.

Total Attributable Agentic Cost = Direct AI Cost + Human Direction Cost + other clearly attributable agentic operating costs. Direct AI costs include API inference, token usage, cloud compute, GPU rental, AI subscriptions, agent platforms, paid connectors, search/tool-call charges and attributable infrastructure. Exclude general company overhead unless attribution is defensible; do not count a charge twice.

Human Direction Cost may be optional for a general Agentic Power Audit. It is required for Agentic Power Efficiency, using known cost or a defensible estimate. If it is unknown, display **"AP\$ pending Human Direction Cost."** Never divide by Direct AI Cost alone.

Current Operating view

Match accepted output, direct AI cost, human direction cost and other attributable operating costs to the selected operating window. This reflects current operating performance.

Fully Loaded / Lifecycle view

Also include attributable setup, implementation, training, failed attempts, initial integration, workflow design, maintenance, system improvement and significant historical build effort. Match output and costs to a defensible lifecycle scope. Do not force historical investment into one short current-performance window or count maintenance twice. Current-performance views show what the mature system does now; full-project views preserve what it cost to get there.

Show Current Operating Agentic Power Efficiency and Fully Loaded Agentic Power Efficiency only where each is defensible. Do not cherry-pick a tiny high-performing window after substantial setup and present it as the total economic story.

Labor-rate evidence

Use the most specific defensible labor-rate benchmark available.

1. **Actual organization cost:** actual fully loaded cost or defensible internal rate for the person or role.
2. **Authoritative occupation + geography benchmark:** for example, BLS occupation + metro/state data.
3. **Credible professional market rate:** for example, legal billing benchmarks or specialized contractor rates.
4. **User-supplied custom rate:** allowed when better evidence is unavailable; explicitly label it.

Do not cherry-pick the highest rate to increase Equivalent Labor Cost. Every rigorous Agentic Power Efficiency result must state labor market, rate basis, source, source date and currency. Keep currencies and incompatible rate bases separate.

AP\$ scenario range

Low = lowest defensible Equivalent Labor Cost ÷ highest defensible Total Attributable Agentic Cost. Central = base labor cost ÷ base attributable cost. High = highest defensible labor cost ÷ lowest defensible attributable cost. Label low / central / high as a **AP\$ scenario range**, never a statistical confidence interval. Missing costs remain pending; a zero denominator does not produce a finite ratio.

Estimated Cost Differential = Equivalent Labor Cost – Total Attributable Agentic Cost. Estimated Cost Differential compares two cost models for comparable work. It is not necessarily realized savings. Claim realized savings or cost avoidance only with evidence that the cost would otherwise have been incurred. Cost per HEH and other derived calculations belong in Audit diagnostics, not additional headline ratios.

1. **Agentic Power** – How much skilled work can one hour of human attention command?
2. **Skilled Human-Equivalent Work** – What comparable work was produced?
3. **Equivalent Labor Cost** – What would comparable skilled-human capacity reasonably cost?
4. **Total Attributable Agentic Cost** – What did the agentic operation actually cost?
5. **AP\$** – How much comparable human labor cost can one dollar of agentic cost command?
6. **Business Outcomes** – Was that power directed somewhere worthwhile?

Labor benchmarks • market, basis, currency and date

The Audit asks for a market and currency before translating HEH into money. Its preferred general basis is **fully loaded employer replacement cost** where reliable evidence exists. Alternatives are **market purchase rate** and **custom internal rate**. Each component records its basis, occupation, skill level, source, source date and low/base/high assumptions. A blend must be explicit; billing rates are not employee compensation.

BLS Occupational Employment and Wage Statistics [<https://www.bls.gov/oes/>] provides occupation wages nationally, by state and metropolitan area. The latest available release checked here covers May 2025, released May 15, 2026. This site does not claim an automatic OEWS lookup or a fetched wage for every occupation.

BLS Employer Costs for Employee Compensation [<https://www.bls.gov/news.release/ecec.nr0.htm>], June 2026 period, released September 9, 2026: private-industry total compensation averaged \$46.89/hour, comprising \$32.82 wages and \$14.07 benefits (wages 70.0%). Salary alone understates employer cost. The general total-to-wage ratio, about 1.43, may be used only as an explicitly labeled estimate when better evidence is unavailable. [Benefit ratios vary by occupation and industry](https://www.bls.gov/news.release/ecec.t04.htm) [<https://www.bls.gov/news.release/ecec.t04.htm>]; do not automatically apply one factor to every occupation.

Clio's 2025 legal billing data [<https://www.clio.com/resources/legal-trends/compare-lawyer-rates/>], page updated March 2026: average U.S. lawyer rate \$349/hour; the detailed table labels \$461/hour "Corporate Litigation" (its takeaway abbreviates this as corporate). [Texas's 2025 lawyer average](https://www.clio.com/resources/legal-trends/compare-lawyer-rates/tx/) [<https://www.clio.com/resources/legal-trends/compare-lawyer-rates/tx/>] was \$366/hour. These historical USD market-purchase benchmarks vary by practice and place and are not employer costs.

For other countries, use credible national statistics, professional datasets or defensible user-supplied rates. No international rate is invented. Keep currencies separate unless an explicit dated exchange-rate method is documented; no automatic purchasing-power-parity adjustment. Low/base/high are scenarios, not confidence intervals.

All sources above accessed September 18, 2026. [Dated benchmark records](https://heroforge.ai/agentic-power/economic-benchmarks.json) [<https://heroforge.ai/agentic-power/economic-benchmarks.json>]. Match a benchmark to the work before using it; presence in this reference list is not a valid skill allocation.

Nathaniel Whittemore / The AI Daily Brief • efficiency AI and opportunity AI

Nathaniel Whittemore, host of The AI Daily Brief, uses the distinction between "efficiency AI" and "opportunity AI." The show's own newsletter, "AI's Great Divergence," April 16, 2026 [<https://aidailybrief.beehiiv.com/p/ai-s-great-divergence>], explicitly uses both terms and contrasts reducing inputs for existing output with pursuing growth, new activities, and business models. Its [official episode link](https://open.spotify.com/episode/7qsaRQmIHRwP08FhrK21) [<https://open.spotify.com/episode/7qsaRQmIHRwP08FhrK21>] identifies the accompanying podcast. Newsletter text inspected September 18, 2026; the [official show listing](https://podcasts.apple.com/us/podcast/the-ai-daily-brief-artificial-intelligence-news/id1680633614) [<https://podcasts.apple.com/us/podcast/the-ai-daily-brief-artificial-intelligence-news/id1680633614>] also identifies Whittemore as host and creator. Audio not independently transcribed. The definitions, economic examples, and Agentic Power connections here are this paper's interpretation, not quotations from Whittemore. The terms are attributed to him; no endorsement of this framework is implied.

Jobs and technological change • research context, not an employment forecast

ILO and NASK, Generative AI and Jobs: A Refined Global Index of Occupational Exposure (2025)

[<https://www.ilo.org/publications/generative-ai-and-jobs-refined-global-index-occupational-exposure>]: the publication summary describes task-level exposure and job transformation. Exposure is potential, not observed displacement. David H. Autor, "Why Are There Still So Many Jobs? The History and Future of Workplace Automation," Journal of Economic Perspectives 29(3), 2015, pp. 3–30 [<https://www.aeaweb.org/articles?id=10.1257/jep.29.3.4>]: the abstract discusses substitution, complementarity, demand, and changing job composition. Publication summaries/abstract inspected September 18, 2026; no claim of full-paper review. Neither source establishes which employment effect will dominate future AI adoption. The economic thought experiment and opportunity examples in this framework are illustrative.