

The Capability-Centric Workforce

Computing human capital from enterprise abilities.

Human capital defined by enterprise abilities can be understood, allocated, developed and anticipated. **No other method does all four.**

Most organisations model the workforce through jobs, skills inventories, processes or value streams. None is a stable object. Skill repositories are fifty-year accumulations, unnormalised, authored by whoever held the pen. Processes and value streams are specific to a line of business, a product, a geography, a use case. **The enterprise ability is the only representation abstracted from a scoped, context-specific sequence of events.**

The evidence. Commercial new business. 569 authored enterprise abilities, 36 named participants, each scored on business value, AI classification, the four AI dispositions and retained human authority.

What the computation demonstrates

-
- 01 **Competency is measurable only when attributed against abilities.**
Product, geography and jurisdiction are attributes of the ability, not competencies in their own right. No second taxonomy is required, and none can be maintained.
-
- 02 **Opportunity is how competency comes into existence.**
Competency develops through executing abilities. When AI changes who executes them, it changes who can develop competency at all. Workforce transformation reaches tomorrow's workforce, not only today's.
-
- 03 **The dispositions move independently.**
The underwriter is augmented on 43 percent of its footprint and absorbed on 20. The support service rep is the least augmented, holds the least authority, and is the most accelerated at 89 percent. A different future, not a milder one.
-
- 04 **The workforce extends beyond the payroll.**
Twenty of the thirty-six named participants belong to the carrier. Effort pushed onto a broker or policyholder is transfer, not efficiency.
-

Once competency and authority are measured against it, **human capital is computed from enterprise abilities rather than inferred from jobs.**

Management decides. The model computes.

A workforce cannot be modeled as a set of jobs. It is the set of participants mapped to the abilities they can perform and may authorize.

Carlson BCM has always carried two authored methods. One models the ability. The other models the competency to perform it. They are not adjacent models. Domain competency is measured against the enterprise's own abilities, which is what turns workforce transformation from a narrative into a computation.

Purpose

Every serious account of AI and the workforce now says the same three things. Routine execution moves to the machine. Human work concentrates on judgment. Accountability has to be redesigned. All three are asserted. None is computed.

This paper establishes what makes them computable and demonstrates the computation on commercial new business. It is a theory-of-the-firm argument, not an engineering one. **It ends where execution begins.**

A note on chronology

The argument is not that AI created a need and a framework was built to meet it. The competency method described here was developed and run in enterprise programs years before the current AI cycle, with assessed populations, manager and self scoring, and integration into workforce management.

What AI has changed is not the method. It is the cost of not having one.

THE OBJECT

Business capability. Structured, relational, and semantic.

THE TWO AXES

Competency is **Can**. Authority is **May**. Both relational to the capability.

THE OUTPUT

Computed workforce impact, before implementation.

Two dimensions, measured and computed separately.

Carlson BCM carries two authored methods. One models what the enterprise can do. The other models who can do it. **The two are measured and computed separately.**

The model classifies dimensions of development, not people.

Domain competency

Mastery of the enterprise's operational abilities, expressed through the capability model. Domain means a domain of the enterprise: product management, new business, policy management, billing and remittance, claims management.

The enterprise models its abilities as capabilities. The individual is assessed against those same abilities.

That assessment is made at Level 3, what the business does, rather than Level 4, how it is currently done. The how matters, but it is conventional. It changes when the platform changes and when the process is redesigned. **Tracking competency there produces a model too large to maintain and too volatile to trust.**

Core competency

How the professional works, not what the enterprise does.

An operations leader holds workflow and process design, exception management, capacity planning, and quality management. An actuary holds modelling and forecasting, assumption setting and defence, and model validation and monitoring. An architect holds business analysis, systems analysis and design, and solution assessment.

The model prescribes neither list. It supplies the structure and the test. **A firm adopting this authors its own definition of competence and inherits the rigour.**

Product is not a competency. It is an attribute of one.

Every skills programme in the market treats product knowledge as a competency in its own right. Does this person know inland marine, workers compensation, general liability. That question forces a second taxonomy into existence, because product knowledge has no natural unit and no defensible boundary. It is why skills inventories decay into self-declared opinion.

Product was never a competency. **It is an attribute of the capability**, alongside business value, AI classification, and retained authority.

So the assertion is never that someone understands workers compensation. It is that they understand workers compensation **in the context of a capability**. Adjudicating a claim in workers compensation. Adjudicating a claim in general liability. Same capability, different product, and a person may hold one and not the other.

GENERAL LIABILITY

Ability to pre-qualify a submission. Held at authoritative competency.

COMMERCIAL AUTO

Ability to pre-qualify a submission. Held at foundation competency.

Two people who both say they know commercial auto hold different ability sets, and the model shows precisely where each of them stops.

Competency is contextualised. It is held against a capability in the context of geography, coverage, jurisdiction and channel.

What a person knows about a product is always seen through the lens of the capability.

Human capital becomes a managed asset.

Because competency is now captured in normalised form across both domain and core, four things become possible that were not possible before.

UNDERSTAND

Who holds which abilities, at what level, and where the depth actually sits.

ALLOCATE

Deploy people against the work by measured fit rather than by availability.

DEVELOP

Target education and assignment where the gap is, not where the budget was last year.

ANTICIPATE


See what AI does to the workforce before any of it is built.

Show me everyone who can:

- adjudicate a claim
- adjudicate a commercial auto liability claim
- adjudicate a commercial auto liability claim in Japan

Each question is answered from the same model. Nothing is added to it, and nothing is rebuilt.

A competency assessment here is not a form. It is a query.



Each point is a person. Each colour is a function of the enterprise. The arcs are competency bands, and progression runs outward through them. The dashed line is balanced development, where most of a population sits.

Because the attribution is relational, this is a rendering of a selection rather than a fixed picture. Narrow the domain axis to one product line or one geography, or the core axis to one profession, and **the same model returns a different population and a different shape.**

What actually broke.

An emerging claim holds that expertise is becoming harder to develop, because the routine work junior people once did is being automated away. The observation is correct. The model of development behind it is not.

Routine work never built judgment. It was simply how the enterprise happened to deliver the three things that do. **Of those three, automation removes only one.**

FORMAL EDUCATION

Survives untouched, and matters more as incidental learning disappears.

MENTORING

Survives, and matters more for the same reason.

OPPORTUNITY

The leg under pressure, and it fails in three separate ways.

Formal education runs on two planes at once

The enterprise plane. Assess competency across a team. Decide which competencies matter most to that team's future. Compare the assessment against that priority. Where the gap sits on what matters most, build programs against it. This is directed education, aimed at a target the enterprise has chosen.

The individual plane. The same person assesses themselves against the same model, sees where they stand, and writes their own plan. What do I want to build. What am I drawn to. Where do I want this career to go, whatever the enterprise has prioritized this year. The individual has a say and directs their own learning.

Both planes run at once without conflict because **they sit on the same coordinates.** A person's chosen path and a company program can be measured against each other instead of being two separate exercises that never meet.

Corporate learning is directed downward. Personal development plans sit beside it with no structure at all. This is the only construction in which both are measured, on the same axes.

Mentoring, and the question of certification.

Mentoring is essential to developing competency, and effective mentoring brings expertise, experience and judgment to it. The mentor's questions are simple. You have chosen this profession, so what do I have to teach you. How do I make your time here worth something. How do I read your plan and help you get there. How do I take what I know and apply it to what you will need next.

This has a long tradition. The trades ran apprentice, journeyman, master for centuries. The apprentice was learning the work under someone else's authority. The journeyman had the methods and could work alone but could not certify work or train anyone. The master held authority, certified the work, and was the only one allowed to teach. **The right to mentor was itself an authority.**

What matters is how that ladder ended. No technology took it. It ended when volume production removed the need for the progression, which is the same substitution now under discussion.

Certification

Certification is credentialing, not development. The model treats it as two separate things.

AS A CONDITION ON AUTHORITY

It belongs in the model. Where a profession or a regulator requires a credential, a person without it cannot act regardless of how capable they are. An actuary who cannot sign an opinion and a representative who cannot transact are facts about authority.

AS EVIDENCE OF COMPETENCY

It earns no place. It shows that a threshold was cleared once, against a syllabus, at a moment. It says nothing about where a person stands on either axis today, which is why no organization can staff from a list of certifications.

Certification is an input to formal education and a condition on authority. It is not a substitute for assessment.

Opportunity is the leg that actually fails.

Formal education survives and matters more, because the incidental learning that used to come with the work is going away. Mentoring survives and matters more for the same reason. Opportunity is the leg under pressure, and it fails in three separate ways that usually get discussed as one.

ONE · THE WORK GOES FIRST

The work that gave people their first opportunities is exactly the work AI absorbs, so the bottom of the ladder goes first. The capability model can name which abilities those are, rather than describing a general erosion. That turns a complaint into a list of what has to be replaced on purpose.

TWO · THE UNKEPT OBLIGATION

A company that pays for education and never gives people the chance to use it has bought nothing. Learning that is not applied fades. Today education is bought centrally while application is left to whatever assignments happen to come along, which is the same hope-based approach this method rejects everywhere else.


THREE · THE CLOSED LOOP

Opportunity goes to whoever is already trusted with it, by proximity and by pedigree, not by who needs it to grow. You cannot get the opportunity without the competency, and you cannot build the competency without the opportunity. That loop is why most talent programs produce so little movement.

The Carlson BCM position is that a company owes opportunity against the individual's plan and not only against its own priority for the year. Once assignments sit on capability coordinates, **who gets which opportunity becomes visible, and therefore arguable.** It can be seen who was never given the ability they asked for.

Opportunity stops being a volume of work and becomes an assignment. A gap is a named set of abilities. Closing it means exposure to those abilities, chosen rather than waited for. Mentoring pairs on the same coordinates, matching someone learning an ability to someone who holds authority over it.

What moves, in what order.



Authority moves as execution enters the machine. Workforce Capacity Intelligence computes what that does to the work. Competency Transformation moves each person to the target.

The capabilities that lose authority retire the old competencies. The capabilities that keep it elevate the new ones. This is why the redistribution, and not the automation, is what transforms a workforce. Lost authority absorbs the role. Partial authority operates it. Retained authority either hands the rule to the machine and elevates the role, or reserves it for a human and preserves it.

Who is in the workforce

The Carlson BCM commercial property and casualty model names thirty-six participants. Twenty belong to the carrier. Fifteen do not.



A workforce model scoped to employees is measuring a fragment. Effort pushed off a carrier's staff onto a broker or a policyholder is not efficiency. **It is transfer, and only a model naming both sides can see it happen.**

The finding is not that AI changes the work. It is that AI changes the indexed competency composition of the workforce in measurable, role-specific ways.

Commercial new business, 569 authored capabilities. Each participant carries a domain competency footprint, the abilities attributed to it. Every ability is scored on the four AI dispositions and on retained human authority. Three participants carry the work of this value chain.

PARTICIPANT	ABILITIES	ABSORB	AUGMENT	ACCELERATE	AUTONOMIZE	AUTHORITY
Commercial Underwriter	403	20%	43%	77%	16%	22%
Commercial Account Technician	385	18%	40%	78%	15%	21%
Commercial Support Service Rep	56	12%	18%	89%	9%	11%

Share of each participant's footprint scored high on each disposition, and the share carrying high retained human authority.

These are not points on a scale

The Support Service Rep is the clearest case. It is the least augmented, holds the least authority, and is the most accelerated of the three. **A participant can hold little authority, be barely augmented, and still be radically faster.**

That is a different future than the underwriter faces, and it is not a milder version of the same one. Absorb, augment, accelerate and autonomize are independent dispositions. The paper is not measuring how much AI affects a role. It is measuring **which dispositions a role receives.**

What the table says.

Judgment is augmented, not absorbed

The underwriter is augmented on forty-three percent of its footprint and absorbed on twenty. The account technician is close behind at forty against eighteen. The work of the people who carry new business is being assisted rather than taken, and the assistance concentrates on the parts that require them.

Authority stays with the participants who hold it

Around a fifth of the underwriter's and account technician's abilities carry high retained human authority. That share does not move to the machine. It is the part of the work where someone has to answer for the outcome, and it stays where it was.

Acceleration is nearly universal, and it is the least structural

Acceleration runs from seventy-seven to eighty-nine percent across all three participants. Almost everything gets faster. Far less is augmented, and less still is absorbed or autonomized.

Acceleration is the loudest disposition and the least structural one.

A participant may keep a fraction of its execution and nearly all of its authority. That is not a smaller version of the same role. It is a different construct.

That contradicts the assumption most operating decisions rest on, which is that less execution means less importance. It does not. Execution and authority are separately attributed and they move separately. **A role can shrink operationally and become more consequential at the same time,** because what remains is the part someone has to answer for.

What this does to development

Mentoring needs someone who holds authority over the ability being learned. Authority concentrates in the participants that carry the judgment work, which means **the ability to mentor concentrates there too.** Development everywhere else has to be sourced deliberately, through education and assigned opportunity, because it will not arrive on its own. That is a computed constraint on the apprenticeship model rather than an assertion about it.

Seven claims the paper makes and defends.

- 01 **Capability and competency are the two authored methods.**
Measured and computed separately, not two models that reference each other.

- 02 **The capability model is the reference.**
Domain competency is measured against the abilities the enterprise has already authored.

- 03 **The ability set already exists.**
An enterprise that has authored its capability model has already authored what domain competency is measured against. There is no second inventory to build.

- 04 **Competency is Can. Authority is May.**
Both relational to the capability. A workforce is the set of participants measured on those two axes.

- 05 **Workforce effect is not authored. It is computed.**
From the intersection of independently authored business value, AI classification, and retained human authority.

- 06 **The dispositions are independent.**
Absorb, augment, accelerate and autonomize move separately. A role can be barely augmented and heavily accelerated.

- 07 **Development has three legs and only one is under threat.**
Education and mentoring survive. Opportunity fails three ways, and indexing makes it designable and its allocation arguable.

The future organization is not the current job architecture with fewer people in it. It is a different distribution of abilities and authority.

What this does not address

This is design-time attribution, not execution-time admissibility. Whether a given action is permitted at the moment an agent attempts it is a runtime question, and it belongs to the execution paper. What the model supplies to that layer is the resolvable identity it has to decide against.

What is measured.

569

authored capabilities in commercial new business, 78 at Level 3 and 491 at Level 4. The Level 3 set is the population for all capability-side findings.

23,669

attribute values across 143 attributes, spanning the strategic, operational, and technology architecture

36

named participants, twenty within the carrier

Competency evidence

Assessed enterprise programs using self and manager scoring. In one population, eighty-five individuals across two hundred thirty-one items and thirteen competency areas, with domain measured as coverage of the enterprise value chains. The two axes correlate at $r = 0.71$.

Discipline

- **One corpus, one paper.** Every capability figure here is drawn from commercial new business.
- **Relational computation.** Every measure reported here can be recomputed against any attributed subset of the enterprise abilities without altering the competency model itself. Only the predicate changes.
- **Ordinal stays ordinal.** Effects are reported as shares of a footprint, never as projected counts. Converting an ordinal score into a quantity would require a scale the model does not contain.
- **Magnitude is out of scope.** Hours, headcount, and cost require enterprise-specific parameters. No client is identified, in any form.

What the capability model **becomes**.

A capability model has always been treated as architecture, maintained beside the work rather than inside it. **Once competency and authority are measured against it, it becomes the computational substrate through which enterprise structure, AI deployment, authority, and workforce transformation are all analysed.**

*Management decides. **The model computes.***