WHITE PAPER WORKFORCE GOVERNANCE FOR ENTERPRISES

Measuring AI Use at Work Without Monitoring Individuals

A company can learn where its AI spend fits the work without keeping a score on any one employee. This paper sets out the design choices that make that possible, and the rules in the European Union and four US states that apply when a measure is tied to a person.

By Mahesh Tatiparthy · November 11, 2026 · 23 min read

Much of what a company spends on AI is counted in tokens, the unit AI vendors bill by, which is roughly a piece of a word. Once the bill is large, finance and HR leaders want to know where the tokens went and whether the spend made sense. Some of the usage reports that answer the first question list each user, as Anthropic's analytics documentation shows, and that makes it easy to drift from a question about spend into a record of how each employee works.

This paper describes how to answer the spend question at the level of roles and teams. It covers why that level is the right one, how large a group should be before it appears in a report, what employees should be told, and which rules apply in the European Union, Germany, California, Colorado, Illinois and New York. We read each rule in its official text on November 11, 2026, and the texts are linked where they are discussed. This is general information and not legal advice, and a company should confirm with its own counsel how any rule applies to its design.

KEY TAKEAWAYS

- A token count shows how much text passed through a model. Amazon's engineering leadership wrote in October 2026 that input measures such as token usage are "deeply flawed". We go a step further in this paper: the count should not be used to rate a person.
- Whether AI spend sits in the right roles is a question about groups, so the answer does not need a report on any one employee.
- A group report should not show groups below a minimum size. Microsoft's Viva Insights documentation says its own minimum group size has to be at least five.
- The EU AI Act lists AI systems intended to monitor and evaluate worker performance among its high-risk uses, and its duties for those systems apply from December 2, 2027 under Regulation (EU) 2026/1744.
- California's privacy regulations give businesses that already use automated decisionmaking technology for decisions about work and pay until January 1, 2027 to comply.

Why measure roles and teams?

The measurement this paper has in mind joins three feeds for each role. The first is token use, read from the usage reports of the AI tools a company pays for. The second is the role's tasks and how much of that work AI can do. The

third is the output measure the team already tracks in the system where the work happens. Placing token use beside the share of a role's work that AI can do gives a simple map. A role with high token use and few tasks that suit AI is spend to question. A role with low use and many suitable tasks is a candidate for training and rollout before others. Where token use and suitable work are both high, the spend is in a sensible place, and the remaining checks are whether the team's output improved and whether the model in use is larger than the tasks need.

Each measure on that map describes a role. How much of an accountant's work suits AI is an estimate for the occupation, to be checked against how the job is done in a given company, and the output measure is chosen for the role as well. The decisions the map supports are group decisions too, such as where to send training and which teams could use a smaller model. Learning that one analyst used more tokens than another adds little to those decisions, because the count says nothing about the quality of the analyst's work.

A count can also be raised without better work. In April 2026, Fortune reported that an employee-built leaderboard at Meta ranked the top 250 token users in the company and came down two days after it was reported, and Meta told the magazine that the employee removed it at their own discretion. In May, Fortune relayed a Financial Times report that some Amazon employees ran an internal AI tool on unimportant tasks to raise their counts on a usage leaderboard. Amazon had not responded to Fortune when that article ran. Amazon's own engineering leadership later put the problem on record, writing on October 27, 2026 that such input metrics are "deeply flawed" because they are disconnected from what matters.

Token tracking is already common, which is why the design question matters now. In a Deloitte survey of 3,500 leaders in 25 countries, fielded in May and June 2026 and released on October 29, 51% said their company tracks token volume.

A TOKEN COUNT IS NOT A PERFORMANCE MEASURE

A count of tokens records volume, so it cannot tell a careful prompt from a wasteful one and it cannot see work done well without AI. Ranking people by it can encourage them to use more tokens, as the reports above describe, and it tells a company little about value. The measures in this paper are for roles and teams.

What does the usage data contain before it is grouped?

Group reporting is a choice a company makes, because the raw data often arrives per person. Anthropic's documentation for its analytics APIs describes "per-user daily metrics" for Claude Enterprise organizations, along with per-user token usage and cost on usage-based plans. Microsoft's Graph report on Microsoft 365 Copilot usage lists a user principal name, which is the account sign-in name and often looks like an email address, together with a display name and the last activity date in each Copilot app.

When a source supplies records per person, the method handles identified data at one point, the moment a usage record is matched to a role. A sound design makes that match once, keeps the role and team, and drops the name and email before anything is stored for reporting. It also keeps what is collected to a minimum. The map needs token counts for each team, by model and by period. It has no use for the text of prompts or answers, and a design that never collects that text has less to protect and less to explain.

How large should a group be before it is reported?

A group report protects individuals only when the group is large enough that no member can be picked out. In a team of two, each member can subtract their own use and see the other's. Small groups also give poor measurements, since one heavy user can move the figure for the whole team.

None of the rules in this paper sets a single number. A public reference point comes from Microsoft, whose Viva Insights documentation says the minimum group size for its organization insights has to be at least five. The same page gives the reason: it is easier to guess information about an individual from results about a smaller group. An administrator can raise that threshold and cannot set it below five.

The practices below are ours, and they are not drawn from any statute. The first is to set one threshold with HR, privacy counsel and worker representatives, and to write it down. Microsoft's floor of five is a number others have used, and a higher one may suit a sensitive output measure. The rest follow from that threshold.

- Apply the threshold to every breakdown. A large role can turn into very small groups once it is split by location and then by level.
- Combine a small role with its job family when it falls under the threshold, and show the family.
- Watch for subtraction. If a report shows a department total and all but one of its teams, the hidden team can be worked out, so hide a second team or show only the total.
- Keep the same rule for exports and for the people who run the system. A threshold that applies to the dashboard and is skipped in the download gives little protection.

Thresholds reduce the risk of singling someone out, and the section on the GDPR explains why they do not make the data anonymous in the legal sense.

What should employees be told?

Employees should hear about the measurement from the company before it starts. Several of the rules in this paper require a notice of some kind, and they differ on timing and content. A plain explanation that covers the following points is a reasonable base to adapt for each jurisdiction with counsel.

- What is collected: for example token counts by model, team and period, and whether prompt text is collected at all.
- The level of reporting: roles and teams above a stated minimum size.
- What it is used for: deciding where to train and which tools and model sizes fit the work.
- What it will not be used for: rating or ranking any person, or decisions on a person's pay, promotion, discipline or termination.
- Who can see the reports, how long the data is kept, and whom to ask.

The fourth point carries the most weight, and it has to stay true. If a company later wants to use the data in decisions about individuals, that is a new purpose. The company would then need to reassess its legal basis and its notice duties in each jurisdiction, and several of the rules below give the employee added rights once an automated output feeds a decision about them.

Which rules apply when AI use is measured at

work?

JURISDICTION	RULE	WHAT IT COVERS	DATE
European Union	AI Act, Annex III point 4(b) and Article 26(7), as amended by Regulation (EU) 2026/1744	AI systems intended to monitor and evaluate the performance and behaviour of people at work. Employers inform workers and their representatives before use.	Applies from December 2, 2027
European Union	GDPR, Articles 5, 6, 13, 22, 35 and 88	Any processing of employees' personal data, including usage records tied to a person	In force
Germany	Works Constitution Act, section 87(1) number 6	Works council co-determination on technical devices designed to monitor employee behaviour or performance	In force
California	Privacy regulations, sections 7200 to 7222	Automated decisionmaking technology used for significant decisions, including work allocation, pay, promotion and termination	Comply by January 1, 2027
California	SB 947, Chapter 859	Automated decision systems used in discipline or termination	Operative July 1, 2027
Colorado	SB 26-189	Automated decision-making technology that materially influences a consequential decision, including employment	Takes effect January 1, 2027

JURISDICTION	RULE	WHAT IT COVERS	DATE
Illinois	Public Act 103-0804	AI used in employment decisions: no discriminatory effect, and notice to employees	In force since January 1, 2026
New York	Civil Rights Law section 52-c	Written notice to employees whose telephone, email or internet use is monitored	In force

Rules that bear on measuring AI use at work, with status as shown on each official source on November 11, 2026. General information, not legal advice.

Source: Official legislative and regulatory texts linked in each row

The dates in this table agree with our 2027 AI and employment compliance calendar, which also covers the rules on hiring tools and layoff notices that this paper leaves out. The sections that follow take the table row by row, starting with the European Union.

What does the EU AI Act say about monitoring worker performance?

The AI Act sorts AI systems by risk, and Annex III lists the uses it treats as high-risk. Point 4(b) covers AI systems intended to be used for decisions on the terms of work relationships, promotion and termination, for allocating tasks based on individual behaviour or traits, and "to monitor and evaluate the performance and behaviour of persons" in those relationships. A system that scores individual employees on how they use AI, and feeds that score into reviews, is close to that description. Article 6(3) allows exceptions for listed systems that pose no significant risk, and it adds that a listed system that profiles people is always high-risk.

The high-risk duties for Annex III systems were due to apply on August 2, 2026. Regulation (EU) 2026/1744, adopted on July 8, 2026 and published on July 24, replaced that date with December 2, 2027. From then, an employer that uses such a system has duties as a deployer, the Act's word for the organization using the system. Article 26(7) requires it to inform workers' representatives and the affected workers before putting the system into use at the workplace. Article 26(2) requires it to assign human oversight to people with the competence, training and authority to carry it out. The amending regulation also rewrites Article 111(2), so that a system already in service before the application date is in general covered only once its design changes significantly. Under Article 99(4) of the Act, breaches of a deployer's Article 26 duties can be fined up to EUR 15,000,000 or, for a company, up to 3% of worldwide annual turnover if that is higher, and Article 99(6) sets lower ceilings for small and medium-sized enterprises.

Article 86(1) of the Act is in a different chapter from the postponed duties, and counsel should confirm when it starts to bite for a given system. It gives a person the right to ask the deployer for a clear explanation of the system's role when a decision based on the output of an Annex III high-risk system has legal or similarly significant effects on them and they consider that it harms their health, safety or fundamental rights.

Article 2(11) of the Act leaves Member States free to keep or adopt laws that are more favourable to workers on the use of AI by employers, so national rules can go further than the Act.

Point 4(b) is written around persons. Our reading, which is not a legal conclusion, is that a report showing only groups, used by no one to evaluate a person, sits further from that wording than a per-person score does. The definition of an AI system in Article 3(1) of the Act also turns on whether the system infers its outputs from its inputs, so a tool that only adds up counts raises a separate question. Both points depend on the intended purpose and

the design of a specific system, and both belong with counsel before anything is switched on in the EU.

What do the GDPR and German works council law add?

Where the GDPR applies, it governs how a company handles personal data, which Article 4(1) defines as any information relating to an identified or identifiable person. A usage record with an email address on it is personal data. Grouping the records later does not change the fact that the company processed personal data to build the groups, so the GDPR's duties attach to that step.

Several articles of the GDPR shape the design. Article 5(1)(b) requires that data be collected for specified purposes and not used in ways incompatible with them, which is why the list of things the data will not be used for has to hold. Article 5(1)(c) limits data to what is necessary for the purpose, which supports collecting counts and leaving prompt text alone. Article 13(1)(c) requires the company to tell people the purposes and the legal basis. Article 35(1) requires an impact assessment before processing that is likely to result in a high risk to people's rights, and Article 22(1) gives a person the right not to be subject to a decision based solely on automated processing that significantly affects them, subject to the exceptions in Article 22(2). Fines for breaching the basic principles or people's rights can reach EUR 20,000,000 or 4% of worldwide annual turnover under Article 83(5).

Consent is rarely the answer at work. The EU data protection authorities, in their 2017 opinion on data processing at work, wrote that consent is "highly unlikely to be a legal basis for data processing at work" unless employees can refuse without adverse consequence. The same opinion says that an employer relying on legitimate interest, the basis in Article 6(1)(f), must choose a method that is necessary, proportionate and as unintrusive as

possible. Measuring groups when groups answer the question supports that argument. It does not complete the balancing test, which each employer has to carry out and document for itself.

Two limits apply to claims about anonymity. Recital 26 of the GDPR says the regulation does not apply to information from which a person is no longer identifiable, and it also says that pseudonymised data, which can be linked back to a person with additional information, is still personal data. Hashing replaces an employee ID with a code derived from it. Where the code can still be linked back to a person, the result is pseudonymised data in the sense of Article 4(5), and it stays inside the GDPR. How much protection hashing gives depends on how it is done and on who holds the means to reverse it. The 2017 opinion adds, in an example about health data, that an employer with more than a thousand employees could still single out individuals by combining data sets. A group report is safer than a personal one, and it should still be described accurately.

Article 88 of the GDPR lets each Member State set more specific rules for employee data, with particular regard to “monitoring systems at the work place”. Germany is one example. Section 87(1) number 6 of the Works Constitution Act gives the works council a right of co-determination over the introduction and use of technical devices designed to monitor the behaviour or performance of employees (our translation). Under section 87(2), if the employer and the works council do not agree, a conciliation board decides. Whether a given measurement design falls under that provision is a question for German counsel. An employer with a works council in Germany should raise the design with the council before it is introduced.

What do California, Colorado, Illinois and New York require?

California

California has two instruments. The California Privacy Protection Agency's regulations cover automated decisionmaking technology, defined in section 7001(e) as technology that processes personal information and uses computation to replace or substantially replace human decisionmaking. The definition says a business substantially replaces human decisionmaking when it uses the output to decide without human involvement, and it sets out what the human reviewer must be able to do. The duties apply when the technology is used for a significant decision, which under section 7001(ddd)(4) includes hiring, allocation or assignment of work, compensation including bonuses, promotion, demotion, suspension and termination.

The agency's rulemaking page shows the regulations were approved on September 22, 2025 and took effect on January 1, 2026. Section 7200(b) gives a business that used the technology for a significant decision before January 1, 2027 until that date to comply. A covered business must give a pre-use notice with a plain language explanation of the specific purpose (section 7220), offer an opt-out unless an exception applies, such as an appeal to a human reviewer with authority to overturn the decision (section 7221), and answer requests for information about how the technology was used (section 7222). Section 7150(b) of the regulations also requires a risk assessment before a business uses automated processing to infer an employee's performance at work from systematic observation, and section 7155(b) gives processing that began before the regulations took effect until December 31, 2027. The regulations use monitoring as their own example: section 7220 describes an employer using productivity monitoring software to decide work allocation, compensation and demotion.

SB 947, approved on September 30, 2026, adds a Labor Code part that becomes operative on July 1, 2027. An employer shall not "rely solely" on an automated decision system when making a disciplinary or termination decision. Where the employer relies primarily on the system's output, a human must corroborate the decision with other information, and the

employee may request a description of their own data that the system used. The civil penalty is $500 per violation.

Colorado

Colorado's SB 26-189 was signed on May 14, 2026. The signed act takes effect on January 1, 2027 and applies to consequential decisions made on or after that date. It covers automated decision-making technology that is used to materially influence a consequential decision, and employment is one of the areas the act covers. The act defines a consumer to include an employee. Before using covered technology in such a decision, a deployer must give a clear and conspicuous notice. If the decision is adverse, the deployer must provide a plain language description of the decision and the technology's role within 30 days. The person can also ask for meaningful human review and for correction of inaccurate personal data. The attorney general enforces the act, which creates no new private right of action.

Illinois

Illinois amended its Human Rights Act through Public Act 103-0804, effective January 1, 2026. It is a civil rights violation for an employer to use AI that has the effect of subjecting employees to discrimination on the basis of a protected class in decisions on recruitment, hiring, promotion, discipline, discharge or the terms and conditions of employment. It is also a violation to fail to give an employee notice that the employer is using AI for those purposes. The statute leaves the timing and form of the notice to rules from the Department of Human Rights, and the compliance calendar linked above tracks where that rulemaking stands.

New York

New York's rule is older and concerns monitoring in general. Civil Rights Law section 52-c applies to private employers with a place of business in the

state. One that monitors employees' telephone, email or internet access or usage must give prior written notice upon hiring, obtain an acknowledgment and post the notice. The section does not apply to processes that manage the type or volume of traffic, are not targeted at a particular individual and are performed solely for system maintenance or protection. The attorney general enforces it, with civil penalties of up to $500 for a first offense, $1,000 for a second and $3,000 for each one after that. The section came from Senate Bill S2628, signed on November 8, 2021 with an effective date 180 days later, which fell in May 2022. Whether an employer's reading of AI usage logs counts as monitoring internet usage under this section is a question for counsel. An employer that already gives this notice should check that its wording covers AI tools.

THE COMMON THREAD

Most of the US duties above attach when an automated output feeds a decision about a person, such as pay, work assignment, promotion, discipline or termination. Two do not wait for a decision: New York's notice rule turns on monitoring itself, and California's risk assessment rule turns on inferring performance from systematic observation. The European rules reach further back, to the handling of personal data and to the introduction of a monitoring system. Keeping the measurement at group level and out of decisions about individuals narrows the exposure under both. It does not remove the duties that come with collecting personal data.

What does this look like in practice?

The steps below pull the earlier sections into an order a company can follow. They are our suggestions and are not a substitute for advice on a specific design.

1. Write down the purpose. State that the measurement exists to decide where AI spend, training and model choice should go by role, and list the uses that are excluded.
2. Bring in HR, privacy counsel and worker representatives before any data is connected. In Germany that means the works council. Under Article 26(7) of the EU AI Act, informing workers and their representatives becomes a legal duty for high-risk systems on December 2, 2027, the date set by Regulation (EU) 2026/1744.
3. Collect the minimum. Token counts for each team, by model and period, answer the question. Leave out prompt text.
4. Match each usage record to a role once, then drop names and emails. Treat hashed identifiers as personal data wherever the GDPR applies.
5. Set the minimum group size and apply it everywhere the data can be seen, including exports.
6. Read output measures at the same group level. When an output measure moves in the same period as token use, treat that as an association to look into, since it does not show that the tokens caused the change.
7. Tell employees in plain words before the first report runs, and give the jurisdiction-specific notices that counsel identifies.
8. Review the design each year and whenever someone proposes a new use. A request to see one person's numbers is a new use.

A company may have good reasons to look at one person's use of a tool, such as a security investigation. Those are separate purposes, and each needs its own assessment of legal basis and notice. They should run through a separate process and stay out of the system that measures spend by role.

How does JobRoute apply this?

JobRoute joins the three feeds described at the start of this paper and reports them by role and team. It reads the usage reports of a company's AI

tools with read access only and groups the usage by role and team. The roster behind the reports has no field for names or emails, and employee IDs are hashed before anything is stored, which means each ID is replaced by a code derived from it. JobRoute does not produce a token score or a ranking for people. The product is described at AI spend by role, and the task scoring behind the second feed is set out in the JobRoute methodology.

Those choices follow the design described in this paper, and they do not settle a customer's legal position. Each company still decides its own purpose and group threshold, gives its own notices, and owes the consultations and assessments that its jurisdictions require.

Method and limits

We checked every legal statement in this paper on November 11, 2026 against the official text linked beside it. The 2017 opinion on data processing at work was issued under the data protection directive that preceded the GDPR, and we cite it for the regulators' reasoning on consent and proportionality.

The paper has four limits a reader should keep in mind. It covers six jurisdictions, and other countries and states have their own rules on employee monitoring. It does not cover collective agreements, which can be stricter than statute. Several of the rules are new, and Colorado's signed act directs the attorney general to adopt rules on its disclosures by January 1, 2027. We did not search for case law or regulator guidance that applies any of these rules to AI usage data. The reading of the EU AI Act's point 4(b) in this paper is ours, and so are the practices on group size and employee notice.

A CAVEAT WORTH NAMING

This paper is general information and not legal advice. Rules on employee data differ by country and state and change often. Before measuring AI use at work, confirm the design with employment and privacy counsel in each place where affected employees work.



ABOUT THE AUTHOR

Mahesh Tatiparthy

Mahesh Tatiparthy is the co-founder and CEO of JobRoute and the author of The Seven Frictions (2026). He founded QueryNow in 2014 and has led its enterprise AI work since.

maheshtatiparthy.com LinkedIn

Sources and further reading

1. Regulation (EU) 2024/1689 (Artificial Intelligence Act): Article 2(11), Article 3(1), Article 6(3), Article 26(2) and 26(7), Article 86(1), Article 99(4) and 99(6), and Annex III point 4(b) EUR-Lex, Official Journal of the European Union, 2024
2. Regulation (EU) 2026/1744 of 8 July 2026 (Digital Omnibus on AI), Official Journal, July 24, 2026: high-risk rules for Annex III systems apply from December 2, 2027; Article 111(2) replaced EUR-Lex, Official Journal of the European Union, 2026
3. Regulation (EU) 2016/679 (General Data Protection Regulation): Recital 26, Article 4(1) and 4(5), Article 5(1)(b) and (c), Article 6(1)(f), Article 13(1)

(c), Article 22(1) and (2), Article 35(1), Article 83(5) and Article 88 EUR-Lex, Official Journal of the European Union, 2016

4. Opinion 2/2017 on data processing at work (WP 249), adopted June 8, 2017 Article 29 Data Protection Working Party, European Commission, 2017

5. Betriebsverfassungsgesetz (Works Constitution Act), section 87: co-determination rights, paragraph 1 number 6 and paragraph 2 (German text) Bundesministerium der Justiz, Gesetze im Internet, 2026

6. Text of Regulations: sections 7001(e), 7001(ddd)(4), 7150(b), 7155(b), 7200(b), 7220, 7221 and 7222 (automated decisionmaking technology) California Privacy Protection Agency, 2025

7. CCPA Updates, Cybersecurity Audits, Risk Assessments, Automated Decisionmaking Technology and Insurance Companies (approved by the Office of Administrative Law September 22, 2025; effective January 1, 2026) California Privacy Protection Agency, 2025

8. Senate Bill No. 947, Chapter 859, Employment: automated decision systems (chaptered text; approved September 30, 2026; operative July 1, 2027) California Legislative Information, 2026

9. SB26-189, Automated Decision-Making Technology (bill page; signed May 14, 2026; Session Laws Chapter 131) Colorado General Assembly, 2026

10. Senate Bill 26-189, signed act: sections 6-1-1701, 6-1-1704, 6-1-1705 and section 5 (effective January 1, 2027) Colorado General Assembly, 2026

11. Public Act 103-0804 (HB 3773), amending section 2-102 of the Illinois Human Rights Act on the use of artificial intelligence in employment (effective January 1, 2026) Illinois General Assembly, 2024

12. Civil Rights Law section 52-c*2, Employers engaged in electronic monitoring; prior notice required New York State Senate, 2022

13. Senate Bill S2628 (2021 to 2022 session), signed November 8, 2021 as Chapter 583; takes effect on the 180th day after becoming law New York State Senate, 2021

14. Configure manager settings: minimum team size and minimum group size in Viva Insights (page dated July 23, 2026) Microsoft Learn, 2026
15. Analytics APIs: Claude Code Analytics API and Claude Enterprise Analytics API (per-user daily metrics) Anthropic, Claude Platform documentation, 2026
16. reportRoot: getMicrosoft365CopilotUsageUserDetail (per-user Copilot activity report; page updated August 26, 2026) Microsoft Learn, Microsoft Graph, 2026
17. An update from Amazon engineering leadership on the role of the developer in the age of AI, October 27, 2026 Amazon, 2026
18. Meta employee token dashboard taken down after report, April 9, 2026 Fortune, 2026
19. Report on Amazon employees inflating token counts, relaying the Financial Times, May 12, 2026 Fortune, 2026
20. Enterprise AI enters an ROI reckoning: State of AI in the Enterprise (3,500 leaders in 25 countries, surveyed May to June 2026), October 29, 2026 Deloitte, 2026

Frequently asked questions

Can a company measure AI token use without tracking individual employees? +

What is a sensible minimum group size for reporting AI use? +

Is a token count a performance measure?

When do the EU AI Act rules on monitoring worker performance apply?

Do California's automated decision rules cover AI usage data?

What should employees be told before AI use is measured?