International Journal of Sustainability and Innovation in Engineering (IJSIE)
https://www.doi.org/10.56830/IJSIE2026
Vol. 4 2026

A Connected Architecture for U.S. Corp-to-Corp Staffing: Representation, Rate Traceability, and Workforce Operations

https://www.doi.org/10.56830/IJSIE202612

Sai Pavan Kumar Gopularam

Giraffy Reach, Minneapolis, Minnesota, United States

Email: saipavankumar.gopularam@gmail.com

Received:16 July, 2026, Accepted:28h October, 2026, Published: 11th November, 2026

Abstract

The architecture described here is targeted at the U.S. market for corp-to-corp (C2C) staffing services. Requirements, representation agreements and negotiation of the rates to be paid are frequently spread out over email threads and messaging groups. Approach. A practitioner-informed conceptual design for such recruitment is developed here from the founder's descriptions of his current workflow, supplemented by a focused narrative discussion of recruitment accountability, explainable AI and human oversight. The design here comprises six stages. Architecture. A candidate signs the right-to-represent (RTR) agreement(s) for that candidate's submission and the employing organization acknowledges representation of the candidate(s) by the submitting organization(s). Where a subsequent rate change is negotiated with the candidate a new signature from that candidate is required. Unreadable screenshots require human review to evaluate agreement. A mismatch between the entered hours on a timesheet and the corresponding evidence uploaded by the employee (e.g. a photo of a time card) to verify those hours will block approval and, consequently, invoice transmission until that mismatch has been corrected by the employee. Status and Limitations. All of the features described here are under current development by the founder and there are currently bugs to be fixed. As far as has been reported, the integrated workflow described here has not yet undergone a structured evaluation and the diagrams provide logical components and controls to be developed and verified as required against a variety of test cases rather than a verified deployment topology. The paper reports no measured efficiency improvements, accuracy improvements, placement improvements or fairness improvements as yet.

Keywords - corp-to-corp staffing; workflow architecture; right to represent; rate traceability; human oversight

1. Introduction

Coordinating C2C staffing for U.S. employers typically involves the candidate, the candidate's employer, recruiters, vendors, and the end client(s). Requirements for open positions are communicated via email and/or messaging groups. Candidates submit themselves for consideration, permission is granted for their representation, and rate negotiation all occur in separate conversations from the rest of the recruitment process. This study proposes a connected architecture that ties together the various communications, the representation agreement(s), the rate(s) of pay, and the various operational records to the appropriate candidate's submission and the corresponding job requirement. As used in this

paper, C2C staffing refers only to corp-to-corp type arrangements and not to all contract employee work.

AI-based recruitment functions are subject to discrimination and accountability issues (Chen, 2003). Such issues can also affect research into explainable AI (Adadi & Berrada, 2018). Distinguishing between generating a recommendation and authorizing an operational action is crucial for a useful staffing architecture. This architecture must always be able to identify who is responsible for a particular action, and which evidence supports it. That evidence can also be required for actions that do not involve AI at all. For example, a timesheet discrepancy requires an explanation for the correction to that discrepancy. However, that explanation for the discrepancy is in no way dependent on the explainability of the platform's matching function.

This contribution outlines a practical, logical architecture that connects representation, rate agreement, placement documentation, time recording and invoicing. It does not claim that equivalent systems are not available on the market. The architecture preserves context across organizational boundaries. In addition, it explicitly outlines where automation is permitted, where consent needs to be reestablished and where an unresolved discrepancy will end the workflow.

This paper presents a practitioner-informed design proposal for a platform to support three functions of a founder of a startup: Giraffy Reach, Giraffy Recruit and C2C Central. The requirements for the platform were derived from the founder's descriptions of these functions and organized as actors, records, actions, approval conditions and exception paths. The literature is used for conceptual background only, not for a systematic review. The platform is under development and thus still contains a number of bugs. All feature descriptions are provided by the founder. No participant study, production-log analysis or controlled experiment has been conducted for this paper.

Section 2 reviews the relevant literature and defines the scope of connected C2C staffing. Section 3 describes the architecture, stakeholder responsibilities, and lifecycle controls. Section 4 maps these controls to explainability and evaluation requirements. Section 5 discusses implementation status, limitations, and future validation. Section 6 concludes the paper.

2. Scope and Related Work

2.1. Digital Labor Intermediation and AI-Augmented Staffing Workflows

The author of this paper describes a model of U.S. corp-to-corp staffing where the candidate is an employee of an employing organization, placed through a vendor (intermediary) that sits in the middle between the employer of the candidate and the end-client. The recruiters from both the vendor and the employing organization negotiate the requirements of the placement and advertise the position for submissions. The author notes that the roles of the organizations involved may change from placement to placement and thus sets out the architecture of the roles described in this paper, as distinct from the use of terms such as 'recruiter', 'vendor', 'employer' and 'client' as synonyms. The architecture described here sets out the continuity of data and approval steps, but does not categorize workers or determine contractual compliance. Such nonstandard work is researched in contrast to temporary agency work and contract-company employment and other forms of nonstandard work (Kalleberg, 2000). While these categories may provide useful background for the C2C

architecture described here, they are not necessarily used as legal definitions of the various forms of nonstandard work.

2.2. Explainable AI Foundations and Recruitment Applications

Explainable AI refers to methods to explain the behavior of black box models (Adadi & Berrada, 2018). In the current architecture, model explanations and workflow explanations have different purposes. On the one hand, a model explanation explains a prediction or a recommendation that the model has produced. On the other hand, a workflow explanation explains which record, rule or approval condition has caused the submission to progress or to stop. An example of the latter are the reported timesheet correction messages. Here a mismatch between the entered hours and the corresponding evidence has been detected. The architecture does not automatically assume faithful explanations of the model's behavior by the components for resume generation, application parsing, and interview-question generation.

2.3. Algorithmic Fairness, Bias, and Accountability Mechanisms

(Chen, 2003) notes there are various forms of potential discrimination when using AI in recruitment that can be addressed by management of data and processes within organizations. However, centralizing candidate data is not in itself sufficient to deal with these potential issues. Fairness evaluation of this future architecture is required, i.e. fairness evaluation of the entire process rather than the individual AI-generated explanation. The founder's feature description and a set of diagrams are not sufficient to establish non-discrimination and/or compliance with laws. (Raghavan, Barocas, Kleinberg, & Levy, 2020) in their paper 'The opaque promises of algorithmic hiring vendors' critically examine the claims made by vendors of algorithmic hiring systems and distinguish between promised safeguarding and evidence of reduction of bias.

2.4. Organizational Implications and Human-AI Interaction

Studies on automation–augmentation investigated scenarios in which human activity was fully replaced by AI and in which human activity was supported by AI (Raisch & Krakowski, 2021). Therefore, in the above workflow an employer could use the comparison for advice or even set up the automated approval of timesheets for specific cases. In order to support employer's decision making by AI, furthermore also the organizational and the implementation-related challenges have been studied (Duan, Edwards, & Dwivedi, 2019). The approval authority as well as the exception handling and the subsequent control of the transfer of the invoice are relevant in this respect. Algorithmic control of the workplace furthermore raises questions concerning the exercise and the contest of authority (Kellogg, Valentine, & Christin, 2020). In this sense, separating information analysis from decision selection and action execution is useful for distinguishing between employer's advice and delegated automation (Parasuraman, Sheridan, & Wickens, 2000).

As opposed to a broad survey of applications of AI, the present paper restricts itself to a complete and operational design. For the purpose of the recruitment process we need to set up controls for accountability, interpretability and the automation vs. human judgment trade-off. As presented in the staffing workflow, the six-stage decomposition, the record relationships and the evaluation cases of the design proposal are to be considered. We use design-science methodology for progressing from problem identification to artifact design and then from design to demonstration and evaluation (Peffers, Tuunanen, Rothenberger, & Chatterjee, 2007). However, this paper remains a design proposal and, consequently, the citation of

design-science methodology does not imply completion of the demonstration and evaluation stages.

3. Architecture and Lifecycle Controls

On one side of the contract, we have candidates and their employing organization (employer), and on the other side, an intermediary vendor or end client. The vendor acts between the employer and the end client. As mentioned previously, recruiters manage requirements and corresponding communication and submissions. This role must not be confused with that of the employer. Negotiation of the submission-linked rate with candidates or their employer may take place prior to agreement of the rate by the employer. The rate is finally decided by the employing organization. Candidates sign the RTR and the employer acknowledges it in respect of candidates that the employer has signed up. Interview and selection decisions are taken by the hiring parties for the relevant stage of the recruitment process. As outlined above, for each business relationship the same contractual signatories for placement documents must be clearly outlined rather than left to be implied by the above-mentioned role names. The architecture comprises six stages: (1) requirement and application; (2) representation and rate agreement; (3) submission, interview and selection; (4) placement documentation; (5) timesheet verification; and (6) invoicing and payment. Table 1 and Figure 1 use this sequence. Communication, correction of errors and handling of disputes occur at various stages of the above mentioned 6 stages of placement. Workflow-pattern research provides a conceptual basis for specifying sequence, branching and synchronization across these stages (van der Aalst, ter Hofstede, Kiepuszewski, & Barros, 2003).

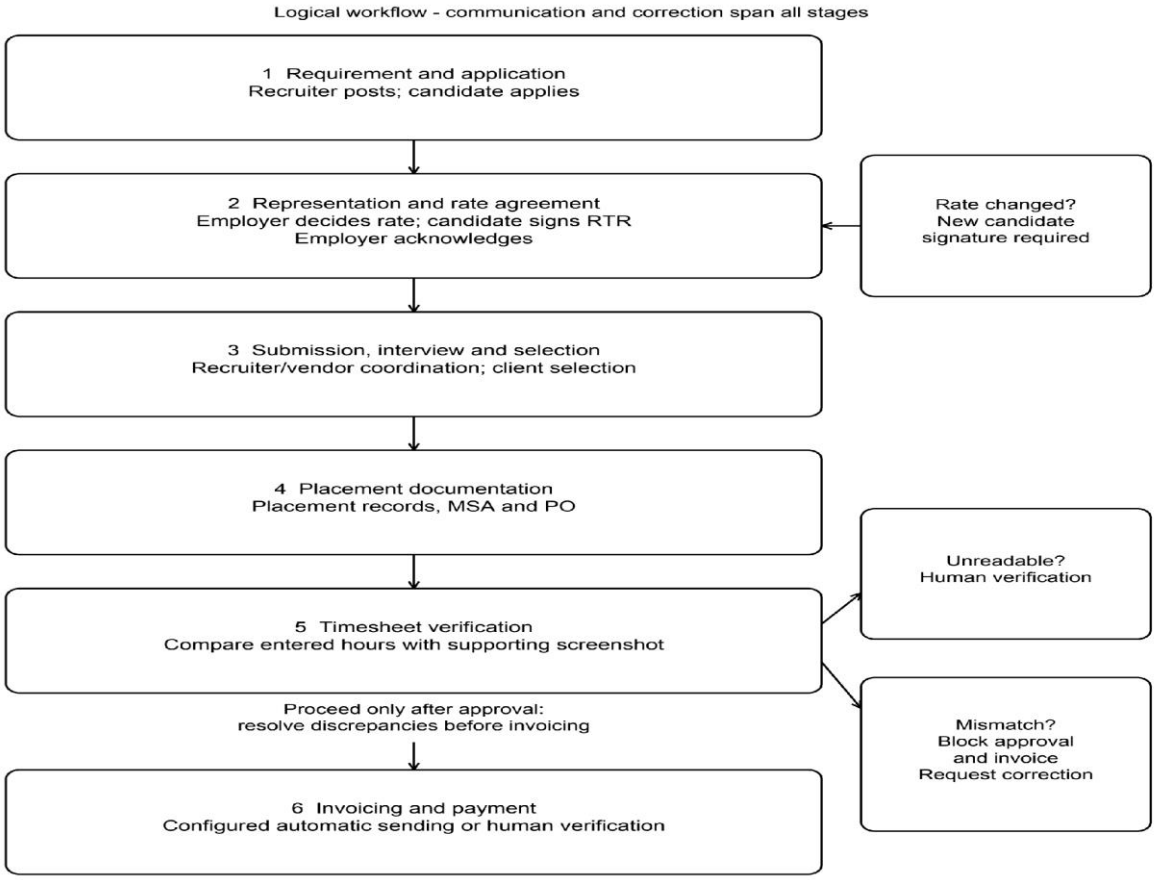


Fig. 1. Proposed six-stage C2C staffing workflow and approval gates. The employer governs rate decisions, while the candidate signs the RTR and signs again after a rate change. Unreadable evidence requires human verification; mismatches block approval and invoice transmission. The figure represents the logical design, not measured performance.

When a recruiter posts a job requirement, a candidate can apply via the platform. In addition to communication about the application via the platform, the rate information is linked to the candidate's application and the job requirement. Candidate's application including representation and commercial agreement are followed through the various stages of the interview and selection process by the recruiter. Once a placement has been approved, the Master Service Agreements (MSAs), Purchase Orders (POs) and other relevant documents can be stored on the platform. The subsequent timesheets and invoices are linked to the workforce administration for the placement.

The negotiation of the rate with the employer of the candidate can be done by the candidate itself or by the employing organization of the candidate. The employer of the candidate decides whether he agrees to the rate proposed by the candidate or not. The agreed rate is then locked to the respective candidate application and job ad. All interested parties (candidate, employer, vendor and end client) can then see the agreed rate for the respective job ad. The candidate signs a right-to-represent (RTR) agreement for the specific job requirement at the agreed rate. This agreement is then acknowledged by the employer. Any changes to the terms and conditions of the rate later agreed by the candidate and the employer with each other has to be signed by the candidate for the revised terms and conditions as well. The prior signed versions of the rate agreement are not verified and have to be validated by the user.

The candidate enters the hours he worked and uploads a supporting screenshot. The employer can choose to enable or disable AI comparison of the candidate's entered hours with the supporting evidence uploaded by the candidate. If the candidate uploaded a supporting screenshot that cannot be read by the AI, human verification of the candidate's uploaded supporting screenshot against the entered hours by the candidate is required. As described above, the approval and transmission of the invoice(s) associated with the candidate's entered hours will be blocked by the system until the mismatch has been resolved by the candidate. The candidate will receive an email with a secure link to the supporting screenshot evidence uploaded by the candidate, and the specific mismatch(s) identified by the AI. The candidate can then update the evidence with correct hours. Invoice transmission can also be configured separately to automatically send the approved invoice(s) or for the transmission to be verified by human(s) prior to sending. These scenarios are currently under development and therefore bugs are being discovered and fixed. No recorded test results exist for the scenario-based validation that has not yet been performed.

3.1. Logical components and record relationships

The information flow of Figure 2, for instance, should not be confused with a verified service or even a complete database schema. Note in particular that a submission consists of Candidate(s) or Candidates and corresponding Job Requirement(s) or Requirements plus the relevant communications plus agreed Rate(s) plus their corresponding RTR(s) or Proof(s) of Rate (s). Also a placement consists of an accepted Submission plus the relevant contractual documents plus corresponding Timesheets plus corresponding Invoices. The Candidate, the

corresponding Requirement, the relevant Submission, the corresponding Placement, the relevant Timesheet and the corresponding Invoice identifiers are proposed as linkage keys for the proof of rate, the representation agreement and the subsequent decisions with respect to the relevant Candidates.

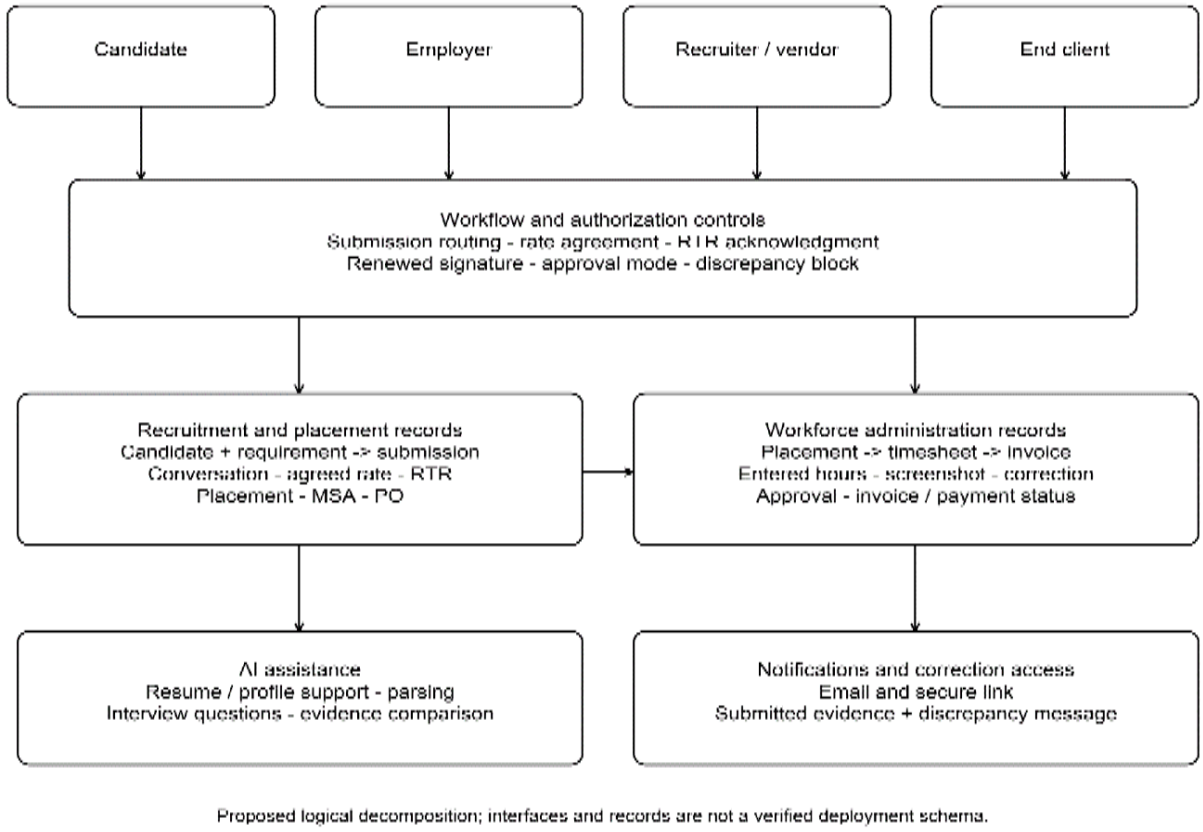


Fig. 2. Proposed logical architecture and information flows. Interfaces communicate with workflow controls and linked recruitment, placement, and administration records. AI assistance and correction notifications support the workflow. Boxes are logical responsibilities, not independently verified services, database tables, or deployed interfaces.

A signed-rate change requires candidate signature after request for new signature for corresponding revised terms which then become signed. Validation requirements for previous agreement versions include retention and retrieval of all versions. Work on Provenance (Buneman, Khanna, & Tan, 2001); (Simmhan, Plale, & Gannon, 2005) distinguishes between the sources contributing to a record and the derivation history of that record (Moreau, et al., 2008). The latter motivates recording the processes that produced the electronic data in question. The above linkage of rates, representation agreements and subsequent decisions with respect to the Candidates is thus informed by concepts for the historical records but does not imply that such historical records are currently retained.

3.2. Interfaces and decision boundaries

For the user interfaces of submission, acknowledgment, correction and administration of respective users, support by AI for resume support, application parsing, interview questions

and timesheet comparison is enabled by feeding task specific input to the AI assistance. In addition to that, workflow controls are implemented to apply the respective approval settings of the employer. Discrepancies for invoices are not transmitted if they remain open. Only AI output is not enabled for actions outside of the approval settings of the employer. Implementation level verification for API contracts, user permissions, signature verification, expiration of online signed secure links, access control and historical retention of online executed agreements is required as well. Role-based access control (RBAC) is used for assigning permissions to organizational roles (Sandhu, Coyne, Feinstein, & Youman, 1996). Collaborative-system research has shown the difficult trade-off between sharing information and protecting information in between participants (Tolone, Ahn, Pai, & Hong, 2005). Workflow authorization constraints are used to distinguish between permission to read a record and permission to perform a certain step in a workflow (Bertino, Ferrari, & Atluri, 1999).

3.3. AI assistance and evidence boundaries

This founder describes functionality such as AI-powered assistance with candidate resumes on LinkedIn, parsing of applications from recruiters, resume-based questions for interviews and evidence compared against timesheets for work completed. The paper does not provide faithful model-level explanations for features described, nor does it evaluate output quality of the features described for independent verification. Instead, operational explanations are provided for specific examples, including a correction to a timesheet and the agreed-upon submission-linked rate that will be visible to all parties in the relevant workflow. That rate, however, is not indicative of all of the commercial agreements and margins in a vendor's chain of labor.

4. Explainability and Evaluation Requirements

This architecture should be viewed as a logical design based on requirements gathered from practitioners combined with a focused narrative use of literature as opposed to a systematic literature search and/or experimental results. Table 1 maps out the various stages of the startup's lifecycle in terms of the various actors, records, decision rights, operational explanations for the various steps in the workflow, proposed fairness checks and corresponding audit fields and points of escalation for each stage. The first column distinguishes between the founder-described workflow rules and the as yet to be verified controls and evaluations. In the case of matching models local attribution methods or interpretable models could be considered if their output and limitations are made explicit (Adadi & Berrada, 2018). The previously described path for the timesheet correction has, however, already a concrete operational explanation for why the entered hours were not correctly assigned: the entered hours and the screenshot do not match. In an evaluation it has to be checked whether the explanation also identifies the real problem, whether it also correctly leads to a correction and whether it is not misused to conclude from matching evidence that the actually worked hours had been entered correctly. The architecture must be able to answer three different questions: 1) do we correctly associate records with submissions and their placements 2) do authorized actions and blocking rules work as expected 3) are decisions made by AI accurate, useful and fair enough. 1 does not automatically imply 3. To prove to have an auditable historical record of events, the proposed event records must contain information such as the responsible actor, affected record, action, time, and reason for the action. Future evaluation must verify completeness and retention of the proposed event records.

To really test our system, we need to try it out with some controlled examples and real workflow records that have been stripped of identifying information. This will help us see how well the system works in key areas. We should test things like negotiating rates, what happens when an agreement is changed after it's been signed, and how the system handles timesheets and images that can't be read. We'll also have two ways to set up invoices. Each test should include what we put into the system, how the system is set up, what we expect to happen, what actually happens, and how it all turns out. We'll also track when the system does something wrong, like letting something happen that shouldn't, or blocking something that should be allowed, and how long it takes to fix mistakes. This testing is still in the planning stages, so we won't have any results to share yet. It would be a good idea to do a separate study to see how well the system's explanations help users find and fix errors. When we do the testing, we should keep track of where our test data comes from, what's in it, and how we plan to use it, like (Gebru, 2021). suggested in2021. We can also use "model cards" to share information about how our AI components are supposed to work, how we tested them, and what their limitations are, as (Mitchell, et al., 2019). But we still need to actually do the testing and report the results, because just documenting things isn't enough.

Table 1. Stage-by-dimension design matrix, presented in three parts. Proposed audit evidence, fairness checks and escalation requirements require validation; they are not reported test results.

Table 1. Stages 1 and 2.

Dimension	1. Requirement and application	2. Representation and rate agreement
Actors	Candidate; recruiter.	Candidate; employing organization; vendor/recruiter.
Records	Requirement, candidate profile, application, communication.	Agreed rate, submission-linked RTR, signature and acknowledgment.
Decision rights	Recruiter posts; candidate applies and corrects information.	Employer decides rate; candidate signs; employer acknowledges. Changed rate requires new signature.
Explanation	Identify the requirement and missing or inconsistent application fields.	Display the current rate, revised terms and required acknowledgment.
Fairness check (proposed)	Compare missing-data and parsing-error patterns across relevant groups where lawful and feasible.	Check whether negotiation access and representation exceptions differ across relevant groups.
Audit evidence (proposed)	Actor, requirement/application IDs, source, timestamp, change reason.	Agreement version, rate, signer, employer acknowledgment, time; prior-version retention to verify.
Escalation	Return incomplete records for correction; route unresolved issues to recruiter review.	Pause treatment of revised terms as signed until a new signature is obtained; employer resolves disputes.

Table 2. Stages 3 and 4.

Dimension	3. Submission interview and selection	4. Placement documentation
Actors	Recruiter/vendor; candidate; responsible hiring parties.	Employer; vendor/client; authorized signatories.
Records	Submission packet, interview material, decision/status.	Placement record, MSA, PO and related documents.
Decision rights	Hiring parties select; generated interview questions assist rather than decide.	Authorized parties approve documents; exact signing authority requires specification.
Explanation	Provide the submission context and recorded human decision reason where available.	Identify missing documents and outstanding approvals.
Fairness check (proposed)	Evaluate selection and review outcomes by relevant groups; assess job-relatedness and data quality.	Check for inconsistent document requirements or completion barriers across comparable cases.
Audit evidence (proposed)	Submission/version IDs, reviewer, decision, time and recorded reason.	Document/version ID, approving actor, signature status, timestamp.
Escalation	Route disputed or incomplete records to responsible human reviewers; do not infer a formal appeal service.	Hold incomplete documentation for authorized review; resolve contractual differences with relevant parties.

Table 3. Stages 5 and 6.

Dimension	5. Timesheet verification	6. Invoicing and payment
Actors	Candidate; employer; designated human reviewer.	Employer; vendor; finance roles.
Records	Entered hours, screenshot, comparison, correction.	Approved timesheet, invoice, payment record.
Decision rights	Employer configures advice or automatic approval. Unreadable evidence goes to humans; mismatch blocks approval.	Configured automatic or human-verified transmission; unresolved discrepancy blocks sending. Settlement controls require validation.
Explanation	Show submitted screenshot and entered-hours mismatch; request corrected evidence through secure link.	State approval status and why transmission is allowed or blocked.
Fairness check (proposed)	Compare extraction and false-block rates across evidence formats and relevant groups where appropriate.	Compare erroneous holds and processing delays across comparable cases; investigate disparities.
Audit evidence (proposed)	Evidence reference, entered hours, comparison outcome, mode, reviewer, correction history.	Invoice/timesheet IDs, approval mode, transmission and payment status, actor, timestamp.
Escalation	Candidate corrects and resubmits; human reviews unreadable or disputed evidence.	Employer or finance reviewer resolves holds and discrepancies; payment-state reconciliation remains to be verified.

5. Implementation Status, Limitations, and Future Validation

From a data provenance and representation standpoint, the connected staffing ecosystem consists of resumes, skill claims, job requirements, representation agreements, work data and billing events, each with varying degrees of accuracy, ownership, frequency of update and sensitivity. The system should never convert a resume into a fixed predictive profile without recording the source, consent/legitimate purpose for processing, full audit trail of processing steps, system-specific retention rules and complete process for correction of errors. The system must implement data minimization and provide users with views relevant to their roles. For example, an employer's view of the candidate must not contain information that is not required for a placement decision, and the finance user's view of recruitment information must not contain information that is not required to ensure that billing of correctly paid invoices is accurate.

Candidates receive email notifications and click on a secure link to access the correction workflow. In the example of a timesheet discrepancy the candidate is then shown the uploaded screenshot and informed that the entered hours for that timesheet do not match the evidence uploaded. The candidate is then asked to re-upload corrected evidence for those hours. This establishes a reported correction mechanism as opposed to a fully specified appeal process for each platform decision. A formal appeal would require a reviewer, a recorded disposition and a documented resolution process.

The notion of fairness cannot be reduced to the removal of protected attributes and to the display of feature importances. Historical labor market data contains structural differences which may have been encoded by proxies, missing values, different job labels, and different referral practices. Therefore, it is necessary to specify the protected/affected groups where applicable, the definition of fairness, the decision threshold, the monitoring frequency, and the action that is taken in case of disparities. Evaluation of fairness has to consider the complete workflow because neutral ranking still can lead to unfair outcomes if the recruiter's review, the representation process, or any downstream compliance rule differentially filters members of one group.

The goal of explaining how something works should be to help the people involved. For example, a recruiter needs a quick way to understand why a job applicant was chosen or not. The applicant wants to know what they did wrong. Technical people might need to see how the system is working. This doesn't have to be fancy or look nice, it just needs to be accurate, stable, and useful. Just because a system can tell us what it's doing, doesn't mean we really understand how it's making decisions or if it's being fair. There are guidelines for how humans and AI systems can work together that can help us design systems that give good feedback and allow for corrections when the AI is wrong or unsure. We need to use these guidelines to design and test our system to make sure it's working well.

During the different stages of implementation, candidate profiles and job matching, the founder reports also about the job posting of recruiter/vendors and RTR signing as well as about the timesheets, the invoicing, the payments and the decision history. Candidates can use the typical job-search functions of a job portal. The RTR and recruiter/vendor workflows are not yet in active use by the founder. Development and bug fixing are ongoing. A structured operational evaluation of the integrated lifecycle has not yet taken place. Instead, a video walkthrough of the typical user behavior has been offered but is not part of the evidence provided here. A demonstration of the system would only show the potential user interactions and does not provide any evidence of improved processing time, accuracy, placement rates or fairness.

6. Conclusion & Future Work

This paper presents a connected architecture to support U.S. corp-to-corp staffing. The architecture supports the six stages of the staff lifecycle and defines corresponding roles within an organization. The paper also describes how linked records (such as those for candidate, job requirements, RTR agreement and rate agreed for that candidate) can be used. Candidate signature of RTR agreement and employer signature of acknowledgment of candidate are supported, as well as subsequent change to signed rate requiring candidate signature of new terms. Employer-configured timesheet processing rules are described to produce unreadable evidence that must be verified by human evaluators. ENTERED HOURS NOT EQUAL TO UPLOADED EVIDENCE rule is described that will prevent approval and

subsequent invoice processing until entered hours equals corresponding evidence uploaded by candidate. As with any staff system described here, verification of normal and exception paths, record retention, access control and quality of explanation for all activities are required before benefits of operational use can be described. Work in development and bug-fixing of connected architecture described in this paper is ongoing and the paper does not claim effective or fair implementation of the staff system described.

References:

Adadi, A., & Berrada, M. (2018). Peeking inside the black-box: A survey on explainable artificial intelligence (XAI). *IEEE Access*, 6, 52138–52160. https://doi.org/10.1109/ACCESS.2018.2870052.

Bertino, E., Ferrari, E., & Atluri, V. (1999). The specification and enforcement of authorization constraints in workflow management systems. *ACM Transactions on Information and System Security*, 2(1), 65–104. https://doi.org/10.1145/300830.300837.

Buneman, P., Khanna, S., & Tan, W.-C. (2001). Why and where: A characterization of data provenance. *In J. Van den Bussche & V. Vianu (Eds.), Database theory—ICDT 2001*, (pp. 316–330). Springer. https://doi.org/10.1007/3-540-44503-x_20.

Chen, Z. (2003). Ethics and discrimination in artificial intelligence-enabled recruitment practices. *Humanities and Social Sciences Communications*, 10(1), Article 567. https://doi.org/10.1057/s41599-023-02079-x.

Duan, Y., Edwards, J. S., & Dwivedi, Y. K. (2019). Artificial intelligence for decision making in the era of Big Data – evolution, challenges and research agenda. *International Journal of Information Management*, 48, 63–71. https://doi.org/10.1016/j.ijinfomgt.2019.01.021.

Gebru, T. M. (2021). Datasheets for datasets. *Communications of the ACM*, 64(12), 86–92. https://doi.org/10.1145/3458723.

Kalleberg, A. L. (2000). Nonstandard employment relations: Part-time, temporary and contract work. *Annual Review of Sociology*, 26, 341–365. https://doi.org/10.1146/annurev.soc.26.1.341.

Kellogg, K. C., Valentine, M. A., & Christin, A. (2020). Algorithms at work: The new contested terrain of control. *Academy of Management Annals*, 14(1), 366–410. https://doi.org/10.5465/annals.2018.0174.

Mitchell, M., Wu, S., Zaldivar, A., Barnes, P., Vasserman, L., Hutchinson, B., et al. (2019). Model cards for model reporting. *In Proceedings of the Conference on Fairness, Accountability, and Transparency*, (pp. 220–229) Association for Computing Machinery. https://doi.org/10.1145/3287560.3287596.

Moreau, L., Groth, P., Miles, S., Vazquez-Salceda, J., Ibbotson, J., Jiang, S., et al. (2008). The provenance of electronic data. *Communications of the ACM*, 51(4), 52–58. https://doi.org/10.1145/1330311.1330323.

Parasuraman, R., Sheridan, T. B., & Wickens, C. D. (2000). A model for types and levels of human interaction with automation. *IEEE Transactions on Systems, Man, and Cybernetics—Part A: Systems and Humans*, 30(3), 286–297. https://doi.org/10.1109/3468.844354.

Peffers, K., Tuunanen, T., Rothenberger, M. A., & Chatterjee, S. (2007). A design science research methodology for information systems research. *Journal of Management Information Systems*, 24(3), 45–77. https://doi.org/10.2753/MIS0742-1222240302.

Raghavan, M., Barocas, S., Kleinberg, J., & Levy, K. (2020). Mitigating bias in algorithmic hiring: Evaluating claims and practices. *In Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, (pp. 469–481). Association for Computing Machinery. https://doi.org/10.1145/3351095.3372828.

Raisch, S., & Krakowski, S. (2021). Artificial intelligence and management: The automation–augmentation paradox. *Academy of Management Review*, 46(1), 192–210. https://doi.org/10.5465/amr.2018.0072.

Sandhu, R. S., Coyne, E. J., Feinstein, H. L., & Youman, C. E. (1996). Role-based access control models. *Computer*, 29(2), 38–47. https://doi.org/10.1109/2.485845.

Simmhan, Y. L., Plale, B., & Gannon, D. (2005). A survey of data provenance in e-science. *ACM SIGMOD Record*, 34(3), 31–36. https://doi.org/10.1145/1084805.1084812.

Tolone, W., Ahn, G.-J., Pai, T., & Hong, S.-P. (2005). Access control in collaborative systems. *ACM Computing Surveys*, 37(1), 29–41. https://doi.org/10.1145/1057977.1057979.

van der Aalst, W. M., ter Hofstede, A. H., Kiepuszewski, B., & Barros, A. P. (2003). Workflow patterns. *Distributed and Parallel Databases*, 14(1), 5–51. https://doi.org/10.1023/A:1022883727209.