

---

# Callable by Whom? Actor–Route Alignment in Public-Sector Agent Evaluation

---

Anonymous Author(s)

Affiliation

Address

email

## Abstract

1       An agent evaluation attributes institutional capability whenever its result is taken  
2       to show what a particular public body or citizen can accomplish. Such an attri-  
3       bution depends on how the actor’s authority relates to the tools supplied by the  
4       evaluator. I propose actor–route alignment as a six-field reporting requirement  
5       and apply it to 43 brokered and nine external Brazilian API pages. The resulting  
6       records identify conditions on different actors, operations and approval layers that  
7       a page-level permission label would collapse; all 52 situated attributions remain  
8       unresolved in the collected evidence. The municipal case makes the consequence  
9       concrete: only 52 of 5,569 administrations meet Conecta’s currently published  
10      capital-or-population condition, although they contain 30.4% of Brazil’s popula-  
11      tion. A bounded UK comparison finds 27 of 35 sampled catalogue descriptions  
12      silent about access conditions, while an Indian specification example conflicts  
13      with its accompanying prose. These findings support treating the authority sup-  
14      plied by an evaluation as part of the capability claim, with explicit counterfactual  
15      grants where institutional correspondence has not been established.

## 16   1 Introduction

17   Brazil’s Conecta gov.br catalogue invites municipal public bodies to seek access to 14 APIs, yet the  
18   programme’s subnational guidance limits that invitation to state capitals and municipalities meeting  
19   a population threshold around 500,000 residents, alongside technical and data-protection require-  
20   ments [Governo Digital, Brasil, n.d.b]. When a benchmark supplies one of these APIs to a generic  
21   “municipal actor”, it can remove an institutional distinction before the agent makes its first decision.  
22   The resulting score remains informative about performance under the supplied configuration, while  
23   its attribution to a particular municipality depends on whether that municipality could act through  
24   the represented route.

25   I call this correspondence *actor–route alignment*: the warrant connecting a represented actor’s au-  
26   thority to the institutional route exercised in an evaluation. Establishing it requires identifying the  
27   principal on whose behalf the agent acts, the delegated operation and purpose, and the conditions  
28   governing that route. Together, these relations define an actor’s *institutional action surface*, the  
29   operations available through specified institutional routes. A callable function describes only part  
30   of this surface, since neither technical availability nor membership in an eligible class establishes  
31   completed accession or delegation.

32   This reporting proposal draws on a Brazilian documentary study and bounded UK and Indian com-  
33   parisons; it does not establish misattribution in the cited benchmarks. Its claim concerns the inter-  
34   pretation of evaluations: when a result is attributed to an institutional actor, the authority supplied  
35   in the harness must be connected to that actor by evidence or declared as a counterfactual grant.

36 The political decision already enters through the design of the task, where the evaluator determines  
37 which institutional differences survive abstraction into a common interface.

## 38 2 Position and prior work

39 Permissions are already part of agent evaluation. WebArena represents roles and permission-limited  
40 tasks, AppWorld models delegated accounts and ownership, and  $\tau$ -bench supplies authentication  
41 and confirmation policies [Zhou et al., 2024, Trivedi et al., 2024, Yao et al., 2025]; EnterpriseBench,  
42 HarnessAudit-Bench and Harness-Bench also represent access or security constraints [Vishwakarma  
43 et al., 2025, Liu et al., 2026, Yao et al., 2026]. Appendix A locates these mechanisms; their insti-  
44 tutional interpretation depends on evidence connecting the modelled permissions to the authority of  
45 the represented actor.

46 Harness disclosure makes configuration part of a score’s interpretation [Zhang et al., 2026], while  
47 public-sector benchmark requirements include task-centred construction, realism and appropriate  
48 metrics [Rystrøm et al., 2026]. Actor–route alignment addresses the further evidential step through  
49 which the authority modelled inside such a configuration is attributed to a particular institution  
50 or person. This also bears on Hagar and Diakopoulos’s distinction between situated delegation  
51 instructions, institutional authorisation and observed compliance [Hagar and Diakopoulos, 2026].

52 The closest conceptual precedent is Armesto and Kolb’s delegation envelope, which specifies when  
53 action is authorised through semantic, evidentiary, procedural and institutional predicates [Armesto  
54 and Kolb, 2026]. Here the object of assessment is the warrant for an evaluation’s attribution: what  
55 evidence connects its proposed or exercised route to the authority of its represented actor? An enve-  
56 lope can contribute to that assessment, but its presence in a benchmark does not itself document the  
57 correspondence. OAuth 2.0 and OpenID Connect delegation, DCAT access-rights metadata, ODRL  
58 policies and OpenAPI security requirements can likewise supply relevant representations [South  
59 et al., 2025, W3C, 2024, OpenAPI Initiative, 2017]; the reporting record assesses their applicability  
60 to a claim and retains unresolved cases.

## 61 3 Documentary study

62 The cases were selected for their different documentary forms: a Brazilian API archive, a structured  
63 UK catalogue and Indian platform guidance. They support comparison of representational problems  
64 rather than country rankings. Coding and author verification, including checks of initially silent  
65 descriptions, are documented in Appendix D; no catalogued API was invoked.

66 **Brazil: integration and accession.** The archive retains raw HTML, text and hashes for 95 detail  
67 pages discovered on 29 August 2026 [Governo Digital, Brasil, 2026b]. All 43 labelled *Disponível*  
68 *via Conecta* direct public bodies to accession: 21 name federal bodies, eight federal and state bodies,  
69 and 14 all three levels. Thirty-nine publish production URLs, including three authentication URLs,  
70 and 33 link Swagger/OpenAPI material. These features coexist with accession instructions, so a  
71 technical description need not displace an institutional one. The index response was not retained,  
72 which prevents reconstructing catalogue-wide completeness.

73 Among the 52 *API externa* pages, nine contain explicit conditions: three institutional-accession  
74 statements, one service-registration requirement, two operation-authentication descriptions, two  
75 credential-input descriptions and one citizen signing-account condition. The remaining 43 contain  
76 no explicit condition under the page-local definitions; linked supplier documentation falls outside  
77 this count. The primary codes distinguish pages for analysis, although several conditions may oper-  
78 ate together on one route (Appendix E).

79 **Municipal scope and participation.** All 14 municipal invitations persist in the 5 September cap-  
80 tures and match the subnational programme’s API list. Its minimum conditions require a state capital  
81 or population  $\geq 500,000$ , while the FAQ says more than 500,000 and qualifies the restriction as *no*  
82 *momento* (currently) [Governo Digital, Brasil, n.d.b]. With this temporal scope, either wording ap-  
83 plied to IBGE’s 1 July 2026 estimates yields 52 of 5,569 municipalities: 26 capitals and 26 others,  
84 together containing 65.1 million residents, or 30.4% of the national population [Instituto Brasileiro  
85 de Geografia e Estatística, 2026]. Thus 99.1% of municipal administrations fail the size-or-capital

condition, a finding about institutional actors that cannot be read as the share of residents denied a service. Four non-capital municipalities lie within 10% below the threshold, the largest being Caxias do Sul at 479,604 residents.

The specific route qualifies the generic FAQ, which still excludes municipalities, and the accession procedure separates organisational registration, an API-specific request and provider approval [Governo Digital, Brasil, 2026c, n.d.a]. The programme also publishes listings for body–dataset approvals and body–API consumption [Ministério da Gestão e da Inovação em Serviços Públicos, 2026b,a]. However, the portal marked the accession dataset as outdated and some resource URLs as unavailable at capture; retrieval attempts did not yield rows. The public record therefore identifies where part of the required connection should be documented, while the collected evidence remains insufficient to assess individual approvals.

**UK and India: access descriptions.** The pinned UK catalogue contains 233 records with no dedicated accession field [UK Government, n.d.]. A reproducible sample of one description per provider yields three requirements, four explicit waivers, 27 descriptions silent about the coded conditions and one with unclear scope. Linked provider references can be more explicit: the Environment Agency waives registration, whereas Companies House documents authentication credentials [Environment Agency, n.d., Companies House, n.d.]. Catalogue structure and the articulation of access conditions are therefore separate properties of these records.

India’s API Setu guidance separates platform registration from publisher approval [National e-Governance Division, 2024, n.d.]. In the 31-page manual, the prose requires both client ID and API key, but the OpenAPI 3.0.1 example places them in separate security requirement objects (pp. 17, 19, 25), which express alternatives under the standard [OpenAPI Initiative, 2017, §4.7.30]. The representations disagree in two of four possible credential-presence assignments. Both linked SOPs were read; this selected inconsistency concerns the documentation, with no claim about live enforcement or its prevalence elsewhere.

## 4 Reporting actor–route alignment

For each consequential route, the evaluator records six items: actor and delegation; operation, data and purpose; institution and route; dated evidence of applicable conditions; harness mechanism and credentials; and attribution status. The first three specify the institutional action; the next two link applicable evidence to the supplied grant, and the last states what that warrants. The grouping reflects the cases’ distinctions, without claiming exhaustive coverage of institutional authority. A claim is *supported* when the evidence warrants the correspondence, *contradicted* when its route conflicts with applicable conditions, and *unresolved* when the connection depends on missing or conflicting evidence. Status concerns the specified claim, so documented access conditions can be supported while executed attribution remains untested (Table 1).

Table 1: Assessments of proposed evaluation designs. Appendix C supplies the six fields; no deployment was run.

| Proposed design                                                              | Applicable evidence                                                        | Assessment                                     |
|------------------------------------------------------------------------------|----------------------------------------------------------------------------|------------------------------------------------|
| Developer reads EA flood alerts by anonymous GET, with attribution           | Provider waives registration; route and data terms specified               | Access condition supported; execution untested |
| Citizen receives a direct Conecta certificate-API tool                       | Programme excludes direct citizen use; separate public consultation exists | Grant through this route contradicted          |
| Eligible public body receives that tool; approval and delegation unspecified | Eligibility and accession procedure documented                             | Situated attribution unresolved                |
| Caxias do Sul receives a municipal-list API tool                             | Non-capital; 479,604 residents                                             | Size eligibility contradicted on this route    |

121 In a proposed Porto Alegre address-validation task, the actor is the municipal administration, with  
122 delegation assumed; the operation is a CEP lookup on synthetic address data for service registration;  
123 the route is Conecta’s CEP API; the evidence is the archived catalogue invitation and guidance  
124 captured on 5 September; the harness is a simulator with a supplied test credential; and attribution  
125 is unresolved because the body’s accession and delegation have not been evidenced. Performance  
126 would be reported under the explicitly assumed grant (Appendix C). Operation scope requires the  
127 same care: the certificate API returns *Nada consta* or a direction to attend the Federal Police, so  
128 granting authority cannot turn it into complete criminal-history retrieval [Governo Digital, Brasil,  
129 2026a].

## 130 **5 The authority supplied by a benchmark**

131 The 52 unresolved Brazilian records identify distinct evidential problems: a citizen’s verified signing  
132 account cannot establish an API integrator’s authority; credential-maintenance inputs concern the  
133 right to modify an account, while DOU publishing combines interface approval, a firewall rule,  
134 publishing-user permissions and SOAP credentials. These conditions concern different actors and  
135 operations even on the same page, distinctions that a single permission label would erase.

136 These records identify claims requiring organisational approval, a particular credential, an operation-  
137 specific permission or evidence of delegation; their page-derived actor classes provide no sample of  
138 completed deployments from which to estimate unauthorised use. A body–dataset approval row  
139 could support organisational accession while leaving the agent’s authority for the proposed opera-  
140 tion unresolved. The reusable schema preserves these dependencies and permits all three statuses;  
141 validation checks completeness and references, leaving the evidence for substantive assessment.

142 An evaluation can proceed when a correspondence is unresolved, provided its capability claim re-  
143 mains conditional on the configuration supplied. The evaluator may publish the relevant warrant,  
144 make restricted evidence reviewable, or specify which authority is assumed counterfactually; in a  
145 task with several routes, unresolved dependencies must remain identifiable alongside performance.  
146 In the Caxias do Sul design, even a successful lookup could support capability under broadened  
147 access but would not establish that municipality’s present capability through the route.

## 148 **6 Discussion and conclusion**

149 The municipal case locates the politics of evaluation in the assignment of tools to actors. Munici-  
150 palities whose administrations fall outside the published condition contain 146.1 million residents,  
151 or 68.2% of Brazil’s population; a generic municipal grant would therefore extend the represented  
152 route to administrations serving roughly seven in ten residents. These counts concern the adminis-  
153 tration through which a service is provided, preserving the distinction between population coverage  
154 and administrative authority. In the terms of Mohamed, Png and Isaac’s decolonial analysis, the  
155 design decision concerns whose institutional knowledge enters the technical representation [Mo-  
156 hamed et al., 2020]. Preserving the restriction makes it available for scrutiny without endorsing its  
157 legitimacy or treating present guidance as permanent.

158 The catalogue comparison also resists a simple hierarchy of technical maturity: explicit Portuguese  
159 accession instructions appear on all 43 Conecta pages, whereas 27 of 35 sampled UK descriptions  
160 are silent about the coded conditions despite their machine-readable format. The different units  
161 preclude a national comparison, but they show why machine readability cannot substitute for in-  
162 stitutional interpretation. *Legibility asymmetry* remains a hypothesis about the additional work,  
163 and possible omissions, involved when conditions fall outside an evaluator’s preferred language  
164 or schema; neither reading costs nor omission rates were measured here.

165 These cases show that institutional authority is distributed across actors, operations and approval lay-  
166 ers whose relationships must survive the construction of an evaluation task. Actor–route alignment  
167 makes those relationships inspectable and identifies the evidence needed to attribute performance to  
168 a public body or citizen. An evaluation may deliberately imagine a different distribution of authority,  
169 including one that broadens access, but its conclusion must then concern capability under that pro-  
170 posed arrangement. Making the grant explicit allows the result to inform both technical assessment  
171 and discussion of who should be able to act.

## References

- Maximiliano Armesto and Christophe Kolb. Toward a science of intent: Closure gaps and delegation envelopes for open-world AI agents. *arXiv preprint arXiv:2604.25000v2*, 2026. URL <https://arxiv.org/abs/2604.25000v2>.
- Companies House. API authentication, n.d. URL <https://developer.company-information.service.gov.uk/authentication>. Accessed 5 September 2026.
- Environment Agency. Real time flood-monitoring API: reference, n.d. URL <https://environment.data.gov.uk/flood-monitoring/doc/reference>. Access and reuse guidance; accessed 5 September 2026.
- Governo Digital, Brasil. Certidão de antecedentes criminais — catálogo de APIs governamentais, 2026a. URL <https://www.gov.br/conecta/catalogo/apis/antecedentes-criminais>. Accessed 29 August 2026.
- Governo Digital, Brasil. Catálogo de APIs governamentais, 2026b. URL <https://www.gov.br/conecta/catalogo/apis>. Accessed 29 August 2026.
- Governo Digital, Brasil. Conecta gov.br, 2026c. URL <https://www.gov.br/governodigital/pt-br/infraestrutura-nacional-de-dados/interoperabilidade/conecta-gov.br>. Accessed 5 September 2026.
- Governo Digital, Brasil. Quero acessar as APIs do conecta, n.d.a. URL <https://www.gov.br/governodigital/pt-br/infraestrutura-nacional-de-dados/interoperabilidade/conecta-gov.br/quero-acessar-as-apis-do-conecta>. Accession procedure; accessed 5 September 2026.
- Governo Digital, Brasil. Conecta para estados e municípios, n.d.b. URL <https://www.gov.br/governodigital/pt-br/infraestrutura-nacional-de-dados/interoperabilidade/conecta-gov.br/conecta-para-estados-e-municipios>. Specific subnational guidance; accessed 5 September 2026.
- Nick Hagar and Nicholas Diakopoulos. Runtime configuration for situated governance of AI agents: a case study in investigative journalism. *AI and Ethics*, 6, 2026. doi: 10.1007/s43681-026-01298-7. URL <https://doi.org/10.1007/s43681-026-01298-7>. Article 456.
- Instituto Brasileiro de Geografia e Estatística. Estimativas da população residente nos municípios brasileiros com data de referência em 1 de julho de 2026, 2026. URL [https://ftp.ibge.gov.br/Estimativas\\_de\\_Populacao/Estimativas\\_2026/estimativa\\_dou\\_2026.xlsx](https://ftp.ibge.gov.br/Estimativas_de_Populacao/Estimativas_2026/estimativa_dou_2026.xlsx). Municipal worksheet; retrieved 5 September 2026.
- Chengzhi Liu, Yichen Guo, Yepeng Liu, Yuzhe Yang, Qianqi Yan, Xuandong Zhao, Wenyue Hua, Sheng Liu, Sharon Li, Yuheng Bu, and Xin Eric Wang. Auditing agent harness safety. *arXiv preprint arXiv:2605.14271v2*, 2026. URL <https://arxiv.org/abs/2605.14271v2>.
- Ministério da Gestão e da Inovação em Serviços Públicos. Consumo conecta gov.br, 2026a. URL <https://dados.gov.br/dados/conjuntos-dados/consumo-conecta-govbr>. Resource listings retrieved 5 September 2026; resource file contents not retrieved.
- Ministério da Gestão e da Inovação em Serviços Públicos. Adesão ao conecta gov.br: programa de interoperabilidade do poder executivo federal, 2026b. URL <https://dados.gov.br/dados/conjuntos-dados/adesao-ao-conecta-govbr>. Resource listings retrieved 5 September 2026; resource file contents not retrieved.

220 Shakir Mohamed, Marie-Therese Png, and William Isaac. Decolonial AI: Decolonial theory as  
 221 sociotechnical foresight in artificial intelligence. *Philosophy & Technology*, 33:659–684, 2020.  
 222 doi: 10.1007/s13347-020-00405-8. URL  
 223 <https://doi.org/10.1007/s13347-020-00405-8>.

224 National e-Governance Division. API Setu: Standard operating procedures, 2024. URL  
 225 <https://cdn.apisetu.gov.in/portal/assets/sop-apisetu-v0.pdf>.  
 226 Version 02/2024; five-page PDF linked as V1 by the SOP landing page; accessed 5 September  
 227 2026.

228 National e-Governance Division. API Setu: User manual, n.d. URL  
 229 <https://cdn.apisetu.gov.in/portal/assets/sop-apisetu-v1.pdf>.  
 230 31-page PDF linked as V2 by the SOP landing page; accessed 5 September 2026.

231 OpenAPI Initiative. OpenAPI specification v3.0.1, 2017. URL  
 232 <https://spec.openapis.org/oas/v3.0.1.html>. Sections 4.7.30 and 4.8; 6  
 233 December 2017; accessed 5 September 2026.

234 Jonathan Rystrom, Chris Schmitz, Karolina Korgul, Jan Batzner, and Chris Russell. Agent  
 235 benchmarks fail public sector requirements. *arXiv preprint arXiv:2601.20617v1*, 2026. URL  
 236 <https://arxiv.org/abs/2601.20617v1>.

237 Tobin South, Samuele Marro, Thomas Hardjono, Robert Mahari, Cedric Deslandes Whitney, Dazza  
 238 Greenwood, Alan Chan, and Alex Pentland. Authenticated delegation and authorized AI agents.  
 239 *arXiv preprint arXiv:2501.09674v1*, 2025. URL  
 240 <https://arxiv.org/abs/2501.09674v1>.

241 Harsh Trivedi, Tushar Khot, Mareike Hartmann, Ruskin Manku, Vinty Dong, Edward Li, Shashank  
 242 Gupta, Ashish Sabharwal, and Niranjan Balasubramanian. AppWorld: A controllable world of  
 243 apps and people for benchmarking interactive coding agents. In *Proceedings of the 62nd Annual  
 244 Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages  
 245 16022–16076. Association for Computational Linguistics, 2024. doi:  
 246 10.18653/v1/2024.acl-long.850. URL  
 247 <https://aclanthology.org/2024.acl-long.850/>.

248 UK Government. API catalogue: data and metadata schema, n.d. URL  
 249 [https://github.com/co-cddo/api-catalogue/tree/  
 250 5ac787bd84b04ac9b3a163e704fc4470c01febb1/data](https://github.com/co-cddo/api-catalogue/tree/5ac787bd84b04ac9b3a163e704fc4470c01febb1/data). Revision  
 251 5ac787bd84b04ac9b3a163e704fc4470c01febb1; accessed 5 September 2026.

252 Harsh Vishwakarma, Ankush Agarwal, Ojas Patil, Chaitanya Devaguptapu, and Mahesh Chandran.  
 253 Can LLMs help you at work? a sandbox for evaluating LLM agents in enterprise environments.  
 254 In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*,  
 255 pages 9167–9201. Association for Computational Linguistics, 2025. doi:  
 256 10.18653/v1/2025.emnlp-main.466. URL  
 257 <https://aclanthology.org/2025.emnlp-main.466/>.

258 W3C. Data catalog vocabulary (DCAT)—version 3, 2024. URL  
 259 <https://www.w3.org/TR/vocab-dcat-3/>. W3C Recommendation; sections 6.4.1  
 260 and 9; accessed 5 September 2026.

261 Shunyu Yao, Noah Shinn, Pedram Razavi, and Karthik Narasimhan.  $\tau$ -bench: A benchmark for  
 262 tool-agent-user interaction in real-world domains. In *The Thirteenth International Conference  
 263 on Learning Representations*, 2025. URL  
 264 <https://openreview.net/forum?id=roNSXZpUDN>. Mechanism locators refer to the  
 265 reviewed arXiv:2406.12045v1 version.

266 Yilun Yao, Xinyu Tan, Chao-Hsuan Liu, Yaoming Li, Zhengyang Wang, Wenhan Yu, Zhewen Tan,  
 267 Yuxuan Tian, Guangxiang Zhao, Lin Sun, Xiangzheng Zhang, and Tong Yang. Harness-Bench:  
 268 Measuring harness effects across models in realistic agent workflows. *arXiv preprint  
 269 arXiv:2605.27922v1*, 2026. URL <https://arxiv.org/abs/2605.27922v1>.

- 270 Yunbei Zhang, Janet Wang, Yingqiang Ge, Weijie Xu, Jihun Hamm, and Chandan K. Reddy. Stop  
 271 comparing LLM agents without disclosing the harness. *arXiv preprint arXiv:2605.23950v1*,  
 272 2026. URL <https://arxiv.org/abs/2605.23950v1>.
- 273 Shuyan Zhou, Frank F. Xu, Hao Zhu, Xuhui Zhou, Robert Lo, Abishek Sridhar, Xianyi Cheng,  
 274 Tianyue Ou, Yonatan Bisk, Daniel Fried, Uri Alon, and Graham Neubig. WebArena: A realistic  
 275 web environment for building autonomous agents. In *International Conference on Learning*  
 276 *Representations*, 2024. URL <https://arxiv.org/abs/2307.13854v4>. Reviewed  
 277 version: arXiv:2307.13854v4.

## 278 A Source-specific comparison

279 The comparison concerns explicit mechanisms in six permission-relevant evaluations. It does not  
 280 code unreported properties as absent or estimate their frequency across agent benchmarks. The  
 281 bibliography identifies the versions used; the supplement records PDF hashes.

Table 2: Mechanisms supporting the comparisons in Section 2. Internally modelled permissions need additional institution-specific evidence when used to support an actual actor’s authority.

| Evaluation         | Mechanism in the full paper                                                                                                   | Locator                          |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------|----------------------------------|
| WebArena           | Explicit user roles and access rights; tasks unachievable because of permissions.                                             | §2.4; §3.2; Appendix A.3         |
| AppWorld           | Delegated account access; authentication, ownership and privacy checks.                                                       | §2.2; Appendices J–L             |
| $\tau$ -bench      | Domain policies require authentication and user confirmation. Terminal-state success does not certify every procedural step.  | §3; Appendix B, policies         |
| EnterpriseBench    | Generated and expert-validated access hierarchies; permission-sensitive unanswerable tasks and an access-constraint ablation. | §3.2.2; Appendices A.2 and A.4.2 |
| HarnessAudit-Bench | Synthetic policies and trajectory-level auditing of permissions, intent and resource access.                                  | §3; Appendices A–D               |
| Harness-Bench      | Security gate and an explicit workplace-policy example covering unauthorised access and secret exposure.                      | §3.4; Appendix D                 |

## 282 B Conecta extraction and validation

283 The anonymous supplement contains the 95 archived HTML pages unchanged, normalised text,  
 284 the original per-page manifest, extracted records, three validation CSVs, an offline parser and a  
 285 verification script. The script checks hashes, replays extraction over all 95 pages, uses separate  
 286 HTML-card extraction for the 43 Conecta-labelled pages, reconstructs the sample and checks the  
 287 reported counts. It makes no network requests. Replaying a parser verifies reproducibility, rather  
 288 than independently validating its interpretation.

289 The offline checks reproduce accession-instruction extraction for all 43 Conecta-labelled pages and  
 290 literal endpoint, credential-term and Swagger/OpenAPI indicators in a fixed 20-page stratified sam-  
 291 ple. The sampling rule, archived inputs and historical comparison CSVs are retained in the sup-  
 292 plement; these comparisons assess extraction consistency. Literal matches, including authentication  
 293 terms in URLs or parameters, do not by themselves determine access policy.

## 294 C Worked reporting records

295 These records apply all six reporting fields to proposed evaluation designs. Public documentation  
 296 supports or contradicts specified access conditions; it does not supply an execution trace or an un-  
 297 observed institution’s approval. The examples deliberately retain those unknowns.

298 **1. Public flood-alert reading. Actor/delegation:** a developer reading public environ-  
299 mental data, with no claim to another person’s account or a privileged institutional role.  
300 **Operation/data/purpose:** read the current flood-alert list for a non-safety-critical display,  
301 preserving provider attribution. **Institution/route:** Environment Agency, HTTP GET on  
302 /flood-monitoring/id/floods. **Evidence/time:** the official reference captured on 5  
303 September 2026 states no registration requirement and gives this read route, attribution terms and  
304 a warning against safety-critical reliance [Environment Agency, n.d.]. **Harness:** proposed single  
305 anonymous HTTP GET; no supplied privileged credential, execution trace or measured response.  
306 **Status:** the stated access condition is supported by provider documentation. An executed actor-  
307 route attribution remains untested. A successful future execution would still not establish the time-  
308 liness or safety of an application built on the data.

309 **2. Citizen with a Conecta certificate tool. Actor/delegation:** a citizen acting on their own behalf,  
310 with no public-body delegation. **Operation/data/purpose:** obtain the documented criminal-record  
311 certificate response; no complete criminal history is claimed. **Institution/route:** Federal Police data  
312 through the Conecta API described by the catalogue. **Evidence/time:** archived API page, 29 Au-  
313 gust, and programme guidance captured 5 September 2026 distinguish the public consultation from  
314 public-body API accession and exclude direct citizen use [Governo Digital, Brasil, 2026a,c]. **Har-**  
315 **ness:** proposed direct Conecta tool with an evaluator-supplied credential. **Status:** the represented  
316 grant is contradicted by those conditions and should be reported as counterfactual authority. This  
317 does not prohibit the documented separate citizen consultation or establish its automation conditions.

318 **3. Eligible body without evidence of completed accession. Actor/delegation:** a federal or state  
319 public body within the certificate page’s stated classes; no particular body’s completed approval or  
320 agent delegation is supplied. **Operation/data/purpose:** certificate response for an unspecified pub-  
321 lic purpose, which remains an evidence gap. **Institution/route:** the same Conecta certificate API.  
322 **Evidence/time:** archived access instruction and the 5 September accession procedure requiring a re-  
323 quest and provider approval [Governo Digital, Brasil, 2026a, n.d.a]. **Harness:** proposed direct API  
324 tool, with credential provenance unspecified. **Status:** situated attribution unresolved. Eligibility, a  
325 plausible purpose and technical possession of a key do not document the missing authorisation.

326 **4. Municipality below the published size threshold. Actor/delegation:** Caxias do Sul, Rio  
327 Grande do Sul, the largest non-capital municipality below the threshold, with an IBGE population  
328 estimate of 479,604 on 1 July 2026 [Instituto Brasileiro de Geografia e Estatística, 2026]; no agent  
329 delegation is supplied. **Operation/data/purpose:** a municipal public-service lookup using the CEP  
330 entry on the programme’s municipal API list. **Institution/route:** Conecta’s documented CEP inte-  
331 gration route. **Evidence/time:** the archived CEP accession invitation and specific subnational guid-  
332 ance captured on 5 September 2026. The population fails both formulations of the programme’s size  
333 condition [Governo Digital, Brasil, n.d.b]. **Harness:** proposed CEP tool with a supplied Conecta  
334 credential. **Status:** size eligibility contradicted under the criteria current at capture; the grant can-  
335 not be treated as established municipal authority. Separate supplier routes, future policy changes or  
336 institutionally evidenced exceptions would require their own evidence. The two wordings differ at  
337 exactly 500,000, but no municipality in the 2026 table has that estimate. Meeting the size condition  
338 would still not establish completed accession.

339 **5. A municipal benchmark with assumed authority. Actor/delegation:** Porto Alegre’s municipi-  
340 pal administration; agent delegation assumed for a proposed benchmark. **Operation/data/purpose:**  
341 CEP lookup on synthetic addresses during simulated municipal service registration. **Institu-**  
342 **tion/route:** Conecta’s municipal CEP route. **Evidence/time:** catalogue invitation archived on 29  
343 August; subnational guidance [Governo Digital, Brasil, n.d.b] and accession procedure [Governo  
344 Digital, Brasil, n.d.a], both captured on 5 September 2026. Capital status meets the capital-or-  
345 population condition, although evidence of approval for this body, purpose and delegation is not  
346 supplied. **Harness:** proposed simulator with synthetic responses and an evaluator-supplied test cre-  
347 dential. **Status:** situated attribution unresolved; any eventual result would be conditional on the  
348 explicitly assumed authority. The example makes no claim about Porto Alegre’s actual approvals  
349 and is excluded from the 52 documentary records.

## D Evidence expansion and reproducible checks

The anonymous supplement’s `evidence-expansion` directory supplies frozen sources, coding decisions, exact excerpts, schemas and offline analysis. Raw captures in `evidence.zip` have a separate manifest of URLs, retrieval dates and SHA-256 hashes, preserving the 95 archived Brazilian pages unchanged alongside the subsequent programme, UK, Indian and standards documents. The analysis verifies 123 source files, checks every coded positive excerpt against its source, reconstructs the UK sample and matches the 14 municipal invitations to the subnational programme’s list. Title matching normalises case, accents, whitespace and punctuation, with a further singular/plural *Endereço/Endereços* difference for CEP; the manifest enumerates all inputs.

**Coding.** The 52 external-page texts and 35 UK descriptions were coded with assistance from Codex (OpenAI) using the page-local definitions below; the exact model identifier was not recorded. Exact excerpts and decisions are retained in the supplement, allowing their relation to the sources to be inspected. Because excerpt matching checks the presence of positive evidence without detecting omitted conditions, the author verification deliberately included initially silent descriptions, where false negatives would undercount documented requirements.

External-page codes distinguish three institutional-accession statements (payment-provider accreditation, DOU integration and WSDenatran), one service-registration requirement (PagTesouro collecting bodies), two operation-authentication descriptions, two credential-input descriptions, and one signing-user account condition. The primary page codes are disjoint; their underlying conditions need not be. DOU, for example, combines interface approval, a host firewall rule, publishing-user permissions and SOAP credentials. Appendix E retains these layers. A password input is not treated as proof of issued credentials. A beneficiary description, data licence, output token or authentication-looking URL alone does not establish an applicable access condition. The 43 remaining pages are coded only as having no explicit statement under these definitions on that page. Linked supplier documents were not exhaustively audited.

**Author verification.** With prior labels withheld, I recorded condition types, supporting phrases and governed actors or operations in two checks: six initial units and 25 expanded units, with one overlap, giving 30 unique units. These comprise all nine positive Brazilian external pages, 19 of 71 initially silent units (12 of 43 Brazilian pages and seven of 28 UK descriptions), and two UK descriptions with explicit conditions, one requiring registration and one waiving API keys. None of the 12 Brazilian silent labels changed, whereas one UK description changed to unclear, leaving 27 UK silent descriptions; all nine positive Brazilian primary codes were retained. Twenty-four of the 25 expanded coarse labels matched their prior labels, with DOU’s publishing-user registration adjudicated separately from organisational-service registration. Original labels, raw responses and adjudications are preserved separately so that the check can be inspected as purposive author verification, without being treated as an estimate of corpus-wide accuracy or independent inter-coder reliability.

**UK sample.** The 233-record CSV and its schema are pinned to revision `5ac787bd84b04ac9b3a163e704fc4470c01febb1`. For each of 35 distinct provider values, select the record with the lexicographically smallest hexadecimal SHA-256 of the UTF-8 string `2026-09-05|` followed by its URL. The sample covers 35 of the 233 records (15.0%) and every distinct provider value in that CSV, giving each provider equal weight irrespective of its number of records. The separate 133-row generated-source CSV is outside this fixed sampling frame and is not pooled into the denominator. All selected descriptions were read. The seven explicit-condition descriptions concern the Environment Agency, Food Standards Agency, Local Government Association, Office for National Statistics, Scottish Government, Transport for London and UK Hydrographic Office. The first, second, fourth and fifth explicitly waive a requirement; the other three state a requirement. The latter’s paid tier is a condition on the described service, not a claim about all hydrographic data. The counts describe this fixed, provider-balanced sample; they are not national prevalence estimates or an exhaustive audit of all linked documentation.

**Participation records and policy scope.** The programme FAQ’s portal link led to two public dataset pages [Ministério da Gestão e da Inovação em Serviços Públicos, 2026b,a]. The accession page lists 13 resources labelled from December 2024 through June 2026, described as identifying

adhering bodies and the datasets whose managing body authorised access. The consumption page lists six resources for January–June 2026, described as transaction counts by body and API. These are resource descriptions, not inspected rows. Download or preview attempts for selected accession resources did not yield file contents; the portal itself marks the dataset as outdated and some resource URLs as unavailable. The dated rendered-page captures preserve these observations in a separate evidence addition. Neither retrieval failure nor catalogue silence establishes that approval evidence is unpublished. If recovered and applicable, a body–dataset approval row could support that part of an attribution; it would require additional evidence to establish agent delegation and the exercised operation.

The subnational FAQ qualifies its capital-or-population restriction as *no momento* and attributes it to programme governance. It does not provide a future admission date or establish a permanent exclusion. The cited pages and the programme’s legislation index did not identify a verified normative instrument establishing this particular Conecta threshold. Related national management-programme ordinances were located but their full texts could not be retrieved. The analysis therefore treats the threshold as current published programme guidance; it makes no finding about its statutory basis, permanence or legality.

**API Setu consistency.** This is a purposively selected consistency check of central platform guidance; no equivalent specification audit was performed across the other two catalogues. The SOP landing page links a five-page document as V1 (filename ending `v0.pdf`) and a 31-page manual as V2 (filename ending `v1.pdf`); both were read completely. The V2 security block on page 19 was transcribed and visually checked against the PDF. Its two objects are `clientId: []` and `apiKey: []`, each a separate array member. The complete sample has no operation-level security override. The prose on page 17 and header comments on page 25 require both credentials. Table 3 applies OpenAPI 3.0.1 semantics assuming each present credential is valid; it does not test credential validity, a server implementation, or other publishers’ authentication methods.

Table 3: Exhaustive documentary truth table for the published API Setu example. A single requirement object containing both schemes is a negative control and agrees with the prose in all four assignments.

| Client ID | API key | Prose: both | Sample: alternatives | Disagreement |
|-----------|---------|-------------|----------------------|--------------|
| Absent    | Absent  | False       | False                | No           |
| Present   | Absent  | False       | True                 | Yes          |
| Absent    | Present | False       | True                 | Yes          |
| Present   | Present | True        | True                 | No           |

**Artefact scope.** The original `x-accession.schema.json` is a study-fixture schema for the 43 catalogue projections. It rejects replacing their unresolved status with supported, asserting unobserved completed accession or dropping a municipal qualification. Those restrictions apply to this evidence set. The separate `x-accession-report.schema.json` defines a reusable reporting record with all three statuses. A supported claim must identify actor, operation/purpose, institution/route and harness; provide dated, locatable evidence; explain authorisation and delegation or their inapplicability; and connect that evidence to the claim. An executed-attribution claim additionally requires execution evidence. Structural conformance checks field completeness and evidence references; a reviewer must still assess authenticity, applicability, scope, conflicts and current validity. Neither schema issues permission.

**Population calculation.** The supplement retains IBGE’s 2026 workbook, its hash, a row-level classification and an offline reproduction script. Of 5,571 coded workbook entries, Brasília (5300108) and Fernando de Noronha (2605459) are excluded from the municipal comparison, leaving 5,569, of which 26 are state capitals matched by state and name. The criterion is capital status OR population at least 500,000; the strict-greater-than variant yields the same set because no included estimate equals the threshold. Fifty-two municipalities meet this condition and 5,517 do not (99.066%).

The 52 municipalities contain 65,056,729 residents, 30.370% of the workbook’s national estimate of 214,211,951. This population denominator includes the Federal District and Fernando de Noronha,

447 whereas the municipal-administration denominator excludes those two entries. The 5,569 municipi-  
 448 palities contain 211,198,591 residents in total, and the complete municipality worksheet sums to the  
 449 national total. The 5,517 municipalities outside the condition therefore contain 146,141,862 resi-  
 450 dents, or 68.223% of the national population. Reporting both denominators distinguishes the dis-  
 451 tribution of administrative eligibility from population coverage without treating either as observed  
 452 service access.

453 Four non-capital municipalities are within 10% below the threshold (450,000 to fewer than 500,000  
 454 residents): Caxias do Sul (479,604), Mogi das Cruzes (472,255), São João de Meriti (466,487) and  
 455 Jundiaí (465,277). Caxias do Sul is the largest below the threshold, falling short by 20,396 residents  
 456 or 4.08%. Population estimates are dated 1 July 2026 and are used here to apply published guidance;  
 457 they are not evidence of the population measure used in a particular accession decision. The cal-  
 458 culation concerns the capital-or-population condition, with other requirements, possible exceptions,  
 459 alternative routes and completed approvals outside its scope.

## 460 **E External cases: conditions with different actors and scopes**

461 All nine positive external pages were assessed using the six fields; the supplement retains full records  
 462 with exact excerpts, source hashes and archive dates. The source-described actor classes leave  
 463 individual delegation and execution unresolved, while the actor and operation needed to assess each  
 464 condition remain identifiable. The table documents these distinctions: a citizen's signing-account  
 465 condition cannot establish integrator access, and credential-maintenance permission cannot stand  
 466 for a data-read grant. Missing evidence therefore has an actor and an operation even when every  
 467 attribution has the same status.

Table 4: Scope distinctions retained in the nine source-derived reporting records. Every row refers to an archived catalogue page dated 29 August 2026.

| Page                          | Condition and actor                                                                                  | What attribution still needs                                                                                     |
|-------------------------------|------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------|
| Consulta de Serviços Públicos | Authentication for the all-services operation                                                        | Principal and credential provenance; other reads have different descriptions                                     |
| Serviços Estaduais            | Login plus token inputs for inserting, editing and consulting services                               | Authority of the represented state-service manager and delegated operation scope                                 |
| Assinatura Digital Avançada   | Citizen's verified or proven gov.br account; basic account excluded                                  | Separate API-integrator eligibility and authority to sign for that citizen                                       |
| SIOP Execução Orçamentária    | Credential-object input; operation described for SDI                                                 | Issued credential and approved scope; audience alone is insufficient                                             |
| SIOP Credencial               | Username/password inputs for credential maintenance                                                  | Right to modify that account's credentials; this is not a data-read grant                                        |
| PagTeseuro PSPs               | Treasury accreditation and Central Bank authorisation of payment provider                            | Actual accreditation and delegation; provider-hosted routes and Treasury notifications have different directions |
| PagTeseuro collecting bodies  | Service registration in SISGRU; users use SIAFI credentials                                          | Registration of the particular service and credential/delegation scope                                           |
| DOU publishing                | Interface approval, host firewall rule, publishing-user registration/permission and SOAP credentials | Evidence for each layer; one approval cannot stand for all four                                                  |
| WSDenatran                    | Denatran authorisation and Serpro contracting; APIs depend on service contracted                     | Particular agreement, contracted service and agent delegation                                                    |