

INDUSTRY INSIGHT

Where Community Codes of Conduct Fall Short

What digital leaders found when they audited their own codes of conduct

An evidence-based audit of community codes of conduct, completed by participants in the DLC Leadership Program pilot cohort using an 18-criterion assessment rubric.

Rachel Kowert, PhD
Program Director
DLC Leadership Program

Kelli Dunlap, PsyD
Director of Learning and Community
DLC Leadership Program

Executive Summary

Most conversations about community safety focus on what platforms should build or what policies they should write. Far less attention is paid to the documents that actually sit between a community and its members: the code of conduct. It is the single document that every member is expected to read, agree to, and be held against. It is a critical component of players' social onboarding to a community, outlining what is expected of them and the consequences of not meeting these expectations.

As part of the DLC Leadership Program pilot, participants audited a code of conduct from a community they lead or belong to, scoring it against an 18-criterion, 54-point rubric built directly out of the EFFECT Framework¹, a practical tool distilling six evidence-based design principles for developing codes of conduct. Total scores place a code in one of five bands: Inadequate, Needs Improvement, Developing, Proficient, Exemplary.

We found that:

- Most codes are “Developing.” The average score for an audited community code of conduct was 33 out of 54, which translates to “adequate, with clear room to improve.” Seven in ten audited codes fell below the “Proficient” threshold, meaning the typical audited code is missing multiple elements the evidence says make a code of conduct work.
- What communities do best is write the document and make it findable. Language, structure, and presentation were the strongest areas overall.
- Communities are not reinforcing or revising their codes. Keeping the document alive after publication (revisiting it, reporting on enforcement, giving members a say) was ranked the weakest area by four in ten participants, more than double any other. In most cases the practices existed; they were simply never written down.
- Training was a deciding variable. Participants who had prior training, particularly training around creating codes of conduct, had higher-scoring codes. Sector, seniority, community size, and years of experience made no reliable difference at all. However, we cannot conclude whether trained participants were working with better documents or were simply better at spotting what was missing in them.
- Leaders cannot fix what they have not been taught to look for. Every participant named a specific area to improve in their code of conduct, and several described gaps they had lived with for years without noticing. This finding is an important reminder that experience cannot teach you what you have never been prompted to look for – training does.

Taken together, these findings point to something that many of us know but do not enact. Community codes of conduct are not legal artifacts or compliance checkboxes – they are the one piece of trust and safety infrastructure that reaches every member of a community and commonly treated as something you write once and move on from. This report is a first attempt at measuring where the field actually stands. It is a baseline, not a verdict, and it exists so that the next attempt can be better.

About the audit

Participants were given a rubric to audit a community code of conduct of their choosing. This rubric was developed to operationalize the EFFECT Framework's six design principles² into 18 scorable criteria, grouped into four sections for assessment. Each criterion is scored 0–3, for a maximum of 54 points.

Section	Criteria	Points	What it assesses
I. Accessibility & Discoverability	4	12	Language clarity, information architecture, visual design, and whether the document can actually be found
II. Content Specificity & Clarity	6	18	Behavioral definitions, contextual examples, scope, consequences, enforcement process, and appeals
III. Positive Community Building	4	12	Desired behaviors, values integration, bystander empowerment, and diversity protections
IV. Implementation & Community Engagement	4	12	Ongoing reinforcement, community agency, operational transparency, and continuous improvement

Participants selected a code of conduct from a community they lead, work on, or belong to, scored it against the rubric, and then answered a short set of reflection questions: 1) what was their total score, 2) what was their performance band (Inadequate, Needs Improvement, Developing, Proficient, Exemplary - outlined in more detail below), 3) to rank the four sections from strongest to weakest, 4) name their three lowest-scoring criteria, 5) identify which gap they would address first, and 6) share what surprised them.

Because participants chose their own documents, this is a self-selected sample of communities rather than a representative survey of the industry. It is best read as a picture of what experienced digital leaders see when they audited the codes governing their own communities.

Who completed the audit

A total of 48 participants opened the community code of conduct audit, and 43 submitted responses. One submission from program staff was a test entry and has been removed. The following data is drawn from the remaining 42 responses. Every respondent in this analytic sample could be matched to a pilot pre-assessment record, so the demographic profile below describes the audit respondents specifically, not the pilot cohort as a whole.

Geography and background

Characteristic	Detail
Countries represented	9 countries across 5 continents
Largest single group	United States (69%), followed by the UK (7%) and Canada (7%)
Age	Mean 34.8 years, median 35, range 19–53
Gender	48% male, 45% female, 2% non-binary or gender non-conforming, 5% preferred not to answer
Education	83% hold a bachelor's degree or higher; 60% hold a postgraduate degree

Roles and experience

Respondents were experienced and, in most cases, wearing more than one hat. Percentages exceed 100% because participants could select multiple roles.

Role	%
Community manager or community specialist	45%
Researcher or academic	38%
Community moderator (volunteer)	24%
People manager or team lead in a gaming/tech context	24%
Content moderator	21%
Game designer or developer	21%
Policy, advocacy, or nonprofit professional	21%
Streamer, content creator, or influencer	19%
Trust and safety professional	17%
Community moderator (paid)	12%

- **Current leadership:** 71% currently lead at least one digital community.
- **Tenure:** 52% have led digital communities for 4–7 years; 19% for 8 years or more.
- **Breadth:** Half hold a leadership position in more than one digital community at the same time.
- **Scale:** 36% lead a community of more than 1,000 members; 17% at 10,000 or more.
- **Platforms:** Discord was the most common platform utilized (71%), followed by YouTube (31%) and Twitch (29%).
- **Organizations:** Academic or research institutions (26%), nonprofits (24%), and game developers or publishers (21%).

Experienced, credentialed, and largely untrained

More than half of respondents (52%) report receiving no formal training, or very little, for their current digital leadership role. On the specific skill of creating community guidelines, 57% rated their formal training at the lowest two points of a five-point scale, and 81% had never taken any third-party training on the topic.

This is the pattern the DLC Leadership Program exists to address: a workforce with real seniority and real responsibility that has built its expertise almost entirely on the job.

Findings

The average code of conduct is adequate, with clear areas for improvement

The average score reported for a community code of conduct among participants was 33 (SD = 9.8) out of 54 points, with a range of 12 to 50. This places the average audited document at the very bottom of the “Developing” band (described as “adequate, with clear areas for improvement”). The overall distribution was wide, but skewed low.

Performance band	Description	Share of documents
Exemplary (49–54)	Industry-leading code of conduct demonstrating best practices across all areas	5%
Proficient (41–48)	Strong code of conduct meeting most best practice standards	23%
Developing (33–40)	Adequate code with clear areas for improvement	26%
Needs Improvement (25–32)	Basic code requiring significant enhancement	28%
Inadequate (0–24)	Code failing to meet minimum standards for effective community governance	18%

Among all of the documents audited, 72% fell below the “Proficient” threshold. Four in ten (46%) landed in the bottom two bands (“Needs Improvement” and “Inadequate”). Only 5% of the codes of conduct audited qualified as “Exemplary”.

These are the documents that govern game studios, platform companies, university programs, professional associations, and Discord servers with tens of thousands of members. And across that range, this audit found that they are, on average, barely adequate.

The audit that could not be completed

One respondent, a researcher with a doctoral degree and more than twelve years of digital leadership experience, set out to audit a gaming Discord she has belonged to for over a decade. She could not find the community guidelines anywhere. She knew the document existed, because every member had agreed to it on joining, and the community had migrated from an older bulletin board where it was originally posted. Nothing was pinned, nothing was linked, and she could not locate the original site to check.

She scored the audit a zero, on the grounds that a code of conduct nobody can locate governs nobody. Her score is excluded from the average below rather than counted as a zero, because it measures a different failure from the one the rubric was built to assess. It is worth naming on its own terms: this is a well-credentialed, highly experienced community member who could not produce the document she had personally agreed to abide by.

Many communities fail with the follow through

Participants were asked to report how their audit highlighted their strongest to weakest categories.

Section	Ranked strongest	Ranked weakest
I. Accessibility & Discoverability	33%	14%
II. Content Specificity & Clarity	29%	21%
III. Positive Community Building	29%	21%
IV. Implementation & Community Engagement	10%	43%

Accessibility & Discoverability was the strongest section – it was ranked first by a third (33%) of participants and last by only 14% of respondents. Communities are, broadly, writing readable documents and putting them somewhere findable.

Implementation & Community Engagement was the weakest by a wide margin. The largest group of participants (43%) ranked this category as their weakest, more than double any other section. This is the section that asks whether the code is reinforced after onboarding, whether members have any say in shaping it, whether enforcement is ever reported on, and whether the document is reviewed and updated.

Positive Community Building and Content Specificity and Clarity landed somewhere in the middle.

Taken together, these rankings describe a familiar lifecycle: a code of conduct gets written with real care, published, and then rarely if ever addressed again.

Note: These rankings were reported by the participants and not computed from criterion scores.

The biggest gaps are around accountability, not rules

When asked to list their three lowest-scoring criteria, 86% named specific rubric criteria and could be coded; the remainder named whole sections rather than criteria, or described their findings generally. The table on page 7 shows the share of coded responses that named each criterion among their lowest-scoring.

The four most-cited lowest-scored criteria are all about what happens after a rule is written: how a member challenges a decision, whether members have any hand in shaping the standards, whether the document is ever revisited, and whether consequences are spelled out. The four least-cited are all about the rule-writing itself: scope, visual design, language clarity, and behavioral definitions.

Community codes of conduct are doing a good job of saying what is not allowed but markedly worse at explaining what happens next, who decides, how to push back, and how any of it changes over time.

It is also worth highlighting that every participant named a specific area they would address first. Some named a criterion, some a whole section, and one named a fix she expected to meet resistance, but nobody finished the audit without a next step. The gaps were not hard to find once there was something to find them with.

Where you work doesn't predict quality; training does.

At the start of the pilot, participants rated their confidence in creating community guidelines on a five-point scale. Among those who completed the audit, 55% rated themselves confident (4 or 5). That confidence correlated positively, albeit moderately, with audit score ($r = .37, p = .022$): participants who came in confident scored their documents higher than those who did not.

Section	Criterion	% Ranked Weakest
Section I. Accessibility and Discoverability	Integration & Discoverability	14%
	Language Clarity	6%
	Information Architecture	6%
	Visual Design	3%
Section II. Content Specificity and Clarity	Transparent Appeal & Recourse Process	36%
	Consequence Clarity	28%
	Transparent Enforcement Process	22%
	Game-Contextual Examples	22%
	Behavioral Definitions	6%
	Scope Clarity	3%
Section III. Positive Community Building	Positive Behavior Emphasis	25%
	Bystander Empowerment	22%
	Diversity & Inclusion	8%
	Community Values Integration	6%
Section IV. Implementation & Community Engagement	Community Agency & Ownership	33%
	Continuous Improvement	31%
	Code Reinforcement	19%
	Operational Transparency	19%

The participants who came in least confident about writing community guidelines were found to be working with the weakest documents. They averaged 27.9 against 36.7 (out of 54) for the confident group. Where confidence was lowest, the codes of conduct were furthest from where they needed to be.

What this data cannot say is why. The audit did not ask about authorship, so whether confident leaders wrote those codes, inherited them, or improved them over time remains open. However, what it does show is a ceiling. Even the confident group averaged only 36.7 out of 54 – the middle of the “Developing” band and short of “Proficient”. The people who felt most equipped to write community guidelines were working with documents that had meaningful gaps, and several only recognized those gaps once they scored the document against the rubric.


Confidence, in other words, is built on something. For most participants, that something was years of practice rather than instruction. For some, it was formal training; however, this was scarce across the board as just under half reported any substantial formal training of any kind for their current role. Only 29% reported meaningful formal training in creating community guidelines specifically, and fewer than one in five had ever taken third-party training on the topic.

That raises an obvious question. If the documents vary this widely in quality, what explains the spread?

The intuitive answer would be resourcing. A well-funded studio with a dedicated trust and safety function should produce a better code of conduct than a volunteer-run Discord. So we tested it. Audited codes were compared across sector, role, community size, organization size, and years of experience. Analyses revealed that none of these variables separated the groups (none reached statistical significance).

The largest gap between any two groups was under four points, and game developers and academic institutions were separated by only 2 points.

However, one variable was found to significantly separate the groups - whether or not the leader had been trained. Participants with some formal training for their role averaged 36.6 out of 54 on their auditing rubric, against 30.3 for those with none or very little, which is a gap of roughly six points. The effect was larger still for training specific to the task. Participants who had taken third-party training on creating community guidelines averaged 39.0, against 31.5 for those who had not. That seven and a half point gap is the largest of any comparison in this analysis, and the highest average of any group in this analysis.



The same caveat that applies to confidence applies here. What this shows is that trained participants gave their codes of conduct higher scores against the rubric. Whether that is because they wrote better documents, inherited responsibility for better ones, had a hand in shaping them over time, or simply read the rubric more accurately, this data set cannot say. The audit asked participants to score a document, not to account for their relationship to it.

The spread matters as much as the average. Scores among trained participants clustered more tightly (SD = 6.7) than among untrained participants (SD = 11.0). Trained participants did not just score higher on average; they also scored more consistently, and none of their codes fell into the Inadequate band. Every code in that band came from someone with little or no formal training.

Training is the key variable

A studio with a trust and safety team and a volunteer-run Discord produced codes of roughly equal quality. Trained and untrained community leaders were not working with codes of equal quality.

Sector, seniority, and scale are fixed features of an organization. Training is the one input that can be changed – and in this data, it is the only one associated with a better document. It isn't the studio. It's the training.

What participants said

As part of this audit, participants also provided written reflections. Some of these quotes are shared below. Quotes are lightly edited for length, attributed by role rather than by name, and have had identifying references to specific organizations removed.

The practice exists. The documentation does not.

This was the single most common realization, and it reframes the whole dataset. Participants repeatedly described mature internal processes that were entirely invisible to the members those processes govern.

“

We have a wonderful appeal bot and system for processing appeals, but we don't disclose that in our guidelines. We include the link when we process a ban, but we have not communicated that this is available. That's an easy fix I can add ASAP.

COMMUNITY LEAD
Game studio, 4–7 years experience

“

There's room to improve everywhere. The gap in the punishment and appeal process was very apparent, because it implies a path for appeal but doesn't really mention or explain how to do it.

COMMUNITY MANAGER & MODERATOR
4–7 years experience

The distinction matters for anyone deciding where to spend resources. Building an appeals system is expensive. Documenting the one you already have is not.

Abstract rules, no examples

Game-contextual examples were named among the lowest-scoring criteria by 22% of coded responses, and the reflections suggest the omission is usually invisible until someone points at it.

“

I'm actually amazed that we had no concrete examples of different behaviors and only had abstract descriptions of them. This seems like a very easy but impactful thing to improve.

COMMUNITY MANAGER
4–7 years experience

“

I hadn't thought about direct game examples at all. We clearly define our game rules versus community rules and which apply where, but we never specifically give examples of what would be a broken rule and what wouldn't. We have certain rules about [in-game activity] that would benefit from examples.

COMMUNITY MANAGER & GAME DESIGNER
4–7 years experience

The document nobody can find

Several participants set out to audit a code of conduct and discovered the more fundamental problem first.

“

I was surprised how hard it was to find the code of conduct when seeking it out.

RESEARCHER
No digital leadership experience

“

There was a whole announcement channel but nothing pinned or anything to the original guidelines post.

RESEARCHER
12+ years experience

When a code of conduct is genuinely unfindable, every other criterion in the rubric is moot — the quality of the writing, the examples, the appeals process, none of it reaches anyone.

The constraints are often organizational

A meaningful minority of participants identified gaps they had no authority to close. This is important context for any reading of the scores as a measure of individual capability.

“

This audit is difficult for a large platform to do, especially one whose specific point of view is that guidelines should be top-down and specifically not community generated. Despite it being research-backed, senior leadership prefers the top-down approach and that informs what we can and can't change.

TRUST AND SAFETY PROFESSIONAL
1-3 years experience

“

We aren't allowed to be that transparent and that should potentially be a consideration to be changed.

RESEARCHER
12+ years experience

Others cited legal constraints on disclosing disciplinary outcomes, institutional inertia, and simple lack of seniority. Training a community leader well does not, on its own, give them the mandate to act. That is an argument for engaging organizations, not only individuals.

You don't know what you don't know

The most consistent thing participants reported was not a specific gap; it was the experience of discovering gaps they had no framework for noticing.



A gap I didn't really consider was how flimsy the actual community code of conduct is. It's really relying on social norms to maintain order within the community.

VOLUNTEER MODERATOR
8-12 years experience



I think the sheer amount of things to think about. This will be very useful to look back at and think what can we do next. Especially for our upcoming game. We do a lot of diverse things and make sure that is important, so it is interesting to see it listed here.

COMMUNITY MANAGER AND GAME DEVELOPER
4-7 years experience



The scope clarity, because the code of conduct did only apply for Discord until now. It became clear to me that it is important to have a code of conduct that also is applicable to social media etc.

COMMUNITY MANAGER
Game studio, 1-3 years experience



Having gone through the code of conduct as it is today, I realized how out of date it currently is. Whilst I haven't been a member of the community for a long period, I can see that, already, the code is becoming increasingly out of date as the months go on.

COMMUNITY MEMBER
Not currently in a leadership role

None of these are failures of effort. A code of conduct that covers Discord but not social media was written by someone who had not been given a reason to think about scope. A code quietly going stale was written by someone with no prompt to schedule a review. A document leaning on social norms to do its enforcement work looks fine right up until the norms shift.

What changed for these participants was not their commitment or their experience. It was that someone handed them a list of things to look for.



It's sort of shocking how easily these things are overlooked and, as a member, you might not even think about it. It was also a lot worse than I would have expected initially, and this caused me to reflect further on how to better structure and maintain my own groups.

GAME DEVELOPER AND RESEARCHER



I was surprised by how much clearer community guidelines become when you just ask the right questions.

TRUST AND SAFETY PROFESSIONAL
8-12 years experience

Good news, everyone!

The most encouraging finding in this dataset is how small most of the fixes are. Participants were not describing communities that need to be rebuilt. They were describing appeal processes that exist but are unpublished, examples that were never added, review cycles that were never scheduled. Give an experienced community leader a rubric and a few hours, and they will find the gaps themselves.

Moving Forward

What you can do now

- **Publicly document what you already do.** Most of what is missing comes after the rules and most of it is missing on paper rather than in practice. Before designing a new appeals process, check whether the existing one is written down where members can see it.
- **Add examples.** Game-contextual examples were among the three weakest criteria for 22% of participants, and they are among the easiest gaps to close. Several described the omission as obvious the moment they looked for it.
- **Put a review date on the document.** Continuous improvement was among the three weakest criteria for 31% of participants; a scheduled review is the smallest possible version of that fix.

What needs organizational support

- **Invest in reinforcement.** Implementation is the weakest layer. There is room to improve the policy itself, but resources aimed at rewriting guidelines are likely to produce less overall improvement than resources aimed at reinforcement, transparency, and review.
- **Adopt a shared standard and build a route to escalate it.** Frontline staff can see the gaps and frequently cannot close them. Several participants named specific improvements blocked by leadership preference or legal caution. Two things are missing in most organizations: a shared standard to assess codes of conduct against and a route for reaching someone with the authority to act. Without the first there is nothing specific to escalate. Without the second, the assessment stops at the person who undertakes it.
- **Change how community codes of conduct are valued.** Codes of conduct are increasingly an accountability surface. Regulators, players, and press all read them as a statement of what a platform will do. Documents that sit at the bottom of the “Developing” band are a weak foundation for that.
- **Train leadership, not just the frontline.** Training, like that provided by the DLC, is important at leadership levels, not only frontline ones. The people who approve a code of conduct need the same grounding as the people who enforce it – what makes these documents usable, what a member actually needs to see, and why a governance decision was made determines whether any of it works on the ground. Several participants named specific fixes they could describe precisely and could not authorize.

What the field needs to build

- **Prioritize training that is specific.** Training was the only significant predictor of higher quality codes of conduct. Sector, seniority, community size, organization size, and years of experience made no reliable difference to how a code of conduct scored; whether the person had been trained did. The effect was largest for training specific to writing community guidelines rather than general to the role.
- **Prioritize the education layer of trust and safety.** It is doing work that tooling and policy cannot. Every participant identified a specific area to improve first; not one came away with nothing to change. A single structured exercise turned forty-two documents into forty-two improvement priorities, at the cost of a rubric and an afternoon.

Concluding Thoughts

This audit began as a teaching exercise. Participants in the pilot were asked to take a document they had probably not read closely in years and look at it properly — not to grade their communities, but to make a point that a community code of conduct is a valuable lever of trust and safety that is not being optimized as such. It is a piece of trust and safety infrastructure that reaches every member of a community and it is almost always the most public-facing community governance document.

Existing research tells us that codes of conduct generally are unevenly adopted³, frequently vague about enforcement⁴, and often poorly understood by the people they govern⁵. What has not been examined is what happens to these documents after publication. It remains unclear whether they are ever revisited, by whom, how often, or in response to what.

The data here reiterated many of these points. In most cases, what participants found missing was not the rules themselves but everything that is meant to come after they are drafted: the reinforcement, the appeals, the revisiting, the reporting back. The weakest section of the rubric, by a wide margin, was the one asking whether anyone had returned to the document since publishing it. That is the write-it-and-forget-it pattern showing up in the numbers, in a cohort of experienced digital leaders who had every reason to know better.

To measure the current state of community codes of conduct, we developed an assessment rubric that operationalizes the six design principles of the EFFECT Framework⁶. To our knowledge it is the first attempt to make community codes of conduct assessable against evidence-based standards rather than legal convention or institutional habit.

That said, the rubric is a new instrument. It has not been validated against independent expert scoring, the performance bands are reasoned rather than empirically derived, and whether eighteen is the right number of criteria remains an open question. Scores should be read as a structured description of a document rather than a calibrated measure of its effectiveness.

Participants also surfaced a tension the rubric does not yet resolve. One put it directly:

“

The rubric itself is great, but I think some guidance on how to include all of this without the community guidelines becoming too verbose would be helpful! I worry about declining literacy rates — I know that our guidelines often don't get read, even when they are very available.

COMMUNITY MANAGER
Game studio, 4-7 years experience

This is certainly the case. Sixteen of the eighteen criteria ask a code of conduct to contain more: more examples, more scope, more process, more transparency. Two questions ask it to remain clear and readable. Scored naively, the rubric rewards length, and length is the enemy of the thing a code of conduct exists to do⁷. A comprehensive document nobody reads is no better than a vague one. Resolving that tension (through layering, progressive disclosure, or delivering rules at the moment they become relevant rather than all at once) is the most important work the next version of this instrument has to do.

The potential verbosity problem is one flaw in a first attempt, and it will not be the only one. This work stands as a baseline, established with a new instrument on a small self-selected sample, in a field that has not previously measured this at all. Everything above should be read that way: not as a verdict on community codes of conduct but as a first attempt at establishing where the field stands.

However, there is one finding here that has nothing to do with the measurement tool itself. Most participants found gaps in their own documents in a single afternoon, many of which they had lived with for years without noticing. This is likely a reflection of the fact that most digital leaders have no formal training in this space and instead built their expertise entirely on their own. Practice teaches you to handle what has already happened to you, but it does not teach you what you have not yet encountered. It gives you no way to ask the questions you do not know to ask.

As discussed earlier, one participant discovered that the code they had been working with covered their Discord server but nothing else — not the social channels or the wider community. When you are running a space, the boundaries of that space feel like the boundaries of the job. Nothing in the day-to-day corrects that.

It is then unsurprising that training was the one variable in this study that separated stronger codes from weaker ones, and the effect was largest when the training was specific to the task rather than general to the role. The people who had been taught what to look for were working with documents that held up better under scrutiny.

The DLC Leadership Program exists for this reason. Community-facing work has professionalized in every respect except preparation: the roles have titles, the teams have headcount, the decisions carry regulatory weight, and the people making them are still expected to figure it out as they go. This audit is a small demonstration of what changes when they are provided education to compliment their experience. One rubric, one exercise, a few hours — and a cohort of experienced digital leaders came away with a specific, named, actionable list of what their communities were missing.

Community codes of conduct have been treated as paperwork. This report is an argument for treating them as infrastructure and infrastructure is only as good as the people trained to build it.

Appendix

Methodological notes and limitations

Raw counts appear in this section only. Percentages throughout the rest of this report were calculated against the denominators recorded here.

Sample

- 48 participants opened the audit assignment; 43 submitted responses; 5 submitted blank. One program-staff test entry was removed, giving an analytic sample of 42.
- All 42 matched to a pilot pre-assessment record by email, so demographic and baseline-confidence figures describe this subgroup rather than the full pilot cohort.
- Denominators vary by question: 39 reported a numeric score, 42 completed the section ranking, and 36 named codeable criteria. Percentages are calculated against the relevant denominator and should not be compared across questions without noting which applies.

Scoring

- Scores are self-assessed. Participants applied the rubric to documents they chose, without inter-rater checks. Both generosity and severity bias are plausible, and one participant's self-reported performance band did not match her own numeric score.
- Three participants described their findings without producing a total. One of those could not locate her community's guidelines at all and assigned a score of zero; that response is described in the findings but excluded from the mean rather than imputed, because it measures the absence of a document rather than a property of one. Including it as a zero would lower the mean to 32.3.
- Section rankings are ordinal and were reported by participants rather than computed from criterion scores. Because Content Specificity & Clarity carries 18 points against 12 for the other three sections, participants who ranked by raw points rather than proportion would have advantaged it. The two sections at the extremes – Accessibility & Discoverability and Implementation & Community Engagement – are unaffected by this, and the criterion-level data independently supports the same conclusion.
- The lowest-scoring criteria question was open text and coded manually against the 18 rubric criteria. 36 of 42 responses were codeable; six named sections rather than criteria or were too general. Four participants named more than three criteria; all named criteria were counted.

Generalizability

- Participants selected their own documents, so this is not a representative sample of gaming communities. It is a sample of communities that experienced digital leaders chose to examine.
- Roughly three in ten respondents (29%) were not currently leading a community and audited a document from a community they belong to or previously worked on. Their reflections are included; the distinction is worth carrying into any interpretation of the implementation findings.
- The sample skews North American (76%) and academically credentialed (83% bachelor's or above), consistent with the pilot cohort overall.

The Code of Conduct Assessment Rubric is proprietary to the DLC Leadership Program. For questions about this brief or the underlying data, contact rachel@dlcleadership.org.

The DLC Leadership Program is the first evidence-based certification program for community-facing professionals and the people who support them: moderators, community managers, trust and safety practitioners, player support and customer service teams, and platform volunteers. Anyone whose work puts them in direct contact with players, or whose decisions shape how a community behaves.

References

1. Kowert R. Community codes of conduct as proactive intervention for toxicity in games. *Front Public Health*. 2026 Jul 14;14:1842010. doi: 10.3389/fpubh.2026.1842010. PMID: 42523818; PMCID: PMC13407757.
2. Kowert R. Community codes of conduct as proactive intervention for toxicity in games. *Front Public Health*. 2026 Jul 14;14:1842010. doi: 10.3389/fpubh.2026.1842010. PMID: 42523818; PMCID: PMC13407757.
3. Singh, V., Bongiovanni, B. and Brandon, W. (2021) 'Codes of conduct in open source software—for warm and fuzzy feelings or equality in community?', *Software Quality Journal*, 30(2), pp. 581–620. doi:10.1007/s11219-020-09543-w.
4. Tourani, Parastou & Adams, Bram & Serebrenik, Alexander. (2017). Code of conduct in open source projects. 24–33. 10.1109/SANER.2017.7884606.
5. Quint, F., Y. Theocharis, M. Nahrgang, N. B. Weidmann, and M. E. Roberts. 2026. "Does the Community Understand the Community Guidelines?." *Policy & Internet* 18: e70045. <https://doi.org/10.1002/poi3.70045>.
6. Kowert R. Community codes of conduct as proactive intervention for toxicity in games. *Front Public Health*. 2026 Jul 14;14:1842010. doi: 10.3389/fpubh.2026.1842010. PMID: 42523818; PMCID: PMC13407757.
7. Quint, F., Y. Theocharis, M. Nahrgang, N. B. Weidmann, and M. E. Roberts. 2026. "Does the Community Understand the Community Guidelines?." *Policy & Internet* 18: e70045. <https://doi.org/10.1002/poi3.70045>.