Message from the NSDI '26 Program Co-Chairs

Welcome to the 23rd USENIX Symposium on Networked Systems Design and Implementation (NSDI '26) held in Renton, Washington, May 4–6, 2026. We are delighted to present a program that reflects both the depth and breadth of the networked systems community. The papers in this year's conference represent significant advances across the networking stack, spanning foundational systems design, networking abstractions, and emerging cross-disciplinary directions. We hope that NSDI '26 leaves you energized by new ideas, inspired by technical excellence, and reconnected with a vibrant and evolving research community.

Growth of the Community and Program Scale

NSDI continues to experience sustained growth in both submissions and topical diversity. In 2026, we received a total of 659 submissions across the Spring and Fall cycles, reflecting continued expansion of the field and strong engagement from the global research community. Of these, 150 papers were accepted, resulting in an overall acceptance rate of 22.8% The acceptance rate is a notable increase relative to prior years and is a consequence of deliberate changes in both process and philosophy. We believe that extremely low acceptance rates are not healthy for the community. A moderate acceptance rate protects authors from participating in submission roulette and protects reviewers from re-reviewing submissions that are largely unchanged. Papers accepted at NSDI generally represent the highest quality networked systems research in the world and we encourage the community to continue to keep up a healthy acceptance rate.

The NSDI conference now operates at a scale that supports three tracks, enabling us to accommodate more high-quality work while preserving depth and rigor in technical discussions. Our goal was not simply to grow the program, but to do so responsibly, maintaining the high standards expected of NSDI while recognizing that the community has grown and is now able to produce more impactful work than ever before. We are proud that this year's program is NSDI's largest to date and that the conference will span three full tracks, an exciting milestone.

Across the two cycles, the Spring deadline received 207 submissions, of which 50 were accepted (24.2%), while the Fall cycle received 452 submissions, with 100 accepted (22.1%). Both cycles included one-shot revisions, with relatively much higher success rates than new submissions: 10 of 13 (77%) in Spring and 8 of 10 (80%) in Fall. These outcomes illustrate the effectiveness of iterative improvement within the review process.

Review Process and Structural Innovations

We made several deliberate changes to the review process this year, informed by lessons from NSDI '25 and inspired in part by Remzi Arpaci-Dusseau's analysis of review and acceptance practices at top systems conferences.

First, we assigned every paper four reviews in Round 1 (R1), up from three in previous years. This provided a stronger signal for early decisions: a paper required three non-positive reviews to be declared an R1 reject, ensuring that every submission received a fair and informed evaluation. Note the emphasis on "three non-positive" reviews: assigning each paper in R1 to four reviewers instead of the usual three was aimed at mitigating the effect of late reviewers (life happens) and, as well, we promoted papers from R1 to R2 if they did not receive enough non-positive reviews.

Second, we adopted a liberal promotion criterion from R1 to Round 2 (R2): any paper receiving at least one weak accept advanced. This inclusive threshold reduced the risk of prematurely rejecting promising work and enabled more thorough consideration of borderline submissions.

Third, we shifted more of the decision-making ahead of the PC meeting. Program Committee members were encouraged to mark papers as #preaccept or #likelyaccept prior to the meeting, reserving live discussion for contested cases. In the Fall cycle, 37 papers were marked #preaccept and 46 as #likelyaccept, allowing the PC meeting to focus on approximately 60 papers, of which 17 were accepted and 15 were recommended for one-shot revision. We deliberated the contentious cases in depth and revisited a few multiple times in the PC meeting and believe that allowing reviewers to talk through their concerns is helpful to steer more submissions towards acceptance. Overall, we accepted roughly half of the papers that were discussed at the PC meetings. We also empowered shepherds to ask for changes and, in almost all cases, we found that the authors were receptive and that the papers improved as a result.

Finally, the Program Committee consisted of 162 members drawn from a broad set of institutions and backgrounds. PC members were asked to complete 11 R1 and 5 R2 reviews in the Fall cycle (16 total) and 6 R1 and 3 R2 reviews in the Spring cycle (9 total), for an overall total of 25 reviews across both cycles. The review load was lower than many recent NSDI cycles, which we believe helped with review quality. We are deeply grateful to every PC member for their dedication and thoughtfulness.

We also benefited from PC assistants and external reviewers, who provided logistical support and domain expertise. Indranil Gupta handled chair-conflict papers in both Spring and Fall cycles. Graduate students from Cornell University (Shuangyu Lei, Salman Abid, Nitika Saran, Yunxi Shen, Keting Chen) attended and assisted with both PC meetings. MIT (Megan Farmer, Mohammad Alizadeh) helped host the in person Spring PC meeting. The Scope Committee, comprising Amar Phanishayee, George Porter, Jon Howell, and James Mickens, played a critical role in evaluating potentially out-of-scope submissions as described below.

Changes to the Scope Determination Process

Maintaining NSDI's scope has been a topic of active discussion within the community and the Steering Committee (SC) for several years. NSDI '25 piloted a scope-enforcement mechanism in which the chairs themselves, working closely with the SC, screened submissions for scope compliance. That approach worked well in Spring but struggled in Fall when submission volume grew substantially.

For NSDI '26, we decentralized the scope decision. Every paper received regular reviews, and PC members were asked to flag potentially out-of-scope submissions using a dedicated checkbox in the review form. Flagged papers were then reviewed by a Scope Committee, who determined whether each paper should be rejected for scope, needed a second scope review, or passed onwards. We found this approach to be more scalable and more consistent: scope decisions were grounded on full technical reviews rather than quick (or desk) judgments, and the chairs did not have to make calls for potentially hundreds of submissions.

In the Spring cycle, 47 of 207 papers (22.7%) were flagged as potentially out of scope. Of these, 10 were accepted, 3 received one-shot revisions, and 34 were rejected, including 13 primarily for scope reasons. In the Fall cycle, 108 of 442 papers (24.4%) were flagged. Of these, 23 were accepted, 2 received one-shot revisions and 83 were rejected, including 17 primarily for scope reasons.

Scope rejections were most often due to lack of a clear networking contribution, including papers focused on compilers, LLM systems, numerical simulation, or unrelated domains. Other categories included sensing/hardware, physical layer work, and computer vision. Notably, several papers that were flagged with scope concerns were discussed in depth and ultimately accepted, underscoring that scope decisions require careful judgment rather than rigid categorization.

Overall, we found this distributed approach effective. The primary challenge was tight turnaround between R1 reviews and scope decisions, which occasionally allowed flagged papers to advance prematurely. We recommend expanding the Scope Committee (e.g., to six members) and aligning timelines more tightly with R1 deadlines.

Program Committee and Community Contributions

The success of NSDI '26 reflects the efforts of a large and dedicated community.

Our Program Committee included 162 members from academia and industry, selected to balance expertise, seniority, and topical coverage. Their thoughtful and constructive reviews are the foundation of this program. We also benefited from additional committees supporting artifact evaluation, awards, posters, mentoring, and scope review. These roles are increasingly essential as the conference grows in scale and complexity.

Artifact Evaluation

NSDI '26 introduced an artifact evaluation process, aligning with broader trends in the systems community. We encouraged authors to submit artifacts, including software, datasets, test suites, proofs, and hardware designs, for evaluation.

We thank the Artifact Evaluation Committee co-chairs, Anny Zheng (Google) and Zhizhen Zhong (Netpreme), for their leadership in establishing this process. We realize that the value of artifact evaluation is being debated in the systems community with concerns that an artifact badge may not mean much (since the artifact can degrade with time and stop being useful) and that artifact evaluation is biased against large-scale production deployments. We concur, to an extent, and are hopeful that continued deliberations may lead to a better mechanism. Nevertheless, we chose to bring back artifact evaluation at NSDI since, even with the above caveats, artifact evaluation is an important step toward improving reproducibility and practical impact.

Awards

We are pleased to recognize three papers with Best/Outstanding Paper Awards and four papers with Community Awards. The Award Committee, comprising Gianni Antichi, Vishal Shrivastav, Andreas Haeberlen, Rodrigo Fonseca, and Noa Zilberman, solicited nominations from the full Program Committee and provided carefully considered recommendations. We are grateful for their diligent work. The Community Award recognizes papers whose code and/or datasets have been made publicly available, encouraging a culture of open and reproducible science within the NSDI community.

Keynote
We are honored to have Amin Vahdat, Chief Technologist and Senior Vice President of AI and Infrastructure at Google, deliver the keynote address. Amin is one of a handful of founders of NSDI and his perspective, looking back on more than two decades of the conference and forward toward the next twenty years, promises to be both historically illuminating and technically inspiring.

Recommendations for the Networked Systems Community and Future Chairs
We offer the following recommendations to the networked systems community and future NSDI chairs:

Scale. Scale is now a defining feature of top-tier systems conferences. Managing this scale requires not only larger committees but also process innovations. Techniques such as pre-acceptances, structured revision pathways, and increased review coverage are essential to maintaining quality.

Submission limits per author. We observed that some authors submitted a large number of papers, effectively increasing their probability of acceptance. This raises broader questions about submission limits and fairness across the community. While NSDI has not yet imposed strict limits, we believe this is an important area for future discussion. For instance, if an author expects a 20% acceptance rate, submitting more papers increases expected acceptances proportionally. This year, several authors submitted nearly 20 papers! Further, their numbers of accepted papers nearly matched the acceptance rate; e.g. submitting 15 papers would expect 3 accepted papers, with 12 papers being rejected. Overall, 14 authors had 8 or more submissions, 156 had 5 or more, and 250 had 4 or more. OSDI, NSDI, and SIGCOMM now impose a limit of 8 submissions per author; we recommend NSDI and the systems community in general, consider even a stricter limit such as 4. Without submission limits, a small number of authors can disproportionately consume significant reviewer bandwidth.

Scope committee process and size. The distributed scope-enforcement model worked well, but a four-member committee felt stretched at peak load. We recommend expanding to six members and establishing clear deadlines for scope reviews that precede R2 assignment.

Conflict declaration and undeclared conflicts. This year we encountered several late-declared conflicts, including advisor-advisee and co-author relationships that were not marked at submission time. Automatic rejection for undeclared conflicts should be enforced consistently. Better tooling such as ORCID integration could help detect co-author conflicts systematically.

Double submission and other violations. We observed a non-trivial number of double submissions and anonymity violations. It is necessary to coordinate with the chairs of related conferences to exchange paper metadata to ensure that paper submissions do not violate submission standards. Standardizing this process will help prevent violators as the community continues to scale.

Operational and experience track bar. We recommend that future CFPs revisit the scope and bar for the operational and experience track papers. The community would benefit from clearer guidance on what constitutes a substantial contribution in this category.

Pre-accept aggressively and early. The strategy of moving papers to #preaccept or #likelyaccept before the PC meeting freed the meeting for genuinely contested papers and improved decision quality. We encourage future chairs to push this approach even further.

Generous review criteria. We recommend maintaining four R1 reviews per paper and a liberal R1—>R2 promotion criterion (any weak accept). This combination provides a strong early signal while ensuring no paper with genuine merit is prematurely eliminated.

One-shot revision. These should complement, not replace, acceptance decisions. We found it effective to first decide to accept/reject and only then consider revision. In particular, we required that a paper be accepted or rejected first, then only after deciding to reject a paper could a one-shot revision be discussed. This process helped to ensure that shepherdable papers were accepted and that one-shot revisions were only used for papers that required significant revision such as new experiments.

Review timeliness. The community does not, in large part, turn reviews in on time. In particular, if we had all four reviews on time, then we advanced a paper to R2 if it had one positive review. But, if it had three reviews at the R1 deadline, then we marked it as an R1 reject if all three were negative. Lastly, if we had papers that had two or less total reviews, then we advanced those to R2, even if the reviews were negative. This was because we had a clear signal to advance to R2 with three or more reviews, but not with two or less. Note that if we marked a paper as an R1 reject because it had three negative reviews and the fourth review came in later as positive, then we double checked the paper with the reviewers. In nearly all cases, it was left as an R1 reject. As a result, late reviews in general caused extra work for the program committee.

Finally, we note that review quality and community culture remain paramount. The success of NSDI depends not only on technical contributions but also on the professionalism, generosity, and integrity of its participants.

Acknowledgments
A program of this scale is possible only because of the extraordinary generosity and hard work of many people. We offer our deepest thanks to the Program Committee members, who gave their expertise and time across both cycles. Special thanks to our Conflict Chair (Indranil Gupta); graduate student assistants; Scope Committee; Award Committee; Artifact Evaluation; Poster (Mina Arashloo, Saksham Agarwal, Grace Liu); and Student Mentoring (Francis Yan) Committee chairs and their respective committees.

We are grateful to the NSDI Steering Committee, particularly Jay Lorch, for thoughtful guidance throughout the year, and to Casey Henderson-Ross and the entire USENIX team for their operational excellence and constant support.

Finally, we thank all authors who submitted to NSDI '26. Every submission represents real intellectual effort, and we hope that even those whose papers were not accepted found value in the reviews they received. NSDI's reputation for constructive, high-quality reviewing is one of the things that makes it a conference worth submitting to. We hope we have honored that tradition.

We hope you enjoy NSDI '26.

Srikanth Kandula, Amazon Web Services
Hakim Weatherspoon, Cornell University
NSDI '26 Program Co-Chairs