TPC® Newsletter

Issue 7, November 2026

Inside: Your Benchmark Briefing

- **Upcoming Tech Talks** – Data quality in the spotlight Dec. 3
- **A New Agentic AI Benchmark** – AMD's proposal for the TPC
- **Tencent Joins the TPC** – SuperSQL at massive scale
- **TPCTC 2026 in Boston** – Eight papers and one invited talk
- **Agentic AI Panel** – From queries to trajectories

Upcoming Tech Talk



The Benchmark Is Never Final: A Living, Demand-Driven Benchmark for Agentic AI with Champion-Relative Evaluation of Context Acquisition

8:30am December 3, 2026

TPC General Council

Dr. Yuqing Zhu of Tsinghua University will give a 20-minute talk by video followed by a 10-minute discussion.

Abstract: Benchmarks have traditionally been written once and frozen, but agentic AI changes faster than any specification, and no ground truth exists for whether an agent found all relevant data. This talk proposes a living benchmark: stable rules, evolving workloads contributed by the demand side, separately tested agent capabilities, and context acquisition scored against a champion record that is re-ranked as data grows. We present the scoring method, its provable properties and its limits.

Bio: Dr. Yuqing Zhu is a researcher at the Beijing National Research Center for Information Science and Technology, Tsinghua University, China. She co-authored BigDataBench, a widely cited big data benchmark suite. She also led the development of IGinX, the system behind AMD's 2025 TPCx-IoT result of 14.3 million IoTps, which ranked first in performance at publication. Her current research focuses on multimodal data management and agentic AI frameworks.

A Proposal for a TPC Benchmark for Agentic AI - by AMD

Agentic systems span nine stages: context preparation, planning, retrieval, reasoning, enterprise tool execution, ephemeral tool execution, verification, response generation, and response delivery. While established benchmarks address many of these stages, ephemeral tool execution remains largely uncovered. These short-lived, task-specific operations include code execution, document processing, compression, cryptographic validation, media processing, testing, and data transformation.

We have developed a simple, reproducible methodology based on recorded-session replay that captures realistic sequences and mixes of ephemeral tool invocations. The initial workload represents five agentic personas: software development, context compaction, security attestation, media processing, and RAG ingestion.

We propose that TPC adopt this as the basis for a Version 1 benchmark, initially focused on CPU-centric ephemeral execution using synthetic but representative workloads, with other stages such as reasoning simulated as needed. The benchmark can evolve to incorporate production-derived workload distributions, accelerator-intensive reasoning, and mixed CPU-GPU execution.

This would establish a vendor-neutral, reproducible standard for measuring the performance, scalability, efficiency, and price/performance of agentic AI infrastructure.

Tencent Joins the TPC Community

Tencent (https://www.tencent.com), a world-leading internet and technology company, joined the Transaction Processing Performance Council (TPC) in September 2026. Tencent's businesses span communication and social services, gaming and digital content, cloud computing, advertising, and FinTech. These businesses make extensive use of SuperSQL, Tencent's analytical platform for interactive analytics and large-scale batch processing. In production, SuperSQL serves more than 100 million interactive queries and runs over 300,000 ETL jobs each day. SuperSQL is also available to external enterprise customers through Tencent Cloud Data Lake Compute (DLC).

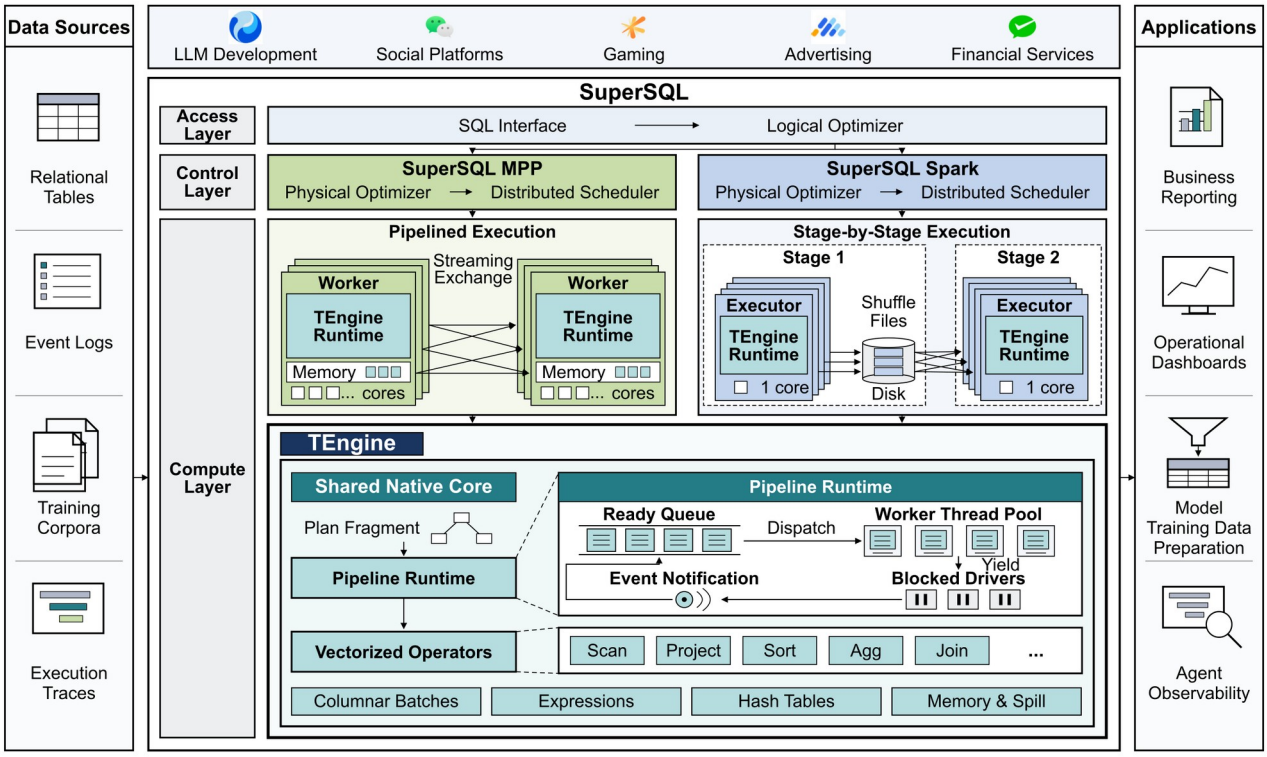


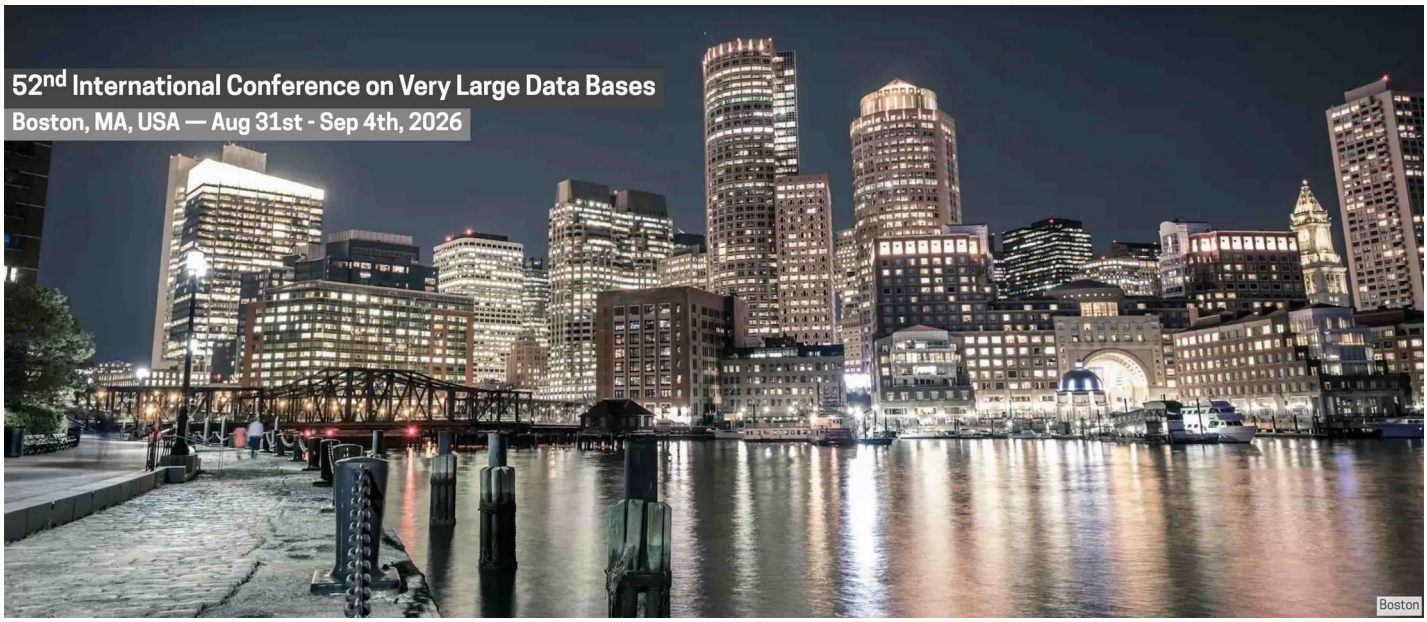
Figure 1. SuperSQL architecture.

At the core of SuperSQL is TEngine, a native vectorized execution engine shared across interactive and batch workloads (Figure 1). Designed to meet Tencent's diverse and rapidly growing analytical demands, SuperSQL brings together **high performance**, **resource efficiency**, and **scalability** in a unified architecture. By combining advanced query optimization, efficient resource management, and distributed execution with vRDMA-accelerated data exchange, SuperSQL delivers fast and reliable analytics across workloads of different sizes and complexities.

The architecture separates query processing into three layers. **(1) Access layer.** A common SQL entry point accepts multiple dialects, uses historical execution statistics to improve plans, and routes queries to SuperSQL MPP for interactive analytics or SuperSQL Spark for batch processing. **(2) Control layer.** Each system uses its own physical optimizer and distributed scheduler to determine how queries are executed and where tasks run. **(3) Compute layer.** TEngine executes the resulting tasks within each worker or executor, providing both systems with shared vectorized operators, columnar data structures, and memory management. Its unified execution architecture enables both interactive queries and large-scale batch processing to benefit from continuous improvements in performance and efficiency, while simplifying system evolution and reducing operational complexity.

Tencent has also evaluated SuperSQL using TPC-DS. Its published TPC-DS Version 4 result for Tencent Cloud Data Lake Compute reports 653,611,141 QphDS@100000GB and a price/performance of CNY 11.04 per kQphDS@100000GB. By joining TPC, Tencent looks forward to sharing lessons from production systems, exchanging ideas with other members, and contributing to transparent, reproducible performance evaluation that helps customers assess analytical systems.

TPC Technology Conference 2026 – TPCTC Boston, United States, August 31st, 2026



TPCTC is TPC's annual technology conference. Its mission is to bring together industry experts and researchers to explore new methodologies for measuring the performance of data-centric applications. Over the last 18 years TPCTC has been recognized as the international event for anyone interested in performance related topics in database technology, including Transaction Processing, Data Warehousing, Big Data Analytics, Internet of Things, Virtualization, Hyper-Converged Infrastructure, and Artificial Intelligence.

Performance evaluation has and still is one of the main differentiators of computer systems. Constant hardware and software improvements require frequent evaluation of existing performance methodologies to allow for them to be technically sound, fair and meaningful. The TPC has always been at the forefront of these developments. While the TPC has focused historically on database centric benchmarks, recent developments include benchmarks for Artificial Intelligence, Internet-Of-Things, Hyper-Converged Infrastructure, Big Data and Virtualization. Many of these benchmarks were sparked by ideas that originated in papers presented at past TPCTC events. These papers were mostly academic papers, that inspired ideas for new benchmarks, identified deficiencies in existing benchmarks and motivated improvements.

At the same time the academic community and industry have benefited from TPCTC to define their own performance methodologies. The last ten years have seen the rise of many new DBMSs, some with unique approaches to traditional problems, like columnar databases, in-memory databases, others specializing on specific applications, such as graph databases, NoSQL databases, Timeseries databases and document databases etc. With them came a flurry of new performance methodologies resulting in customized benchmarks, many of which were based on methodologies originally developed in the TPC.

TPCTC has always served as a venue where both practitioners with real world performance expertise meet with highly innovative academics to discuss performance methodologies for emerging technologies. The results, both in academia and TPC benchmarks, are a direct result of this knowledge exchange.

TPCTC26, chaired by Raghu Nambiar and Meikel Poess, took place in Boston, United States, August 31st, 2026, in conjunction with the 52nd conference on Very Large Databases (VLDB), a top tier academic conference for databases. TPCTC26 marks TPCTC's 18th anniversary.

Starting last year TPCTC has been offering authors the option to submit the tools and scripts used to conduct performance measurements together with their papers. The TPC refers to the tools and scripts as a workload artifact. Accepted workload artifacts will be hosted in TPC's GitHub repository and promoted on the TPC's website together with author names. The TPC will also provide a mechanism for results from workload artifacts to be listed on the TPC's webpage.

This year we received 19 regular paper submissions in a variety of research topics and accepted 8 papers in the following research areas:

General enhancements to TPC workloads (2 Papers):	DataGenX V2: Automated Privacy-Preserving Synthetic Relational Data Generation from Database Statistics and DeepChainBench: A Multi-model Benchmark for Deep Recursive Lineage and Resource Elasticity
GenAI (3 Papers):	The Reasoning Tax: Token Economics of LLM Reasoning Across Task Types and Deployment Contexts, TexBench: Harnessing LLMs for Efficient Key-Value Benchmarking and MLPerf Endpoints: AI Inference Performance Evaluation for the Age of Generative AI.
Internet of Things (1 Paper):	SECBench: A Benchmarking Framework for Stream Processing Systems in the Sensor-Edge-Cloud Continuum.
Database Optimizations (2 Papers):	Intel Linear Tool: Diagnosing Multicore Scaling Bottlenecks Through End-to-End Cache Line Contention Analysis for Microsoft SQL Workloads and BufBench: A Benchmarking Framework for Buffer Pool Analysis in PostgreSQL.

The following are summaries of each of the accepted papers:

SECBench: A Benchmarking Framework for Stream Processing Systems in the Sensor-Edge-Cloud Continuum (Taha Tekdogan, Lukas Schwerdtfeger, Tilmann Rabl, Steffen Zeuch and Volker Markl)

Taha Tekdogan is a PhD student and research associate in the Database Systems and Information Management (DIMA) Group at TU Berlin, as well as a member of the BIFOLD graduate school. Supervised by Prof. Volker Markl, his research focuses on benchmarking, workload characterization, and performance evaluation of distributed stream processing systems (such as NebulaStream) across the sensor-edge-cloud continuum.

Stream processing systems (SPSs) process large volumes of data in real time by distributing computation across multiple machines. Many established SPS deployments focus on centralized cloud environments. However, as more data are generated outside of the cloud, forwarding all the raw data can increase the network load and limit the responsiveness of

applications. Modern SPSs address this limitation by distributing the computation across sensor, edge, and cloud resources, forming the sensor-edge-cloud (SEC) continuum. Although the performance of SPSs has been studied extensively in cloud-only settings, benchmarks provide limited support for evaluating SPSs under the hardware heterogeneity and network diversity of SEC deployments. To address this gap, we propose SECBench, a benchmarking framework for evaluating SPSs in the SEC continuum. SECBench generates workloads, executes benchmark runs, and measures SPS performance while deploying systems in configurable SEC environments that model hardware and network properties. We demonstrate the applicability of SECBench using two real-world use cases with different workload and deployment characteristics. SECBench provides a foundation for studying how SEC configurations affect SPS performance and for comparing cloud-only and SEC processing models in reproducible benchmark environments.

The Reasoning Tax: Token Economics of LLM Reasoning Across Task Types and Deployment Contexts (Sachin Gopal Wani, Ajay Dholakia and David Ellison)

Sachin Gopal Wani is a Staff Data Scientist at Lenovo specializing in end-to-end artificial intelligence and machine learning solutions, with expertise spanning large language models (LLMs), computer vision, and system sizing. A Rutgers University graduate and J.N. Tata Scholar, he frequently publishes on AI infrastructure, token economics, and LLM inference optimization.

Accuracy-only benchmarking of reasoning-capable large language models misses a central deployment question: when do extended thinking tokens earn their cost? We introduce the Token Economy Score (TES), a marginal benchmarking metric that measures the accuracy gain of a reasoning model over a non-reasoning baseline, normalized by the generated-token multiplier. We define paired and approximated TES variants for model families with reasoning toggles and frontier models without direct non-reasoning counterparts. We then conduct an empirical benchmarking analysis across 151 model-benchmark evaluation runs on seven benchmarks spanning mathematics, code generation, science reasoning, instruction following, expert knowledge, knowledge recall, and research-level physics. The analysis examines three deployment-facing dimensions: which task structures yield positive marginal reasoning efficiency, how increasing reasoning effort changes TES within model families, and how deployment context changes economic viability. Results show that task structure predicts reasoning efficiency better than nominal difficulty: sequential inference-chain tasks such as AIME 2025 and LiveCodeBench show high TES, while knowledge-recall tasks such as MMLU-Pro show low TES despite their difficulty. We also find systematic diminishing returns at higher reasoning effort levels, including cases where additional thinking reduces accuracy. Finally, Reasoning Cost Share (RCS) shows that inference spend is often dominated by internal thinking, while Deployment Cost Multiplier (DCM) shows how on-premises deployment can change the economics of otherwise costly reasoning workloads. These findings support a benchmarking-driven model-selection rule: enable reasoning selectively by task type, effort level, and deployment context rather than treating it as a universally beneficial mode.

DataGenX V2: Automated Privacy-Preserving Synthetic Relational Data Generation from Database Statistics (Ahmad Ghazal, Sree Harsha Ramanavarapu, Hanumath Maduri, Qizhi Wang, Ankit Kapoor, Xianfei Deng and Yu Dong)

Ahmad Ghazal is a database systems architect and researcher at SingleStore, with prior engineering leadership roles at PingCAP and Teradata. A prolific contributor to the database community and TPC benchmarking efforts, his work centers on query optimization, HTAP architectures, synthetic data generation, and vector database systems.

Synthetic relational data is essential for database testing, benchmarking, and optimizer debugging when production rows cannot be shared. Optimizer statistics—row counts, distinct-value counts, histograms, and top-N frequencies—are collected by databases to guide plan selection; they also constitute a compact and precise characterization of the data's distribution, making them a natural specification for synthetic data generation. DataGenX V1 introduced annotated SQL

DDL for generating synthetic relational data from schema-level specifications and statistics, but required users to manually inspect schemas, encode annotations, and tune generation rules. We present DataGenX V2, an end-to-end automation framework that turns relational catalog metadata and optimizer-visible statistics into executable DataGenX templates, generating annotated DDL and synthetic data without copying source rows.

DataGenX V2 is implemented as a metadata extraction and template-generation pipeline. It builds a foreign-key dependency graph, topologically orders table generation, and selects column-specific expression strategies. Database-specific adapters for MySQL, TiDB, and SingleStore handle differences in catalog metadata, histogram access, and statistics collection. We evaluate DataGenX V2 on TPC-H and TPC-DS across MySQL, TiDB, and SingleStore, validating row- and column-level statistics, FK referential integrity, and privacy preservation. Generated databases preserve schema structure, table cardinalities, distinct-value properties, referential relationships, and optimizer-relevant column statistics, while containing no copied source data values.

Invited Talk: Benchmarking Agentic AI
Raghunath Nambiar

Raghunath (Raghu) Nambiar is a Corporate Vice President and Chief Technology Officer at AMD, leading datacenter ecosystem architecture, application engineering, and AI/HPC solutions. A veteran technologist and long-time leader within the Transaction Processing Performance Council (TPC), his work focuses on systems architecture, big data analytics, and performance benchmarking standards.

Intel Linear Tool: Diagnosing Multicore Scaling Bottlenecks Through End-to-End Cache Line Contention Analysis for Microsoft SQL Workloads (Swathi Kovvuri, Janardhana Yoga Narasimhaswamy, Thierry Fevrier and Sharanyan Srikanthan)

Modern server platforms increasingly adopt disaggregated, multi-die architectures with growing core densities, widening coherence domains and amplifying cache-line ownership costs across NUMA and die boundaries. These effects create scalability blind spots that software diagnostics and hardware profilers cannot bridge independently. We present Linear, a cache-line-centric analysis tool that combines hardware sampled latency, data source attribution, instruction context, and symbol metadata into a unified per-address contention profile. Linear ranks the cache lines along with corresponding data structures and fields whose sharing patterns dominate stall time, exposing context unavailable from database wait statistics or conventional hotspot profiles. Using OLTP and OLAP benchmarks on production class Intel multi-socket Xeon® systems running Microsoft® SQL Server® presented two case studies where Linear identified layout-driven page cache metadata sharing and columnar store cache population synchronization as dominant scaling bottlenecks. These diagnoses guided targeted mitigations that reduced coherence overhead, closed cold/warm execution gaps, and improved throughput scaling.

BufBench: A Benchmarking Framework for Buffer Pool Analysis in PostgreSQL (Aroma Hoque and Tarikul Islam Papon)

Aroma Hoque is a computer science researcher and graduate student at the University of Massachusetts Boston, working under Dr. Tarikul Islam Papon. Her research focuses on database systems, page replacement policies, and database performance, notably co-authoring *BufBench*, a benchmarking framework for PostgreSQL buffer pool analysis presented at TPCTC.

The buffer pool manages all data movement between main memory and disk, making it one of the most central components of any database management system. Yet popular benchmarking tools like pgbench and BenchBase offer no visibility into buffer pool internals such as page eviction, buffer utilization, or background process interference. In this paper, we present BufBench, a purpose-built, open-source benchmarking framework that provides systematic analysis of PostgreSQL buffer pool behavior and its impact on overall performance. BufBench makes four core contributions. First, it provides a unified suite of six standard benchmarks: TPC-B, TPC-C, TPC-H, YCSB, CH-benCHmark, and TATP, covering a broad spectrum of OLTP, OLAP, skewed-access, and mixed workload characteristics. Second, it supports a pluggable page replacement policy interface, enabling direct comparison of PostgreSQL's default Clock-Sweep against alternative policies (currently including LRU, CFLRU, and LRU-WSR). Third, it incorporates a fine-grained buffer pool monitoring layer that captures 34 metrics per second from PostgreSQL's internal statistics views. Fourth, it automatically detects and quantifies the interference caused by checkpointing and vacuum events on throughput, a capability absent from all existing tools. Together, these capabilities make BufBench a ready-to-use diagnostic tool that database researchers and practitioners can deploy to analyze the whole design space of buffer pool behaviors.

TexBench: Harnessing LLMs for Efficient Key-Value Benchmarking (Shubham Kaushik, Abhishek Chanda and Subhadeep Sarkar)

Shubham Kaushik is a PhD candidate in computer science at Brandeis University, advised by Prof. Subhadeep Sarkar. His research centers on data systems, storage engine internals, and benchmarking, with a particular focus on LSM-tree optimizations, buffer designs, and workload generation tools such as Tectonic and TexBench.

Key-value stores are widely adopted to serve modern heterogeneous workloads, yet selecting the best key-value store for a given workload and target performance remains an incredibly complex and time-consuming task, even for expert DBAs. In practice, this entails extensive, large-scale benchmarking and analysis involving dozens of candidate key-value stores. Comparing databases using existing key-value benchmarks, such as YCSB, KVBench, db_bench, and Tectonic, has several challenges. (A) It requires a substantial effort in setting up the benchmarking environment and connecting the databases, along with a thorough understanding of benchmark-specific configurations and tuning parameters. Further, existing benchmarks either (B) are completely unable to capture the dynamic and complex nature of real-world workloads, or (C) rely on fixed, templated benchmark workloads, with limited customization options. Thus, benchmarking key-value stores against application-specific, production-like workloads is difficult and does not scale, leading to suboptimal performance. We introduce TexBench, a unified key-value benchmarking suite that abstracts the complexities of workload configuration, benchmark setup, and database connections behind a unified interface, enabling users to evaluate multiple key-value stores against production-like workloads, and visually compare their performance side-by-side. To emulate the dynamic and complex properties of real-world workloads, TexBench uses Tectonic – a Rust-based, highly configurable key-value benchmark – under the hood. To facilitate benchmarking against custom workloads, TexBench integrates a large language model (LLM) plugin that captures benchmarking specifications expressed in natural language and translates them into a JSON format interpretable by Tectonic. TexBench currently supports four key-value stores – ScyllaDB, RocksDB, Cassandra, and Redis, and performs up to 16.7× faster in terms of workload generation compared to raw LLM prompting. The source code of TexBench is available at https://github.com/SSD-Brandeis/TexBench.

MLPerf Endpoints: AI Inference Performance Evaluation for the Age of Generative AI (Miro Hodak and Meena Arunachalam)

Miro (Miroslav) Hodak is an AI systems architect and researcher—formerly an AI Architect at Lenovo Infrastructure Solutions Group and now a Senior Member of Technical Staff at AMD. A former computational physicist, his work centers on AI benchmarking, edge inference architectures, and performance optimization across industry standards such as

MLPerf and TPCx-AI.

The rapid adoption of generative AI has exposed a gap between established AI inference benchmarks and the metrics needed to plan generative AI deployments. Traditional benchmarks such as MLPerf Inference produce single-number performance results that do not capture the complex trade-offs inherent in serving large language models to multiple users. This paper presents MLPerf Endpoints, a new benchmark developed by MLCommons that combines serving native metrics—throughput, interactivity, and latency evaluated as functions of concurrency—with real datasets, strict accuracy requirements, and community governance. Its key design choices, including a decoupled server/client architecture, Pareto-curve-based results display, and rolling submissions, are designed to provide insights needed for generative AI performance evaluation and deployment planning. We also present a comparative analysis of MLPerf Endpoints alongside MLPerf Inference and InferenceX, examining how the three suites differ in metrics, governance, and design philosophy. Our analysis shows that MLPerf Endpoints bridges the gap between raw performance evaluation and practical deployment needs, bringing AI inference benchmarking into the generative AI era.

DeepChainBench: A Multi-model Benchmark for Deep Recursive Lineage and Resource Elasticity
(Jorge Martins, Nelson Tenório and Jorge Bernardino)

Jorge Martins is a computer science researcher and engineer associated with the Polytechnic University of Coimbra (ISEC). His work focuses on data engineering, multi-model data management, and systems evaluation, contributing to benchmark developments such as DeepChainBench for recursive data lineage and elasticity.

The emerging paradigm of Multi-model Database Management Systems promises to combine the flexibility of NoSQL ecosystems with the rigor of relational systems. However, contemporary benchmarking tools are unable to evaluate the behavior of these platforms under conditions of extreme structural depth and hardware elasticity. Many of the current solutions are limited to evaluating horizontal data varieties, resorting to shallow graph queries that hide the execution engines' physical limits. This paper proposes DeepChainBench, a conceptual benchmarking framework and methodological specification that evaluates deep recursive lineages and infrastructure resilience in industrial supply chain scenarios. DeepChainBench models a synthetic dataset that forces a Semantic Collision between graph, semi-structured document, and key-value models. A deterministic algorithmic generator circumvents the small-world phenomenon by introducing a 5-Zone Hybrid Topology, featuring extreme linear structures and memory diamonds. The goal is to stress execution stacks and quantify their impact on essential hardware resources. To this end, we introduce the Resource Footprint Index (RFI) alongside tail latency metrics.

Get the TPCTC Proceedings

Want the TPCTC papers on your bookshelf? The proceedings of the 17th TPC Technology Conference (TPCTC 2025) are published by Springer as *Performance Evaluation and Benchmarking*, Lecture Notes in Computer Science volume 16261 (ISBN 978-3-032-18069-8). You can order the softcover from Springer or major booksellers, or download the eBook or individual chapters from SpringerLink. If your company or university subscribes to Springer, your library may be able to get you the papers at no cost. Volumes from earlier years are on the TPCTC series page.

Panel: From Queries to Trajectories: Benchmarking Closed-Loop Agentic AI Systems

AI systems are rapidly evolving from single-turn generative models into closed-loop agents that plan, invoke tools, coordinate with other agents, maintain state, and operate within simulated or learned environments. This shift challenges conventional benchmarking methods that focus on isolated prompts, model accuracy, latency, or throughput. For agentic systems, the natural unit of evaluation is no longer a query, but a trajectory: the end-to-end execution path of a task, including model calls, tool use, memory updates, inter-agent communication, retries, artifacts, costs, and final outcomes. This panel will examine how performance evaluation must evolve for such systems, with emphasis on workload design, trace-based measurement, token economics, run-to-run variance, reliability, reproducibility, and energy-aware cost-performance. The discussion will also consider the emerging role of world models, both as benchmark targets and as controlled environments for evaluating autonomous agents. Panelists from academia and industry will debate what metrics, reporting standards, and audit mechanisms are needed for fair comparison of closed-loop AI systems. The goal is to identify practical directions toward a TPC-style benchmarking framework, including the possibility of a standardized agent transaction log for reproducible evaluation.



Panelist: Dr. Ajay Dholakia (moderator, Lenovo)

Dr. Ajay Dholakia is Principal Engineer, AI Leader and Chief Technologist for Software & Solutions Development in Lenovo Infrastructure Solutions Group. In this role, he is leading the development of customer solutions in the areas of AI / ML, Generative AI, AIOps, Big Data & Analytics, Cloud Computing, Edge Computing and Blockchain. Most recently, he is driving new projects using emerging AI technologies such as Agentic AI systems, RAG-assisted Large Language Models and domain-specific Small Language Models. In his career spanning over 30 years, he has led diverse projects in research, technology, product and solution development as well as business/technical strategy. Ajay holds more than 70 patents and has authored over 70 technical publications including a book.



Sachin Wani (Lenovo)

Sachin Gopal Wani is a Staff Data Scientist at Lenovo specializing in end-to-end Artificial Intelligence solutions, from high-throughput LLM inference to real-time edge computer vision projects. A Rutgers University gold medalist and J.N. Tata Scholar, he actively works on researching latest in GenAI and publishes whitepapers and conference papers that impact practitioners. Sachin is passionate about translating complex AI research into scalable, practical machine learning applications that solve real-world challenges across diverse industries.



David Ellison (Lenovo)

David Ellison is the Chief Data Scientist & Director of AI & HPC Engineering for Lenovo ISG. Through Lenovo's US and European AI Discover Centers, he leads a team that uses cutting-edge AI techniques to deliver solutions for external customers while internally supporting the overall AI strategy for the World Wide Infrastructure Solutions Group. Before joining Lenovo, he ran an international scientific analysis and equipment company and worked as a Data Scientist for the US Postal Service. Previous to that, he received a PhD in Biomedical Engineering from Johns Hopkins University. He has numerous publications in top tier journals including two in the Proceedings of the National Academy of Sciences.



Miro Hodak (AMD, MLPerf)

Miro Hodak is the Principal Member of Technical Staff in AI Performance Engineering at AMD and serves as the MLPerf Inference Workgroup Co-Chair at MLCommons. He spearheads AMD's AI software optimization and leads the engineering teams driving AMD's MLPerf datacenter benchmark performance and submissions.



Debo Dutta (Nutanix)

Debo works for Nutanix and is also active in MLPerf. As Chief AI Officer of Nutanix he defines Nutanix's enterprise AI strategy and leads the engineering teams behind platforms like Nutanix Enterprise AI and GPT-in-a-Box, enabling organizations to run large language models and agentic applications on private, hybrid-cloud infrastructure. MLCommons & MLPerf: A founding member of MLCommons, he serves as a board advisor and helped co-incubate MedPerf, an open-source federated benchmarking framework designed to validate medical AI models securely across institutions without exposing patient data.



Tim Kraska (MIT)

Tim Kraska is an Associate Professor of Electrical Engineering and Computer Science at MIT, where he is part of the Data Systems Group in the Computer Science and Artificial Intelligence Laboratory (CSAIL). He is also a Director of Applied Science at Amazon Web Services (AWS) and co-directs MIT's Generative AI Impact Consortium (MGAIC) and the Data Systems and AI Lab (DSAIL@CSAIL).

TPCTC 2026 Organization

General Chairs and Contacts

Raghunath Nambiar, AMD, USA, raghu.nambiar@amd.com
Meikel Poess, Oracle, USA, meikel.poess@oracle.com

Publication Chair

Rodrigo D. Escobar, Univ. Texas at San Antonio, USA

Program Committee

Andrew Bond, Red Hat, USA
Michael Brey, Oracle, USA
Paul Cao, HPE, USA
Ajay Dholakia, Lenovo, USA
Rodrigo Escobar, Univ. Texas at San Antonio, USA
Ahmad Ghazal, SingleStore, USA
Shahram Ghandeharizadeh, University of Southern California, USA
Miro Hodak, AMD, USA
Swathi Kovvuri, Intel Corporation, USA
Klaus-Dieter Lange, Hewlett Packard Enterprise, USA
Hanumath Maduri, Workday, USA
Raghunath Nambiar, AMD, USA
Meikel Poess, Oracle Corporation, USA
Tilmann Rabl, Hasso Plattner Institute, Germany
Steve Shaw, HammerDB, United Kingdom
Pengcheng Zheng, Timecho Europe GmbH, Germany