Zestawienie konferencji USENIX: https://www.usenix.org/conferences/past

Dla każdej konferencji wskazuję przykładowe prace, co nie oznacza, że należy się ograniczać tylko do tych wskazanych.

2025

1. The 2026 Linux Storage, Filesystem, Memory-Management, and BPF Summit, **May 2026**, Zagreb. https://events.linuxfoundation.org/lsfmmbpf/

2. FAST'26. The 24th USENIX Conference on File and Storage Technologies will take place on **February 24–26, 2026**
   https://www.usenix.org/conference/fast26

3. Systor 2025, The ACM International Systems and Storage Conference. **Sept 2025**.
   https://www.systor.org/2025/

4. The 2025 Linux Storage, Filesystem, Memory-Management, and BPF Summit, March 2025, Montreal.
   https://lwn.net/Articles/lsfmmbpf2025/
   Mnóstwo ciekawych prezentacji.

5. Reports from Power Management and Scheduling in the Linux Kernel (OSPM), March 2025
   a. Day 1,
   b. Day 2,
   c. Day 3.

6. FAST'25. 23rd USENIX Conference on File and Storage Technologies, February 2025.
   https://www.usenix.org/conf%C3%A9rence/fast25/technical-sessions
   - Boosting File Systems Elegantly: A Transparent NVM Write-ahead Log for Disk File Systems, https://www.usenix.org/system/files/fast25-wang.pdf
   - Fast, Transparent Filesystem Microkernel Recovery with Ananke, https://www.usenix.org/conference/fast25/presentation/liu-jing – Best paper
   - DJFS : Directory-Granularity Filesystem Journaling for CMM-H SSDs, https://www.usenix.org/system/files/fast25-yoo.pdf
   - ScaleLFS: A Log-Structured File System with Scalable Garbage Collection for Commodity SSDs, https://www.usenix.org/conference/fast25/presentation/ha
   - Rethinking the Request-to-IO Transformation Process of File Systems for Full Utilization of High-Bandwidth SSDs, https://www.usenix.org/conference/fast25/presentation/zhan
   - Cloudscape: A Study of Storage Services in Modern Cloud Architectures, https://www.usenix.org/conference/fast25/presentation/satija
   - Mooncake: Trading More Storage for Less Computation — A KVCache-centric Architecture for Serving LLM Chatbot, https://www.usenix.org/conference/fast25/presentation/qin, Best paper
   - ShiftLock: Mitigate One-sided RDMA Lock Contention via Handover, https://www.usenix.org/conference/fast25/presentation/gao
   - Liquid-State Drive: A Case for DNA Block Device for Enormous Data, https://www.usenix.org/conference/fast25/presentation/zhou-jiahao

7. HotStorage 2025, The 17th ACM Workshop on Hot Topics in Storage and File Systems, July 2025, https://www.hotstorage.org/2025/program.html
    - SnapBPF: Exploiting eBPF for Serverless Snapshot Prefetching (best paper)
    - Can a Client–Server Cache Tango Accelerate Disaggregated Storage?
    - Why Paying for Storage Beats Free Networking in Cloud Bursting
    - Revisiting Memory Hierarchies with CMM-H: Using Device-side Caching to Integrate DRAM and SSD for a Hybrid CXL Memory
    - ByteExpress: A High-Performance and Traffic-Efficient Inline Transfer of Small Payloads over NVMe - Best paper
    - PaperCache: In-Memory Caching with Dynamic Eviction Policies

8. USENIX ATC'25. 2025 USENIX Annual Technical Conference, July 2025 https://www.usenix.org/conference/atc25/technical-sessions
    - PageFlex: Flexible and Efficient User-space Delegation of Linux Paging Policies with eBPF, https://www.usenix.org/conference/atc25/presentation/yelam

9. OSDI'25. 19th USENIX Symposium on Operating systems Design and Implementation. July 2025. https://www.usenix.org/conf%C3%A9rence/osdi25/technical-sessions
    - FineMem: Breaking the Allocation Overhead vs. Memory Waste Dilemma in Fine-Grained Disaggregated Memory Management, https://www.usenix.org/conference/osdi25/presentation/wang-xiaoyang
    - Enabling Efficient GPU Communication over Multiple NICs with FuseLink, https://www.usenix.org/conference/osdi25/presentation/ren
    - Skybridge: Bounded Staleness for Distributed Caches, https://www.usenix.org/conference/osdi25/presentation/lyerly
    - Tintin: A Unified Hardware Performance Profiling Infrastructure to Uncover and Manage Uncertainty, https://www.usenix.org/conference/osdi25/presentation/li
    - Building Bridges: Safe Interactions with Foreign Languages through Omniglot, https://www.usenix.org/conference/osdi25/presentation/schuermann
    - XSched: Preemptive Scheduling for Diverse XPUs, https://www.usenix.org/conference/osdi25/presentation/shen-weihang

10. NSDI'25. 22nd USENIX Symposium on Networked Systems Design and Implementation, April 2025. https://www.usenix.org/conf%C3%A9rence/nsdi25/technical-sessions
    - Beehive: A Scalable Disaggregated Memory Runtime Exploiting Asynchrony of Multithreaded Programs, https://www.usenix.org/conference/nsdi25/presentation/li-quanxi
    - Making Serverless Pay-For-Use a Reality with Leopard, https://www.usenix.org/conference/nsdi25/presentation/cao
    - GRANNY: Granular Management of Compute-Intensive Applications in the Cloud, https://www.usenix.org/conference/nsdi25/presentation/segarra
    - One-Size-Fits-None: Understanding and Enhancing Slow-Fault Tolerance in Modern Distributed Systems, https://www.usenix.org/conference/nsdi25/presentation/lu
    - White-Boxing RDMA with Packet-Granular Software Control, https://www.usenix.org/conference/nsdi25/presentation/zhao-chenxingyu

- ONCache: A Cache-Based Low-Overhead Container Overlay Network, https://www.usenix.org/conference/nsdi25/presentation/lin-shengkai
- The Benefits and Limitations of User Interrupts for Preemptive Userspace Scheduling, https://www.usenix.org/conference/nsdi25/presentation/guo
- Mitigating Scalability Walls of RDMA-based Container Networks, https://www.usenix.org/conference/nsdi25/presentation/liu-wei
- Eden: Developer-Friendly Application-Integrated Far Memory, https://www.usenix.org/conference/nsdi25/presentation/yelam
- ODRP: On-Demand Remote Paging with Programmable RDMA, https://www.usenix.org/conference/nsdi25/presentation/wang-zixuan
- Understanding and Profiling NVMe-over-TCP Using ntprof, https://www.usenix.org/conference/nsdi25/presentation/kang

11. Inne

2024

1. The 2024 Linux Storage, Filesystem, Memory-Management, and BPF Summit, May 13
   https://lwn.net/Articles/973683/
   https://www.youtube.com/playlist?list=PLbzoR-pLrL6oj1rVTXLnV7cOuetvjKn9q
2. Systor 2024, The ACM International Systems and Storage Conference. Sept 2024.
   Accepted papers:
   https://www.systor.org/2024/accepted-papers/#1630482835121-31ef82a8-bd1f
   Na przykład:
   - ~~TwinPilots: A New Computing Paradigm for GPU-CPU Parallel LLM Inference (best paper)~~
     ~~https://dl.acm.org/doi/10.1145/3688351.3689164~~
   - Performance Characterization of SmartNIC NVMe-over-Fabrics Target Offloading
     https://dl.acm.org/doi/10.1145/3688351.3689154
   - Offloading Datacenter Jobs to RISC-V Hardware for Improved Performance and Power Efficiency
     https://dl.acm.org/doi/10.1145/3688351.3689152
   - Can OS Specialization give new life to old carbon in the cloud?
     https://dl.acm.org/doi/10.1145/3688351.3689158 aaa
3. OSDI 2024, 18th USENIX Symposium on Operating Systems Design and Implementation
   List of papers: https://www.usenix.org/conference/osdi24/technical-sessions
   Full proceedings: https://www.usenix.org/system/files/osdi24_full_proceedings.pdf
   Na przykład:
   - Nomad: Non-Exclusive Memory Tiering via Transactional Page Migration
     https://www.usenix.org/conference/osdi24/presentation/xiang
   - ~~Harvesting Memory-bound CPU Stall Cycles in Software with MSH~~
     ~~https://www.usenix.org/conference/osdi24/presentation/luo~~
   - ~~DRust: Language-Guided Distributed Shared Memory with Fine Granularity, Full Transparency, and Ultra Efficiency~~
     ~~https://www.usenix.org/conference/osdi24/presentation/ma-haoran~~
   - Llumnix: Dynamic Scheduling for Large Language Model Serving

https://www.usenix.org/conference/osdi24/presentation/sun-biao aa

4. USENIX ATC 2024, 2024 USENIX Annual Technical Conference
   List of papers: https://www.usenix.org/conference/atc24/technical-sessions
   Na przykład:
   - ~~Limitations and Opportunities of Modern Hardware Isolation Mechanisms~~
     ~~https://www.usenix.org/conference/atc24/presentation/chen-xiangdong~~
   - ExtMem: Enabling Application-Aware Virtual Memory Management for Data-Intensive Applications
     https://www.usenix.org/conference/atc24/presentation/jalalian
   - ~~PeRF: Preemption-enabled RDMA Framework~~
     ~~https://www.usenix.org/conference/atc24/presentation/lee~~
   - ~~An Empirical Study of Rust-for-Linux: The Success, Dissatisfaction, and Compromise~~
     ~~https://www.usenix.org/conference/atc24/presentation/li-hongyu~~
   - Scalable and Effective Page-table and TLB management on NUMA Systems
     https://www.usenix.org/conference/atc24/presentation/gao-bin-scalable
   - SuperBench: Improving Cloud AI Infrastructure Reliability with Proactive Validation
     https://www.usenix.org/conference/atc24/presentation/xiong
   - StreamCache: Revisiting Page Cache for File Scanning on Fast Storage Devices
     https://www.usenix.org/conference/atc24/presentation/li-zhiyue
   - ScalaCache: Scalable User-Space Page Cache Management with Software-Hardware Coordination
     https://www.usenix.org/conference/atc24/presentation/peng
   - Every Mapping Counts in Large Amounts: Folio Accounting
     https://www.usenix.org/conference/atc24/presentation/hildenbrand

5. HotStorage 2024, The 16th ACM Workshop on Hot Topics in Storage and File Systems
   https://www.hotstorage.org/2024/
   Accepted papers: https://www.hotstorage.org/2024/accepted.html
   Na przykład:
   - Dictionary Based Cache Line Compression
     https://dl.acm.org/doi/10.1145/3655038.3665941
   - Improving Virtualized I/O Performance by Expanding the Polled I/O Path of Linux
     https://dl.acm.org/doi/10.1145/3655038.3665944
   - Secure Archival is Hard... Really Hard
     https://dl.acm.org/doi/10.1145/3655038.3666093
   - Can Storage Devices be Power Adaptive?
     https://dl.acm.org/doi/10.1145/3655038.3665945
   - Can ZNS SSDs be Better Storage Devices for Persistent Cache?
     https://dl.acm.org/doi/10.1145/3655038.3665946
   - Quantitative Analysis of Storage Requirement for Autonomous Vehicles
     https://dl.acm.org/doi/10.1145/3655038.3665948
   - Asymmetric RAID: Rethinking RAID for SSD Heterogeneity
     https://dl.acm.org/doi/10.1145/3655038.3665952
   - Breaking Barriers: Expanding GPU Memory with Sub-Two Digit Nanosecond Latency CXL Controller
     https://dl.acm.org/doi/10.1145/3655038.3665953 zz

6. FAST 2024, 22nd USENIX Conference on File and Storage Technologies
   https://www.usenix.org/conference/fast24
   Full proceeedings: https://www.usenix.org/system/files/fast24-full_proceedings.pdf
   Na przykład:
   - Lessons Learnt in Trying to Build New Storage Technologies
     https://www.usenix.org/conference/fast24/presentation/rowstron
   - ~~TeRM: Extending RDMA-Attached Memory with SSD~~
     ~~https://www.usenix.org/conference/fast24/presentation/yang-zhe~~
   - ~~Combining Buffered I/O and Direct I/O in Distributed File Systems~~
     ~~https://www.usenix.org/conference/fast24/presentation/qian~~
   - Optimizing File Systems on Heterogeneous Memory by Integrating DRAM Cache with Virtual Memory Management
     https://www.usenix.org/conference/fast24/presentation/liu-yubo
   - I/O Passthru: Upstreaming a flexible and efficient I/O Path in Linux
     https://www.usenix.org/conference/fast24/presentation/joshi
   - Storage Systems in the LLM Era
     https://www.usenix.org/conference/fast24/presentation/panel-storage-systemsaaa
7. NSDI 2024, 21st USENIX Symposium on Networked Systems Design and Implementation
   Full proceedings: https://www.usenix.org/sites/default/files/nsdi24-full_proceedings.pdf
   List of papers: https://www.usenix.org/conference/nsdi24/technical-sessions
   Na przykład:
   - ~~Horus: Granular In-Network Task Scheduler for Cloud Datacenters~~
     ~~https://www.usenix.org/conference/nsdi24/presentation/yassini~~
   - DINT: Fast In-Kernel Distributed Transactions with eBPF
     https://www.usenix.org/conference/nsdi24/presentation/zhou-yang
   - ~~Netcastle: Network Infrastructure Testing At Scale~~
     ~~https://www.usenix.org/conference/nsdi24/presentation/sherwood~~
   - SIEVE is Simpler than LRU: an Efficient Turn-Key Eviction Algorithm for Web Caches
     https://www.usenix.org/conference/nsdi24/presentation/zhang-yazhuo
   - ~~Load is not what you should balance: Introducing Prequal~~
     ~~https://www.usenix.org/conference/nsdi24/presentation/wydrowski~~
   - ~~MobileConfig: Remote Configuration Management for Mobile Apps at Hyperscale~~
     ~~https://www.usenix.org/conference/nsdi24/presentation/guo~~

Różne

- Memory deduplication (od Marcina Copika z ETH Zurich)
  Artykuł: https://mcopik.github.io/assets/pdf/2023_bigdata_upm.pdf
  Slajdy: https://mcopik.github.io/assets/pdf/2023_bigdata_upm_slides.pdf

- ‘Mind-blowing’ IBM chip speeds up AI, https://www.nature.com/articles/d41586-023-03267-0?utm_source=Live+Audience&utm_campaign=2437e496f0-briefing-dy-20231020&utm_medium=email&utm_term=0_b27a691814-2437e496f0-51328812
(trzeba poszukać w sieci jakiegoś sensownego artykułu i najlepiej zrobić przegląd najnowszych architektur wspomagających AI)

- Flame graphs:
https://www.brendangregg.com/flamegraphs.html,

https://queue.acm.org/detail.cfm?id=2927301
https://queue.acm.org/downloads/2016/Gregg4.svg
https://www.youtube.com/watch?v=VMpTU15rIZY
https://www.youtube.com/watch?v=HKQR7wVapgk

Storag trends for 2025, https://www.youtube.com/watch?app=desktop&v=lKBWYKRNJ-4
Storage trends for 2024, https://www.youtube.com/watch?v=l8qPrFtY6qE
Storage trends for 2023 and beyond, https://www.youtube.com/watch?v=QmmVkz5zhgg

---

Od Andrzeja Jackowskiego

- The Fungible DPU A New Category of Microprocessor

What is New in the Linux kernel
https://lwn.net/Kernel/
https://lwn.net/Kernel/Index/ i wyszukać po ostatnim roku (2025) – przykłady:

- Taking BPF programs beypond one-million instructions
- Improving control over transparent huge page use
- Per-CPU memory for user space
- The state of the page in 2025
- Improving the merging of anonymous VMAs
- Three ways to rework the swap subsystem
- Cache awarness for the CPU scheduler
- Scheduler medley: time-slice extension, sched_ext deadline servers, and LRU batching

Prace z Research at Google, Distributed Systems and Parallel Computing
https://research.google/research-areas/distributed-systems-and-parallel-computing/
Zdecydowana większość prac bardzo dobrze pasuje do tematyki seminarium, oczywiście lepsze są te najnowsze, z ostatniego roku.
Pełna lista publikacji: https://research.google/pubs/?area=distributed-systems-and-parallel-computing