



nKV in Action

Accelerating KV-Stores on Native Computational Storage with Near-Data Processing



Tobias Vinçon, Arthur Bernhardt, Ilia Petrov
Data Management Lab
Reutlingen University, Germany


Lukas Weber, Andreas Koch
Embedded Systems and Applications Group
TU Darmstadt, Germany



Demo Setup



COSMOS+



PCIe

Host

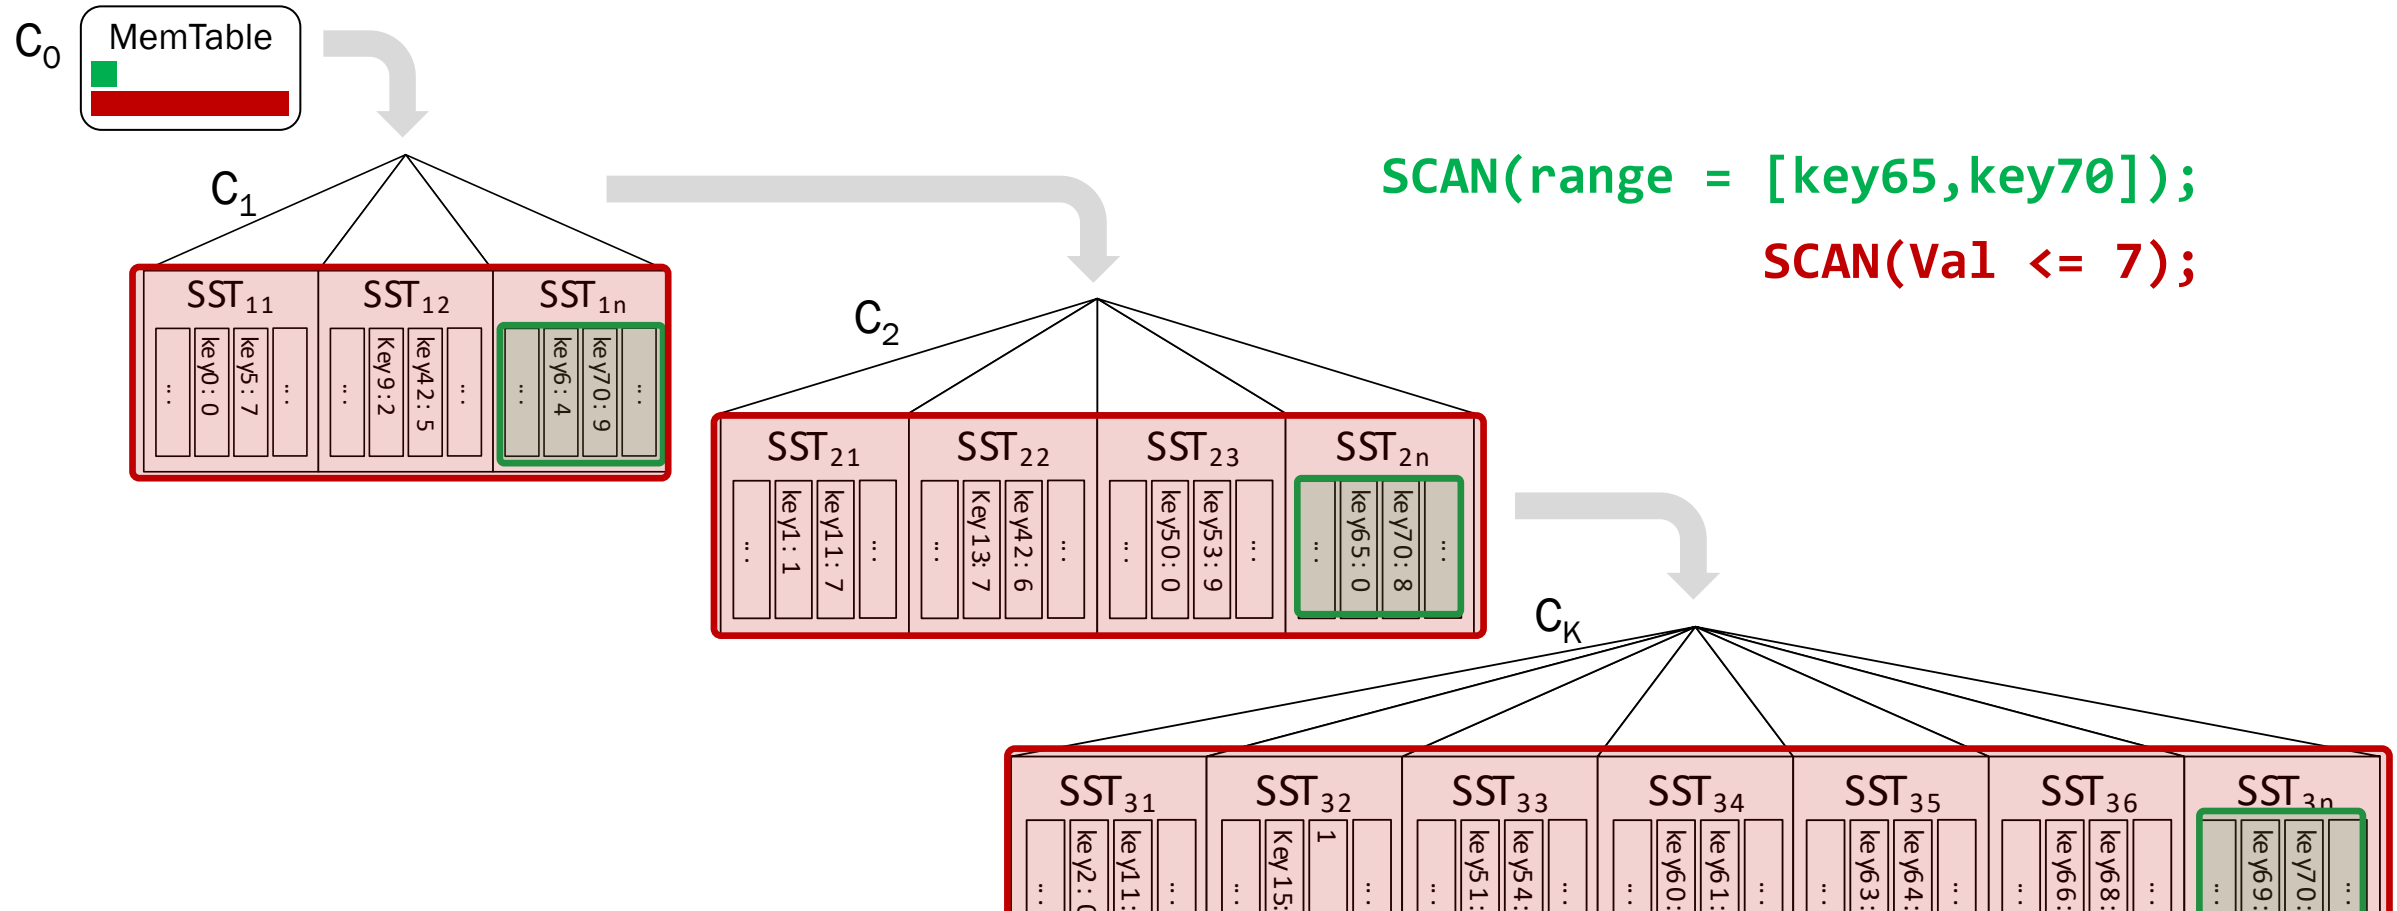


Switch to nKV VLDB Demo GUI

1. Execute BC (Duration 2-4 min): Best Paper Award Check

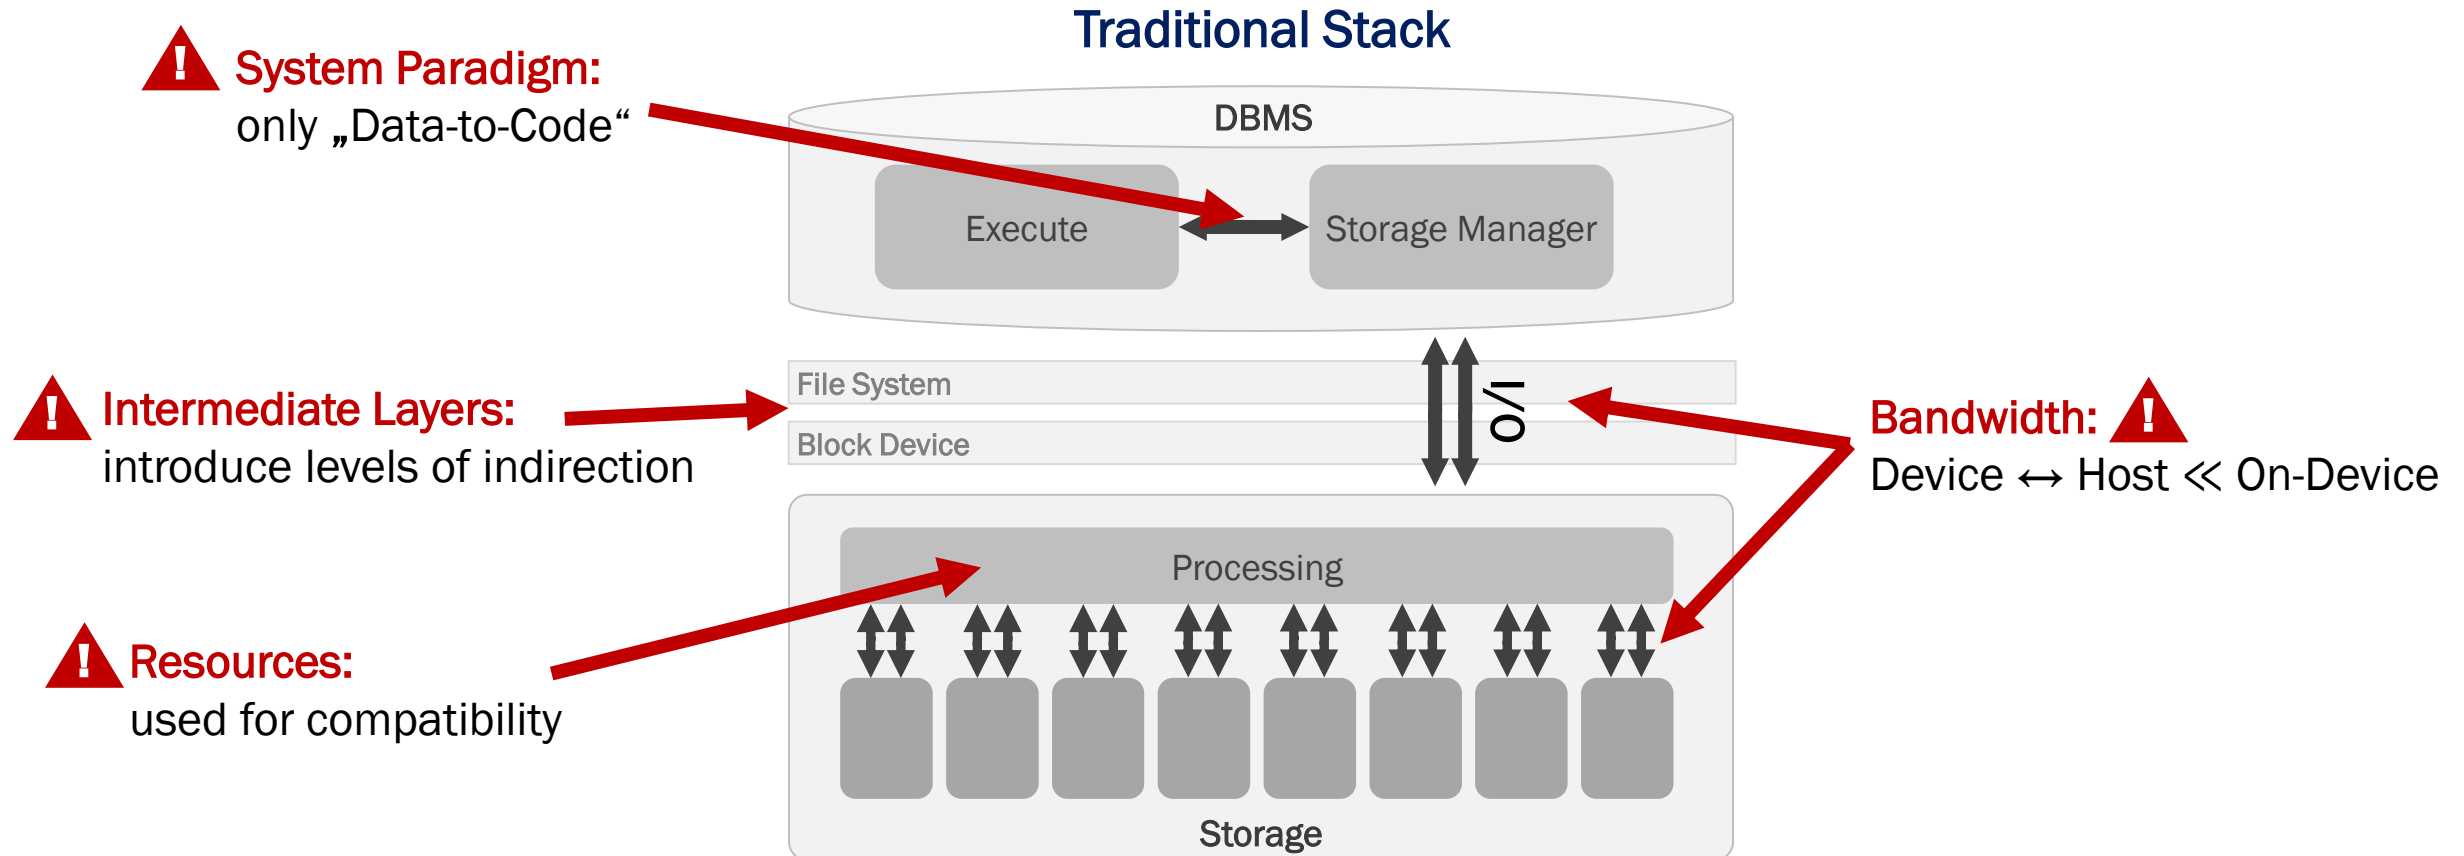
Switch to presentation during execution to explain the issues we solve with nKV

Read Amplification of LSM-Trees



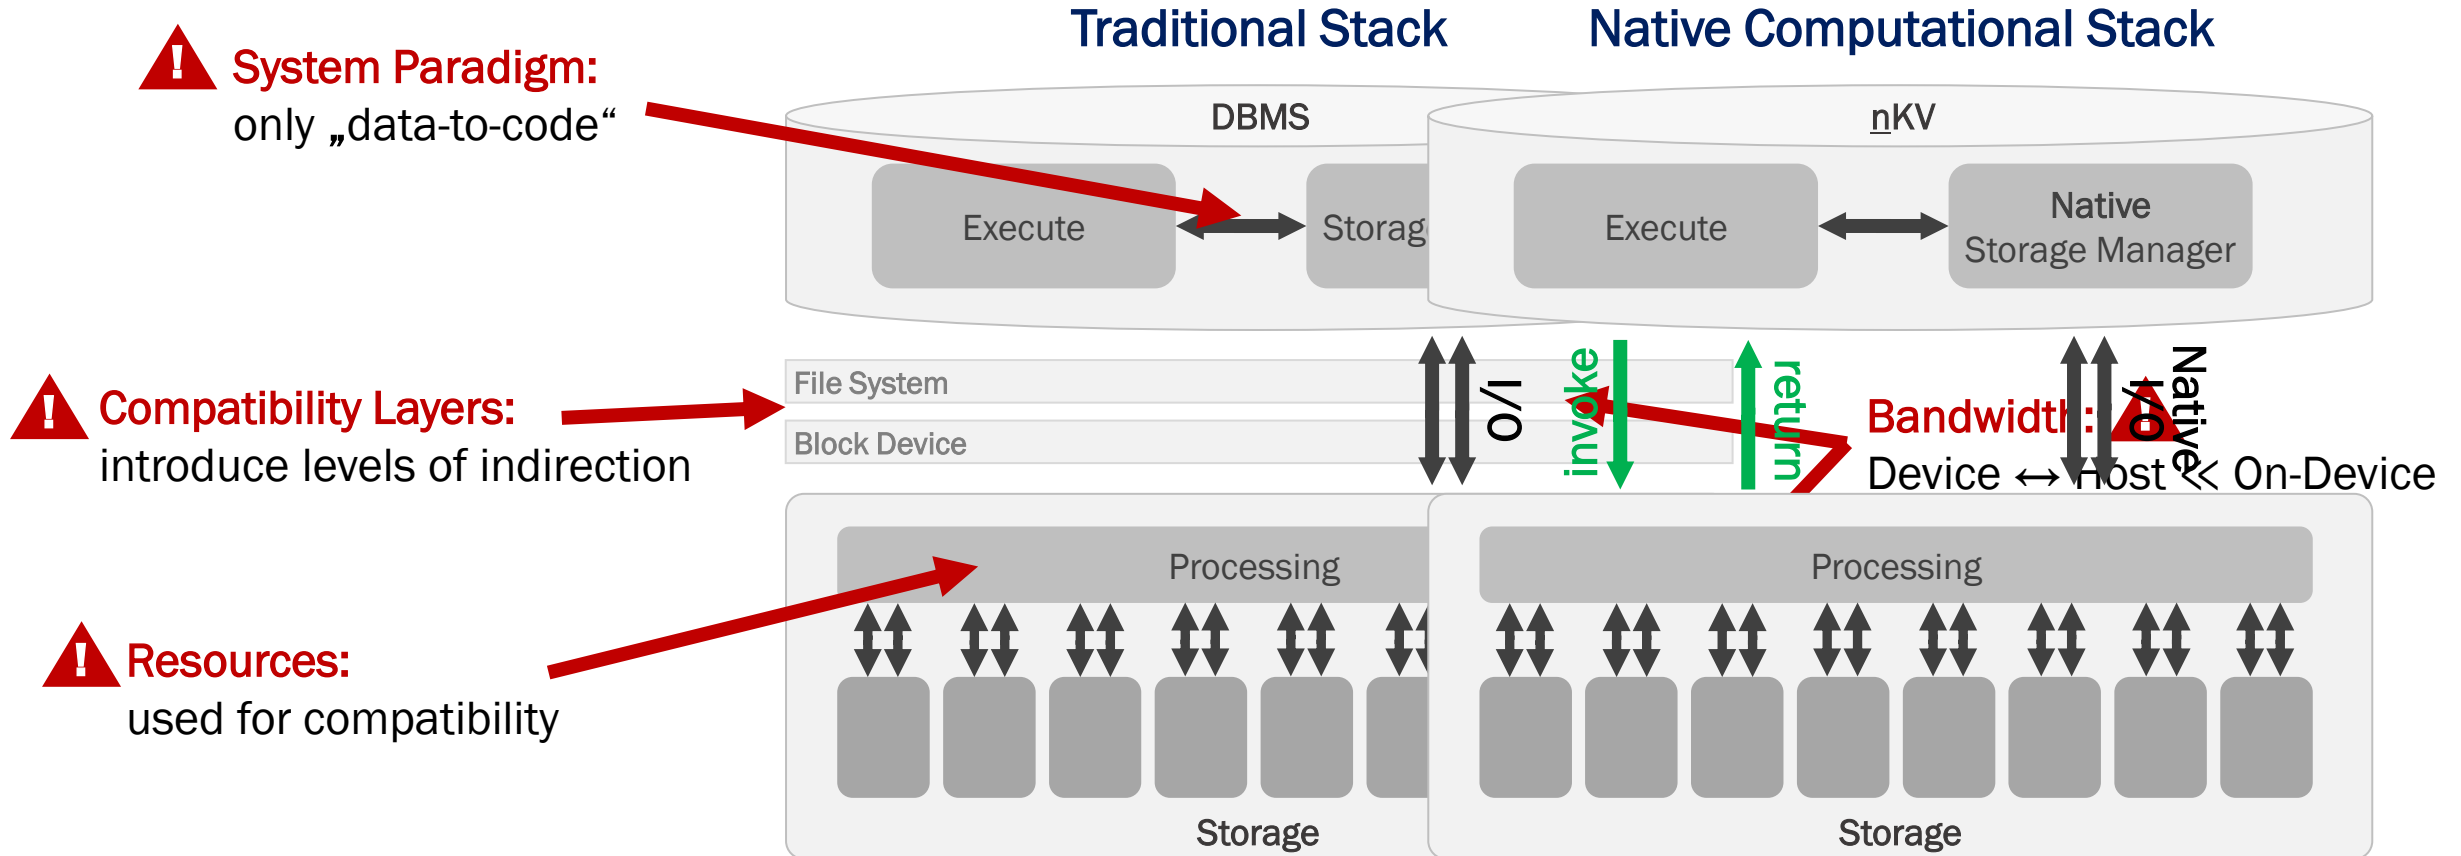
One request might end up in reading the entire data, but is that all?

Drawbacks of Traditional Stacks



Several elements impair performance – let's do it differently!

Native Computational Stack



nKV uses NDP to utilise on-device resources and to reduce data transfers



Switch to nKV VLDB Demo GUI

1. Discuss Runtimes BC
2. Execute Get on best BC Result (Duration 1-2 second)
 - Discuss Runtimes Get
 - Discuss Best Paper Award Check?
3. Execute Scan on Paper same year (Duration few seconds)
 - Discuss Runtimes Scan



nKV in Action

native Computational Storage

- Elimination of intermediary layers
- Physical data placement
- Embedded NVMe interface



near-Data Processing

- NDP interface extension
- Data format parsers and accessors
- Flexible scheduling

Demo Showcases:

- Paper detail retrieval: **1.4x**
- Scan for papers: **>2x**
- BC calculation: **~2.7x**

More Details:

nKV: Near-Data Processing with KV-Stores on Native Computational Storage

T. Vincon, L. Weber, A. Bernhardt, A. Koch, I. Petrov.

In Proc. [DAMON 2020](#).



<https://dblab.reutlingen-university.de>
<https://www.esa.informatik.tu-darmstadt.de>



TECHNISCHE
UNIVERSITÄT
DARMSTADT



Hochschule
Reutlingen
University

