Marcin Chrapek, Siyuan Shen, Patrick Iff, Tiancheng Chen, Mikhail Khalilov, Marcin Copik, Maciej Besta, Torsten Hoefler

SecPerf: Demystifying Cost of Confidential HPC



Two trends in HPC

Two trends in HPC



HPC converges with the cloud

Two trends in HPC



HPC converges with the cloud

Big Data and HPC Convergence: The Cutting Edge and Outlook

Sardar Usman[1(✉)], Rashid Mehmood[2], and Iyad Katib[1]

[1] Department of Computer Science, FCIT, King Abdulaziz University, Jeddah 21589, Saudi Arabia
usmansardar@hotmail.com, iakatib@kau.edu.sa
[2] High Performance Computing Center, King Abdulaziz University, Jeddah 21589, Saudi Arabia
RMehmood@kau.edu.sa

Two trends in HPC



HPC converges with the cloud

Big Data and HPC Convergence: The Cutting Edge and Outlook

Sardar Usman[1(✉)], Rashid Mehmood[2], and Iyad Katib[1]

[1] Department of Computer Science, FCIT, King Abdulaziz University, Jeddah 21589, Saudi Arabia
usmansardar@hotmail.com, iakatib@kau.edu.sa
[2] High Performance Computing Center, King Abdulaziz University, Jeddah 21589, Saudi Arabia
RMehmood@kau.edu.sa

XaaS: Acceleration as a Service to Enable Productive High-Performance Cloud Computing

Torsten Hoefler
ETH Zurich & Swiss National Supercomputing Centre (CSCS)
Switzerland

Marcin Copik
ETH Zurich
Switzerland

Pete Beckman
Argonne National Laboratory
USA

Andrew Jones
Microsoft
United Kingdom

Ian Foster
Argonne National Laboratory
USA

Manish Parashar
Utah University
USA

Daniel Reed
Utah University
USA

Matthias Troyer
Microsoft
USA

Thomas Schulthess
Swiss National Supercomputing Centre (CSCS)
Switzerland

Dan Ernst
NVIDIA
USA

Jack Dongarra
University of Tennessee
USA

Two trends in HPC

HPC converges with the cloud

Big Data and HPC Convergence: The Cutting Edge and Outlook

Sardar Usman[1(✉)], Rashid Mehmood[2], and Iyad Katib[1]

[1] Department of Computer Science, FCIT, King Abdulaziz University, Jeddah 21589, Saudi Arabia
usmansardar@hotmail.com, iakatib@kau.edu.sa
[2] High Performance Computing Center, King Abdulaziz University, Jeddah 21589, Saudi Arabia
RMehmood@kau.edu.sa

XaaS: Acceleration as a Service to Enable Productive High-Performance Cloud Computing

Torsten Hoefler, ETH Zurich & Swiss National Supercomputing Centre (CSCS), Switzerland
Marcin Copik, ETH Zurich, Switzerland
Pete Beckman, Argonne National Laboratory, USA
Andrew Jones, Microsoft, United Kingdom
Ian Foster, Argonne National Laboratory, USA
Manish Parashar, Utah University, USA
Daniel Reed, Utah University, USA
Matthias Troyer, Microsoft, USA
Thomas Schulthess, Swiss National Supercomputing Centre (CSCS), Switzerland
Dan Ernst, NVIDIA, USA
Jack Dongarra, University of Tennessee, USA

One Step Closer to Converged Computing: Achieving Scalability with Cloud-Native HPC

Daniel J. Milroy, *Lawrence Livermore National Laboratory*, Livermore, CA, USA, milroy1@llnl.gov
Claudia Misale, *IBM T.J. Watson Research Center*, Yorktown Heights, NY, USA, c.misale@ibm.com
Giorgis Georgakoudis, *Lawrence Livermore National Laboratory*, Livermore, CA, USA, georgakoudis1@llnl.gov
Tonia Elengikal, *IBM T.J. Watson Research Center*, Yorktown Heights, NY, USA, tonia.elengikal@ibm.com
Abhik Sarkar, *Lawrence Livermore National Laboratory*, Livermore, CA, USA, sarkar6@llnl.gov
Maurizio Drocco, *IBM T.J. Watson Research Center*, Yorktown Heights, NY, USA, maurizio.drocco@ibm.com
Tapasya Patki, *Lawrence Livermore National Laboratory*, Livermore, CA, USA, patki1@llnl.gov
Jae-Seung Yeom, *Lawrence Livermore National Laboratory*, Livermore, CA, USA, yeom2@llnl.gov
Carlos Eduardo Arango Gutierrez, *Red Hat*, Raleigh, NC, USA, carangog@redhat.com
Dong H. Ahn, *NVIDIA Corporation*, Santa Clara, CA, USA, donga@nvidia.com
Yoonho Park, *IBM T.J. Watson Research Center*, Yorktown Heights, NY, USA, yoonho@us.ibm.com

Two trends in HPC

HPC converges with the cloud

The Convergence of Hyperscale Data Center and High-Performance Computing Networks

Torsten Hoefler, ETH Zurich
Ariel Hendel, Scala Computing
Duncan Roweth, Hewlett Packard Enterprise

Big Data and HPC Convergence: The Cutting Edge and Outlook

Sardar Usman[1(✉)], Rashid Mehmood[2], and Iyad Katib[1]

[1] Department of Computer Science, FCIT, King Abdulaziz University, Jeddah 21589, Saudi Arabia
usmansardar@hotmail.com, iakatib@kau.edu.sa
[2] High Performance Computing Center, King Abdulaziz University, Jeddah 21589, Saudi Arabia
RMehmood@kau.edu.sa

XaaS: Acceleration as a Service to Enable Productive High-Performance Cloud Computing

Torsten Hoefler, ETH Zurich & Swiss National Supercomputing Centre (CSCS), Switzerland
Marcin Copik, ETH Zurich, Switzerland
Pete Beckman, Argonne National Laboratory, USA
Andrew Jones, Microsoft, United Kingdom
Ian Foster, Argonne National Laboratory, USA
Manish Parashar, Utah University, USA
Daniel Reed, Utah University, USA
Matthias Troyer, Microsoft, USA
Thomas Schulthess, Swiss National Supercomputing Centre (CSCS), Switzerland
Dan Ernst, NVIDIA, USA
Jack Dongarra, University of Tennessee, USA

One Step Closer to Converged Computing: Achieving Scalability with Cloud-Native HPC

Daniel J. Milroy, *Lawrence Livermore National Laboratory*, Livermore, CA, USA, milroy1@llnl.gov
Claudia Misale, *IBM T.J. Watson Research Center*, Yorktown Heights, NY, USA, c.misale@ibm.com
Giorgis Georgakoudis, *Lawrence Livermore National Laboratory*, Livermore, CA, USA, georgakoudis1@llnl.gov
Tonia Elengikal, *IBM T.J. Watson Research Center*, Yorktown Heights, NY, USA, tonia.elengikal@ibm.com
Abhik Sarkar, *Lawrence Livermore National Laboratory*, Livermore, CA, USA, sarkar6@llnl.gov
Maurizio Drocco, *IBM T.J. Watson Research Center*, Yorktown Heights, NY, USA, maurizio.drocco@ibm.com
Tapasya Patki, *Lawrence Livermore National Laboratory*, Livermore, CA, USA, patki1@llnl.gov
Jae-Seung Yeom, *Lawrence Livermore National Laboratory*, Livermore, CA, USA, yeom2@llnl.gov
Carlos Eduardo Arango Gutierrez, *Red Hat*, Raleigh, NC, USA, carangog@redhat.com
Dong H. Ahn, *NVIDIA Corporation*, Santa Clara, CA, USA, donga@nvidia.com
Yoonho Park, *IBM T.J. Watson Research Center*, Yorktown Heights, NY, USA, yoonho@us.ibm.com

Two trends in HPC

HPC converges with the cloud

HPC ➔ Cloud

"[…] data center networks […] incorporate […] HPC […] "

The Convergence of Hyperscale Data Center and High-Performance Computing Networks

Torsten Hoefler, ETH Zurich
Ariel Hendel, Scala Computing
Duncan Roweth, Hewlett Packard Enterprise

Big Data and HPC Convergence: The Cutting Edge and Outlook

Sardar Usman[1(✉)], Rashid Mehmood[2], and Iyad Katib[1]

[1] Department of Computer Science, FCIT, King Abdulaziz University, Jeddah 21589, Saudi Arabia
usmansardar@hotmail.com, iakatib@kau.edu.sa
[2] High Performance Computing Center, King Abdulaziz University, Jeddah 21589, Saudi Arabia
RMehmood@kau.edu.sa

XaaS: Acceleration as a Service to Enable Productive High-Performance Cloud Computing

Torsten Hoefler, ETH Zurich & Swiss National Supercomputing Centre (CSCS), Switzerland
Marcin Copik, ETH Zurich, Switzerland
Pete Beckman, Argonne National Laboratory, USA
Andrew Jones, Microsoft, United Kingdom
Ian Foster, Argonne National Laboratory, USA
Manish Parashar, Utah University, USA
Daniel Reed, Utah University, USA
Matthias Troyer, Microsoft, USA
Thomas Schulthess, Swiss National Supercomputing Centre (CSCS), Switzerland
Dan Ernst, NVIDIA, USA
Jack Dongarra, University of Tennessee, USA

One Step Closer to Converged Computing: Achieving Scalability with Cloud-Native HPC

Daniel J. Milroy, *Lawrence Livermore National Laboratory*, Livermore, CA, USA, milroy1@llnl.gov
Claudia Misale, *IBM T.J. Watson Research Center*, Yorktown Heights, NY, USA, c.misale@ibm.com
Giorgis Georgakoudis, *Lawrence Livermore National Laboratory*, Livermore, CA, USA, georgakoudis1@llnl.gov
Tonia Elengikal, *IBM T.J. Watson Research Center*, Yorktown Heights, NY, USA, tonia.elengikal@ibm.com
Abhik Sarkar, *Lawrence Livermore National Laboratory*, Livermore, CA, USA, sarkar6@llnl.gov
Maurizio Drocco, *IBM T.J. Watson Research Center*, Yorktown Heights, NY, USA, maurizio.drocco@ibm.com
Tapasya Patki, *Lawrence Livermore National Laboratory*, Livermore, CA, USA, patki1@llnl.gov
Jae-Seung Yeom, *Lawrence Livermore National Laboratory*, Livermore, CA, USA, yeom2@llnl.gov
Carlos Eduardo Arango Gutierrez, *Red Hat*, Raleigh, NC, USA, carangog@redhat.com
Dong H. Ahn, *NVIDIA Corporation*, Santa Clara, CA, USA, donga@nvidia.com
Yoonho Park, *IBM T.J. Watson Research Center*, Yorktown Heights, NY, USA, yoonho@us.ibm.com

Two trends in HPC

HPC converges with the cloud

HPC ➔ Cloud
"[…] data center networks […] incorporate […] HPC […] "

Cloud ➔ HPC
"[…] we […] bring innovative cloud techniques into HPC"

The Convergence of Hyperscale Data Center and High-Performance Computing Networks

Torsten Hoefler, ETH Zurich
Ariel Hendel, Scala Computing
Duncan Roweth, Hewlett Packard Enterprise

Big Data and HPC Convergence: The Cutting Edge and Outlook

Sardar Usman[1(✉)], Rashid Mehmood[2], and Iyad Katib[1]

[1] Department of Computer Science, FCIT, King Abdulaziz University, Jeddah 21589, Saudi Arabia
usmansardar@hotmail.com, iakatib@kau.edu.sa
[2] High Performance Computing Center, King Abdulaziz University, Jeddah 21589, Saudi Arabia
RMehmood@kau.edu.sa

XaaS: Acceleration as a Service to Enable Productive High-Performance Cloud Computing

Torsten Hoefler, ETH Zurich & Swiss National Supercomputing Centre (CSCS), Switzerland
Marcin Copik, ETH Zurich, Switzerland
Pete Beckman, Argonne National Laboratory, USA
Andrew Jones, Microsoft, United Kingdom
Ian Foster, Argonne National Laboratory, USA
Manish Parashar, Utah University, USA
Daniel Reed, Utah University, USA
Matthias Troyer, Microsoft, USA
Thomas Schulthess, Swiss National Supercomputing Centre (CSCS), Switzerland
Dan Ernst, NVIDIA, USA
Jack Dongarra, University of Tennessee, USA

One Step Closer to Converged Computing: Achieving Scalability with Cloud-Native HPC

Daniel J. Milroy, *Lawrence Livermore National Laboratory*, Livermore, CA, USA, milroy1@llnl.gov
Claudia Misale, *IBM T.J. Watson Research Center*, Yorktown Heights, NY, USA, c.misale@ibm.com
Giorgis Georgakoudis, *Lawrence Livermore National Laboratory*, Livermore, CA, USA, georgakoudis1@llnl.gov
Tonia Elengikal, *IBM T.J. Watson Research Center*, Yorktown Heights, NY, USA, tonia.elengikal@ibm.com
Abhik Sarkar, *Lawrence Livermore National Laboratory*, Livermore, CA, USA, sarkar6@llnl.gov
Maurizio Drocco, *IBM T.J. Watson Research Center*, Yorktown Heights, NY, USA, maurizio.drocco@ibm.com
Tapasya Patki, *Lawrence Livermore National Laboratory*, Livermore, CA, USA, patki1@llnl.gov
Jae-Seung Yeom, *Lawrence Livermore National Laboratory*, Livermore, CA, USA, yeom2@llnl.gov
Carlos Eduardo Arango Gutierrez, *Red Hat*, Raleigh, NC, USA, carangog@redhat.com
Dong H. Ahn, *NVIDIA Corporation*, Santa Clara, CA, USA, donga@nvidia.com
Yoonho Park, *IBM T.J. Watson Research Center*, Yorktown Heights, NY, USA, yoonho@us.ibm.com

Two trends in HPC

HPC converges with the cloud

HPC ➔ Cloud
"[…] data center networks […] incorporate […] HPC […] "

Cloud ➔ HPC
"[…] we […] bring innovative cloud techniques into HPC"

The Convergence of Hyperscale Data Center and High-Performance Computing Networks

Torsten Hoefler, ETH Zurich
Ariel Hendel, Scala Computing
Duncan Roweth, Hewlett Packard Enterprise

High Performance Serverless for HPC and Clouds

Doctoral Thesis

Author(s):
Copik, Marcin

Big Data and HPC Convergence: The Cutting Edge and Outlook

Sardar Usman[1(✉)], Rashid Mehmood[2], and Iyad Katib[1]

[1] Department of Computer Science, FCIT, King Abdulaziz University, Jeddah 21589, Saudi Arabia
usmansardar@hotmail.com, iakatib@kau.edu.sa
[2] High Performance Computing Center, King Abdulaziz University, Jeddah 21589, Saudi Arabia
RMehmood@kau.edu.sa

XaaS: Acceleration as a Service to Enable Productive High-Performance Cloud Computing

Torsten Hoefler, ETH Zurich & Swiss National Supercomputing Centre (CSCS), Switzerland
Marcin Copik, ETH Zurich, Switzerland
Pete Beckman, Argonne National Laboratory, USA
Andrew Jones, Microsoft, United Kingdom
Ian Foster, Argonne National Laboratory, USA
Manish Parashar, Utah University, USA
Daniel Reed, Utah University, USA
Matthias Troyer, Microsoft, USA
Thomas Schulthess, Swiss National Supercomputing Centre (CSCS), Switzerland
Dan Ernst, NVIDIA, USA
Jack Dongarra, University of Tennessee, USA

One Step Closer to Converged Computing: Achieving Scalability with Cloud-Native HPC

Daniel J. Milroy, *Lawrence Livermore National Laboratory*, Livermore, CA, USA, milroy1@llnl.gov
Claudia Misale, *IBM T.J. Watson Research Center*, Yorktown Heights, NY, USA, c.misale@ibm.com
Giorgis Georgakoudis, *Lawrence Livermore National Laboratory*, Livermore, CA, USA, georgakoudis1@llnl.gov
Tonia Elengikal, *IBM T.J. Watson Research Center*, Yorktown Heights, NY, USA, tonia.elengikal@ibm.com
Abhik Sarkar, *Lawrence Livermore National Laboratory*, Livermore, CA, USA, sarkar6@llnl.gov
Maurizio Drocco, *IBM T.J. Watson Research Center*, Yorktown Heights, NY, USA, maurizio.drocco@ibm.com
Tapasya Patki, *Lawrence Livermore National Laboratory*, Livermore, CA, USA, patki1@llnl.gov
Jae-Seung Yeom, *Lawrence Livermore National Laboratory*, Livermore, CA, USA, yeom2@llnl.gov
Carlos Eduardo Arango Gutierrez, *Red Hat*, Raleigh, NC, USA, carangog@redhat.com
Dong H. Ahn, *NVIDIA Corporation*, Santa Clara, CA, USA, donga@nvidia.com
Yoonho Park, *IBM T.J. Watson Research Center*, Yorktown Heights, NY, USA, yoonho@us.ibm.com

Two trends in HPC



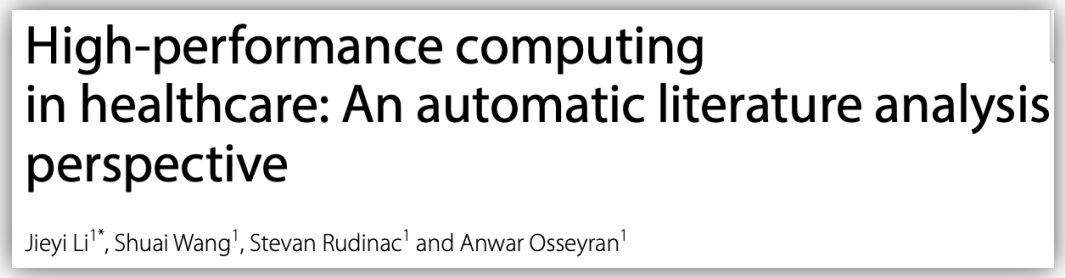
Operating on sensitive data & produce company secrets

High-performance computing in healthcare: An automatic literature analysis perspective

Jieyi Li[1*], Shuai Wang[1], Stevan Rudinac[1] and Anwar Osseyran[1]

Two trends in HPC

Operating on sensitive data & produce company secrets



High-performance computing in healthcare: An automatic literature analysis perspective

Jieyi Li[1*], Shuai Wang[1], Stevan Rudinac[1] and Anwar Osseyran[1]



Bringing Big Data to Personalized Healthcare: A Patient-Centered Framework

Nitesh V. Chawla, PhD[1] and Darcy A. Davis, PhD[2]

[1]Department of Computer Science and Engineering, Interdisciplinary Center for Network Science and Applications (iCeNSA), University of Notre Dame, Notre Dame, IN, USA; [2]University of Notre Dame, Notre Dame, IN, USA.

Two trends in HPC



Operating on sensitive data & produce company secrets

High-performance computing in healthcare: An automatic literature analysis perspective

Jieyi Li[1*], Shuai Wang[1], Stevan Rudinac[1] and Anwar Osseyran[1]

Bringing Big Data to Personalized Healthcare: A Patient-Centered Framework

Nitesh V. Chawla, PhD[1] and Darcy A. Davis, PhD[2]

[1]Department of Computer Science and Engineering, Interdisciplinary Center for Network Science and Applications (iCeNSA), University of Notre Dame, Notre Dame, IN, USA; [2]University of Notre Dame, Notre Dame, IN, USA.

LawGPT: Knowledge-Guided Data Generation and Its Application to Legal LLM

Zhi Zhou[1], Kun-Yang Yu[1,2], Shi-Yu Tian[1,2], Xiao-Wen Yang[1,2], Jiang-Xin Shi[1,2], Peng-Xiao Song[1], Yi-Xuan Jin[1], Lan-Zhe Guo[1,3,*], Yu-Feng Li[1,2,*]
[1]National Key Laboratory for Novel Software Technology, Nanjing University, China
[2]School of Artificial Intelligence, Nanjing University, China
[3]School of Intelligence Science and Technology, Nanjing University, China

Two trends in HPC

Operating on sensitive data & produce company secrets

High-performance computing in healthcare: An automatic literature analysis perspective

Jieyi Li[1*], Shuai Wang[1], Stevan Rudinac[1] and Anwar Osseyran[1]

Bringing Big Data to Personalized Healthcare: A Patient-Centered Framework

Nitesh V. Chawla, PhD[1] and Darcy A. Davis, PhD[2]

[1]Department of Computer Science and Engineering, Interdisciplinary Center for Network Science and Applications (iCeNSA), University of Notre Dame, Notre Dame, IN, USA; [2]University of Notre Dame, Notre Dame, IN, USA.

LawGPT: Knowledge-Guided Data Generation and Its Application to Legal LLM

Zhi Zhou[1], Kun-Yang Yu[1,2], Shi-Yu Tian[1,2], Xiao-Wen Yang[1,2], Jiang-Xin Shi[1,2], Peng-Xiao Song[1], Yi-Xuan Jin[1], Lan-Zhe Guo[1,3,*], Yu-Feng Li[1,2,*]
[1]National Key Laboratory for Novel Software Technology, Nanjing University, China
[2]School of Artificial Intelligence, Nanjing University, China
[3]School of Intelligence Science and Technology, Nanjing University, China

FinGPT-HPC: Efficient Pretraining and Finetuning Large Language Models for Financial Applications with High-Performance Computing

Xiao-Yang Liu[1], Jie Zhang[2], Guoxuan Wang[3], Weiqing Tong[2], and Anwar Walid[1]
[1]Department of Electrical Engineering, Columbia University
[2]School of Computer Engineering and Science, Shanghai University
[3]Department of Computer Science, Whiting School of Engineering, Johns Hopkins University
Emails: {XL2427,aie13}@columbia.edu, {zhangjie,wqtong}@shu.edu.cn, gwang69@jhu.com

Two trends in HPC

Operating on sensitive data & produce company secrets

High-performance computing in healthcare: An automatic literature analysis perspective

Jieyi Li[1*], Shuai Wang[1], Stevan Rudinac[1] and Anwar Osseyran[1]

Bringing Big Data to Personalized Healthcare: A Patient-Centered Framework

Nitesh V. Chawla, PhD[1] and Darcy A. Davis, PhD[2]

[1]Department of Computer Science and Engineering, Interdisciplinary Center for Network Science and Applications (iCeNSA), University of Notre Dame, Notre Dame, IN, USA; [2]University of Notre Dame, Notre Dame, IN, USA.

LawGPT: Knowledge-Guided Data Generation and Its Application to Legal LLM

Zhi Zhou[1], Kun-Yang Yu[1,2], Shi-Yu Tian[1,2], Xiao-Wen Yang[1,2], Jiang-Xin Shi[1,2], Peng-Xiao Song[1], Yi-Xuan Jin[1], Lan-Zhe Guo[1,3,*], Yu-Feng Li[1,2,*]
[1]National Key Laboratory for Novel Software Technology, Nanjing University, China
[2]School of Artificial Intelligence, Nanjing University, China
[3]School of Intelligence Science and Technology, Nanjing University, China

FinGPT-HPC: Efficient Pretraining and Finetuning Large Language Models for Financial Applications with High-Performance Computing

Xiao-Yang Liu[1], Jie Zhang[2], Guoxuan Wang[3], Weiqing Tong[2], and Anwar Walid[1]
[1]Department of Electrical Engineering, Columbia University
[2]School of Computer Engineering and Science, Shanghai University
[3]Department of Computer Science, Whiting School of Engineering, Johns Hopkins University
Emails: {XL2427,aie13}@columbia.edu, {zhangjie,wqtong}@shu.edu.cn, gwang69@jhu.com

Artificial Intelligence Index Report 2024

Stanford University Human-Centered Artificial Intelligence

*"[…] OpenAI's GPT-4 used an estimated **$78 million worth of compute to train, while Google's Gemini Ultra cost $191 million** […]"*

Security and privacy when converging HPC and cloud

Security and privacy when converging HPC and cloud

CLOUD COMPUTING IN HEALTHCARE AND BIOMEDICINE

BARBARA CALABRESE* AND MARIO CANNATARO†

Security and privacy when converging HPC and cloud

CLOUD COMPUTING IN HEALTHCARE AND BIOMEDICINE

BARBARA CALABRESE* AND MARIO CANNATARO†

"The power of HPC […] is […] parallel […] computing."

↓

Need HPC for performance and scalability

Security and privacy when converging HPC and cloud

CLOUD COMPUTING IN HEALTHCARE AND BIOMEDICINE

BARBARA CALABRESE* AND MARIO CANNATARO†

"The power of HPC […] is […] parallel […] computing."

"[…] allocation of resources is flexible and elastic […]"

Need HPC for performance and scalability

Want cloud for efficiency, easy of access, flexibility

Security and privacy when converging HPC and cloud

CLOUD COMPUTING IN HEALTHCARE AND BIOMEDICINE

BARBARA CALABRESE* AND MARIO CANNATARO†

"The power of HPC […] is […] parallel […] computing."

"*[…] allocation of resources is flexible and elastic […]*"

"*[…] security and privacy […] are the main challenges […] on the cloud […]*"

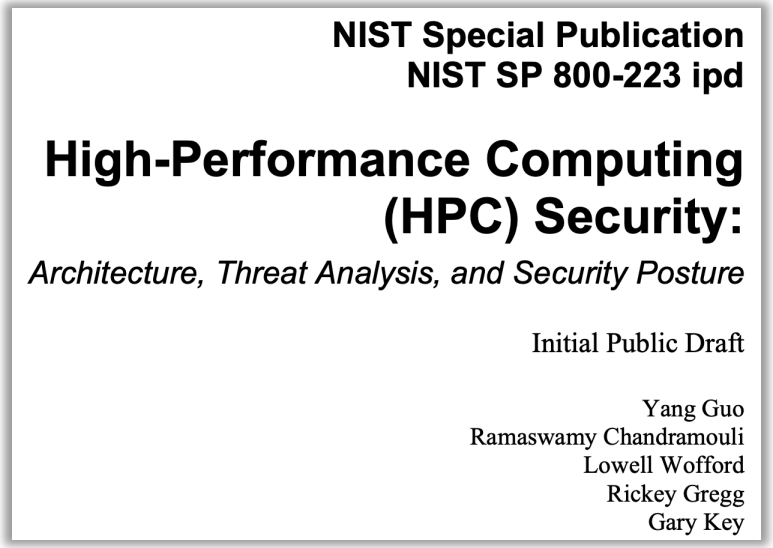
Need HPC for performance and scalability

Want cloud for efficiency, easy of access, flexibility

Trusted HPC cluster becomes untrusted cloud

The emergence of confidential HPC (cHPC)

The emergence of confidential HPC (cHPC)



NIST Special Publication
NIST SP 800-223 ipd

High-Performance Computing (HPC) Security:

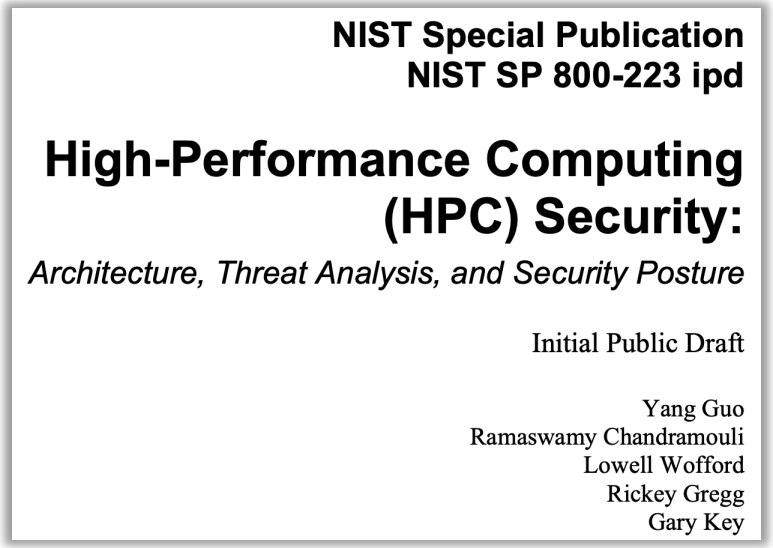
Architecture, Threat Analysis, and Security Posture

Initial Public Draft

Yang Guo
Ramaswamy Chandramouli
Lowell Wofford
Rickey Gregg
Gary Key

> *"Security is essential to achieving the [...] benefits of HPC [...]"*

The emergence of confidential HPC (cHPC)



NIST Special Publication
NIST SP 800-223 ipd

High-Performance Computing (HPC) Security:

Architecture, Threat Analysis, and Security Posture

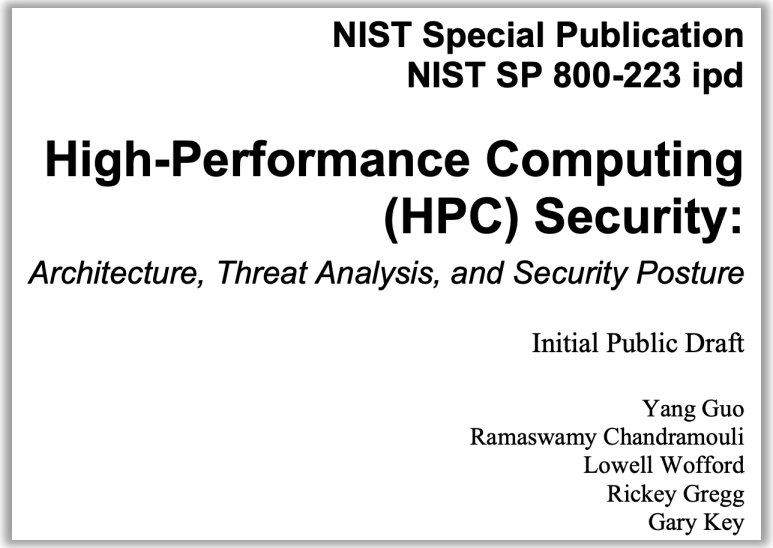
Initial Public Draft

Yang Guo
Ramaswamy Chandramouli
Lowell Wofford
Rickey Gregg
Gary Key

> *"Security is essential to achieving the [...] benefits of HPC [...]"*

> *"[...] At Anthropic, we [...] ensure [...] trust is cryptographically guaranteed. [...]"*

The emergence of confidential HPC (cHPC)



NIST Special Publication
NIST SP 800-223 ipd

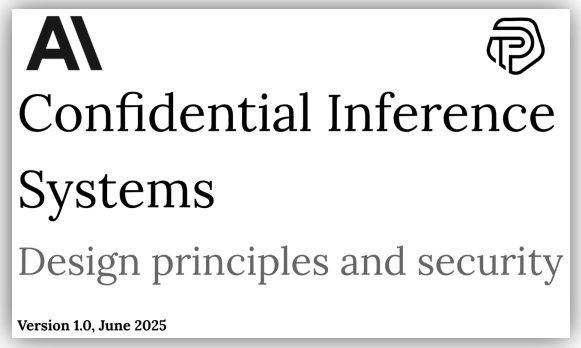
High-Performance Computing (HPC) Security:

Architecture, Threat Analysis, and Security Posture

Initial Public Draft

Yang Guo
Ramaswamy Chandramouli
Lowell Wofford
Rickey Gregg
Gary Key

> *"Security is essential to achieving the [...] benefits of HPC [...]"*



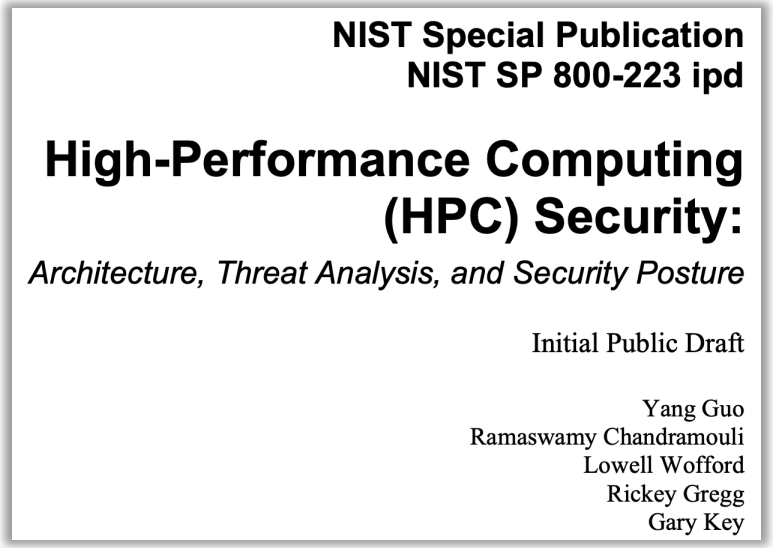
Confidential Inference Systems

Design principles and security

Version 1.0, June 2025

> *"[...] At Anthropic, we [...] ensure [...] trust is cryptographically guaranteed. [...]"*

The emergence of confidential HPC (cHPC)



NIST Special Publication
NIST SP 800-223 ipd

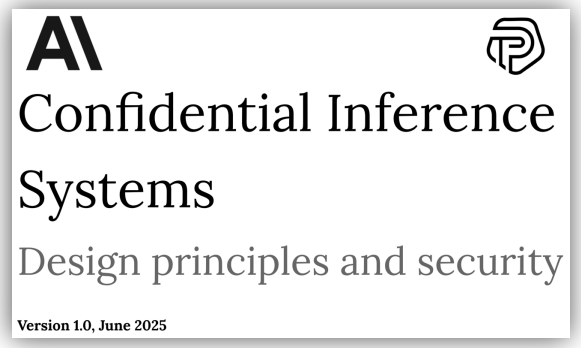
High-Performance Computing (HPC) Security:

Architecture, Threat Analysis, and Security Posture

Initial Public Draft

Yang Guo
Ramaswamy Chandramouli
Lowell Wofford
Rickey Gregg
Gary Key

"Security is essential to achieving the [...] benefits of HPC [...]"



Confidential Inference Systems

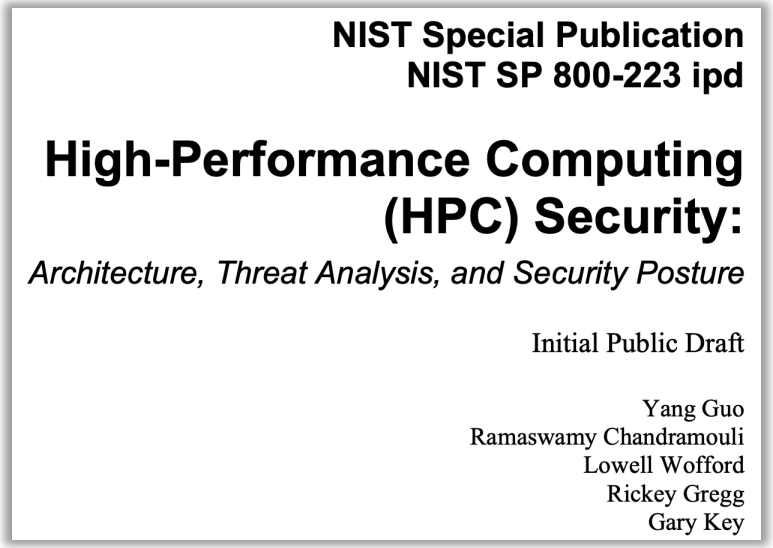
Design principles and security

Version 1.0, June 2025

"[...] At Anthropic, we [...] ensure [...] trust is cryptographically guaranteed. [...]"

HPC hardware releases introduce security (H100, B100, Graphcore IPUs)

The emergence of confidential HPC (cHPC)



NIST Special Publication
NIST SP 800-223 ipd

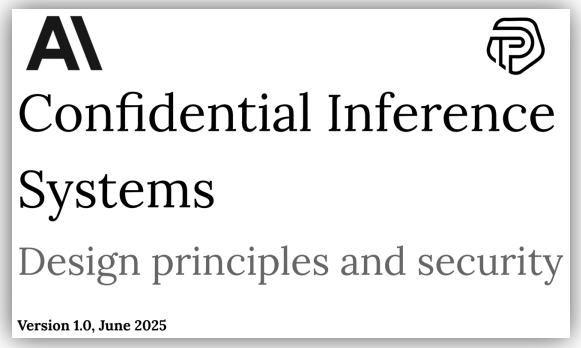
High-Performance Computing (HPC) Security:

Architecture, Threat Analysis, and Security Posture

Initial Public Draft

Yang Guo
Ramaswamy Chandramouli
Lowell Wofford
Rickey Gregg
Gary Key

"Security is essential to achieving the [...] benefits of HPC [...]"



Confidential Inference Systems

Design principles and security

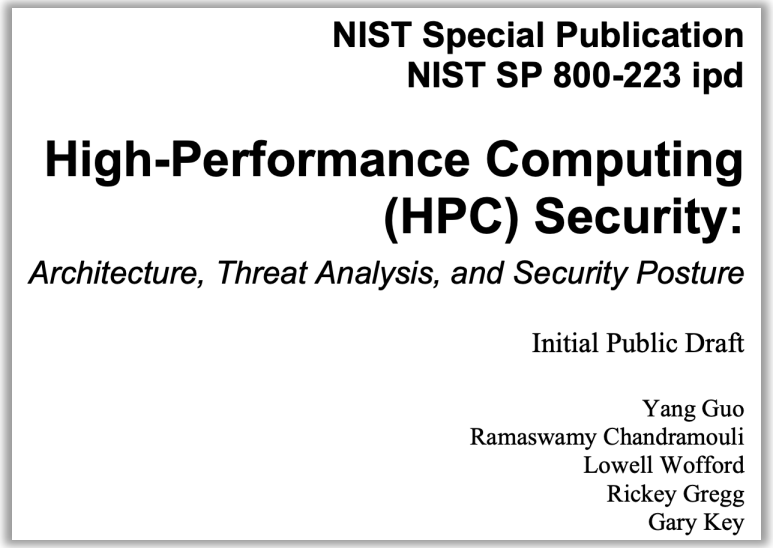
Version 1.0, June 2025

"[...] At Anthropic, we [...] ensure [...] trust is cryptographically guaranteed. [...]"

Confidential Computing on NVIDIA H100 GPUs for Secure and Trustworthy AI

HPC hardware releases introduce security (H100, B100, Graphcore IPUs)

The emergence of confidential HPC (cHPC)



NIST Special Publication
NIST SP 800-223 ipd

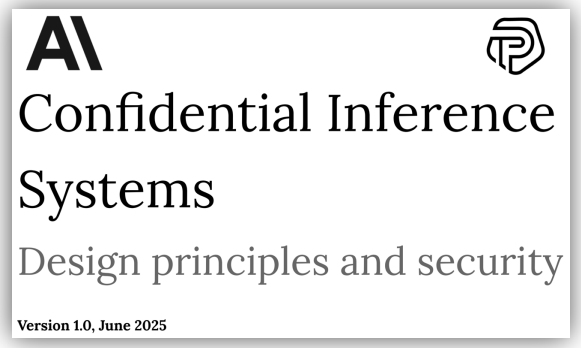
High-Performance Computing (HPC) Security:

Architecture, Threat Analysis, and Security Posture

Initial Public Draft

Yang Guo
Ramaswamy Chandramouli
Lowell Wofford
Rickey Gregg
Gary Key

> *"Security is essential to achieving the [...] benefits of HPC [...]"*



Confidential Inference Systems

Design principles and security

Version 1.0, June 2025

> *"[...] At Anthropic, we [...] ensure [...] trust is cryptographically guaranteed. [...]"*

Confidential Computing on NVIDIA H100 GPUs for Secure and Trustworthy AI

Secure AI

NVIDIA Blackwell includes NVIDIA Confidential Computing, which protects sensitive data and AI models from unauthorized access with strong hardware-based security. NVIDIA Blackwell is the first TEE-I/O capable GPU in the industry,

HPC hardware releases introduce security (H100, B100, Graphcore IPUs)

The emergence of confidential HPC (cHPC)

NIST Special Publication
NIST SP 800-223 ipd

High-Performance Computing (HPC) Security:

Architecture, Threat Analysis, and Security Posture

Initial Public Draft

Yang Guo
Ramaswamy Chandramouli
Lowell Wofford
Rickey Gregg
Gary Key

"Security is essential to achieving the [...] benefits of HPC [...]"

Confidential Inference Systems

Design principles and security

Version 1.0, June 2025

"[...] At Anthropic, we [...] ensure [...] trust is cryptographically guaranteed. [...]"

Confidential Computing on NVIDIA H100 GPUs for Secure and Trustworthy AI

Secure AI

NVIDIA Blackwell includes NVIDIA Confidential Computing, which protects sensitive data and AI models from unauthorized access with strong hardware-based security. NVIDIA Blackwell is the first TEE-I/O capable GPU in the industry,

Confidential Machine Learning within Graphcore IPUs

Kapil Vaswani,[1] Stavros Volos,[1] Cédric Fournet,[1] Antonio Nino Diaz,[1] Ken Gordon,[1] Balaji Vembu,[0] Sam Webster,[1] David Chisnall,[1]
Saurabh Kulkarni,[2] Graham Cunningham,[2] Richard Osborne,[2] Dan Wilkinson[2]
[1]Microsoft Research [2]Graphcore

HPC hardware releases introduce security (H100, B100, Graphcore IPUs)

The emergence of confidential HPC (cHPC)

NIST Special Publication
NIST SP 800-223 ipd

High-Performance Computing (HPC) Security:

Architecture, Threat Analysis, and Security Posture

Initial Public Draft

Yang Guo
Ramaswamy Chandramouli
Lowell Wofford
Rickey Gregg
Gary Key

"Security is essential to achieving the [...] benefits of HPC [...]"

Confidential Inference Systems

Design principles and security

Version 1.0, June 2025

"[...] At Anthropic, we [...] ensure [...] trust is cryptographically guaranteed. [...]"

Confidential Computing on NVIDIA H100 GPUs for Secure and Trustworthy AI

Secure AI

NVIDIA Blackwell includes NVIDIA Confidential Computing, which protects sensitive data and AI models from unauthorized access with strong hardware-based security. NVIDIA Blackwell is the first TEE-I/O capable GPU in the industry,

Confidential Machine Learning within Graphcore IPUs

Kapil Vaswani,[1] Stavros Volos,[1] Cédric Fournet,[1] Antonio Nino Diaz,[1] Ken Gordon,[1] Balaji Vembu,[0] Sam Webster,[1] David Chisnall,[1]
Saurabh Kulkarni,[2] Graham Cunningham,[2] Richard Osborne,[2] Dan Wilkinson[2]
[1]Microsoft Research [2]Graphcore

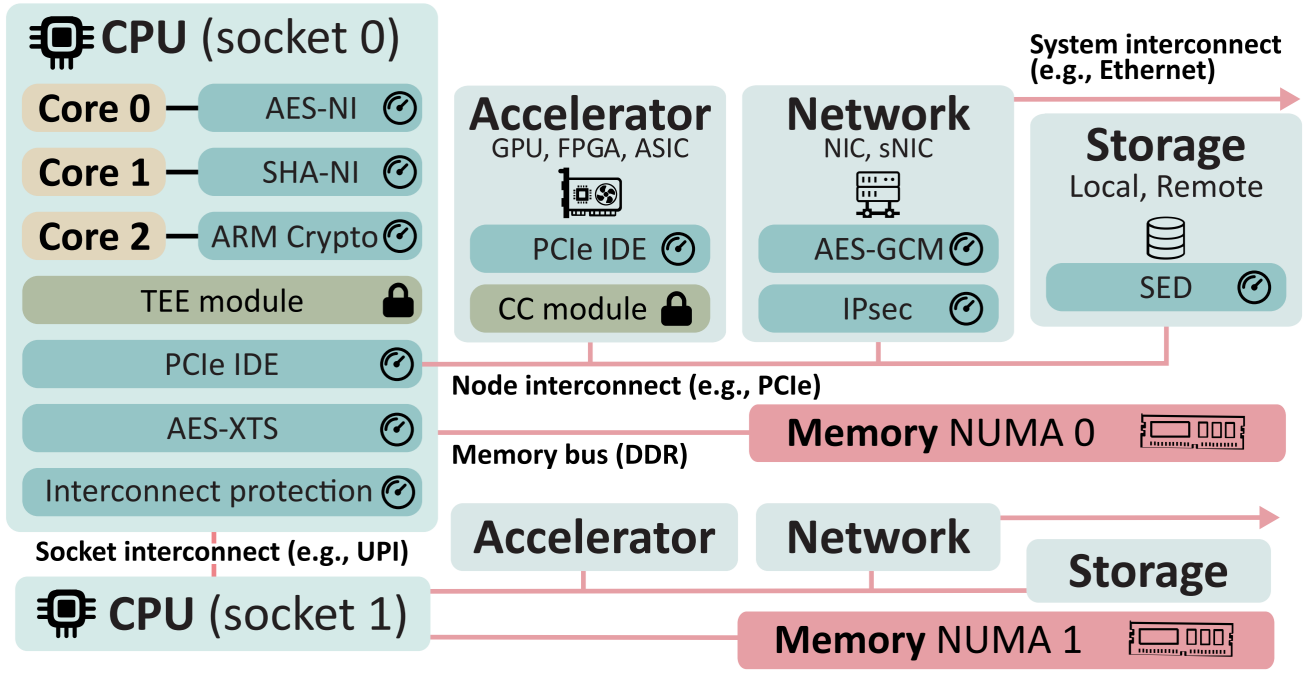
Security is anchored by the NVIDIA BlueField Advanced Secure Trusted Resource Architecture, enabling service providers to create secure, bare-metal compute instances with zero-trust tenant isolation and full software-defined infrastructure control.

HPC hardware releases introduce security (H100, B100, Graphcore IPUs)

Modern Secure Systems from an HPC perspective

Modern Secure Systems from an HPC perspective

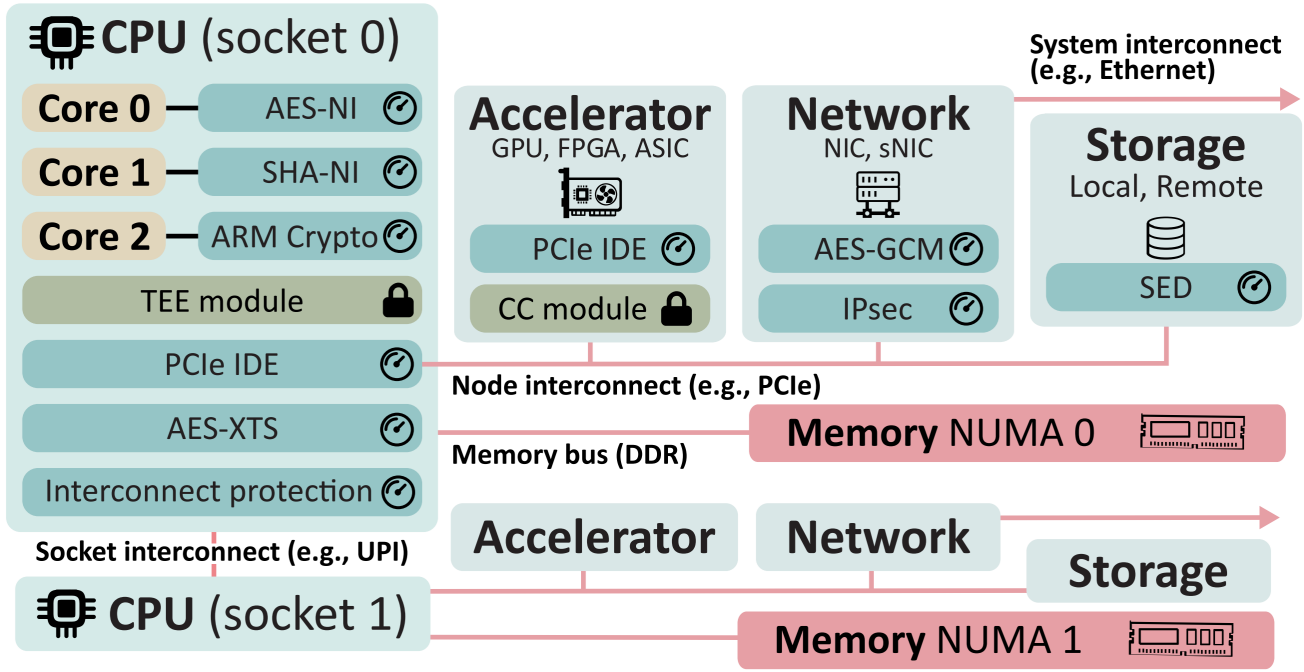
Heterogeneous HPC nodes have a lot of secure systems



Modern Secure Systems from an HPC perspective

Heterogeneous HPC nodes have a lot of secure systems

What are the overheads of these systems?



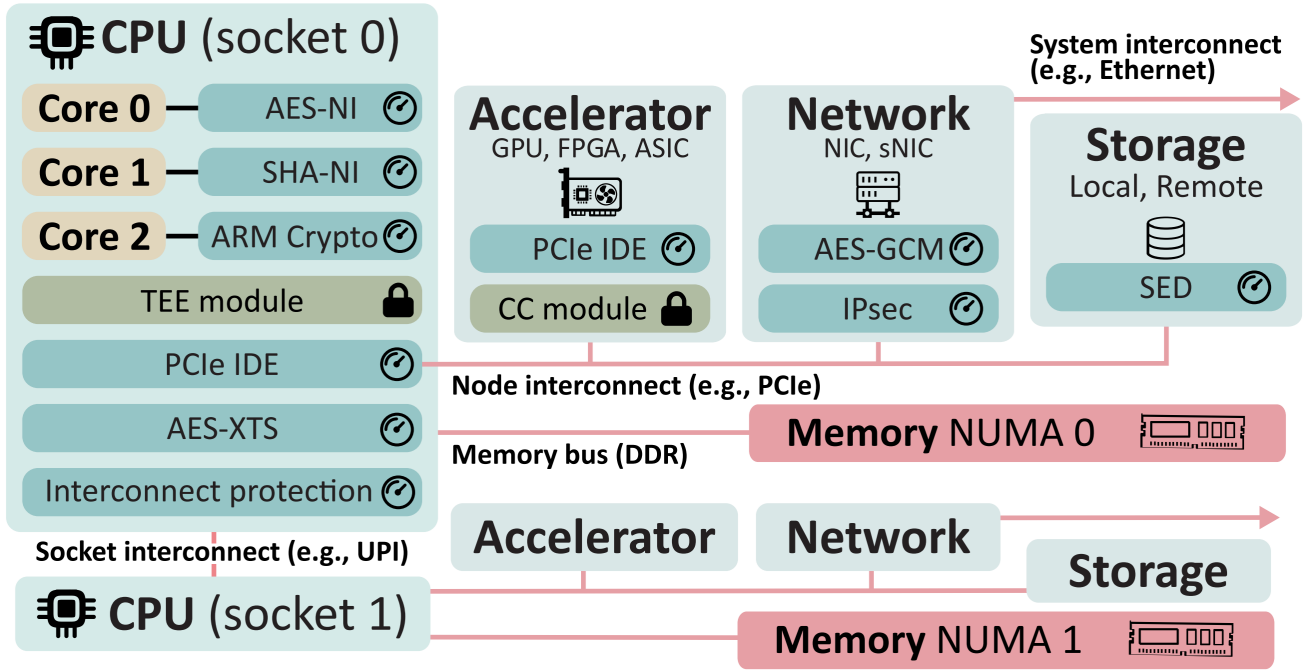
Modern Secure Systems from an HPC perspective

Heterogeneous HPC nodes have a lot of secure systems

What are the overheads of these systems?

cHPC needs to secure not only compute

How do tradeoffs for these areas look like?



Modern Secure Systems from an HPC perspective

Heterogeneous HPC nodes have a lot of secure systems

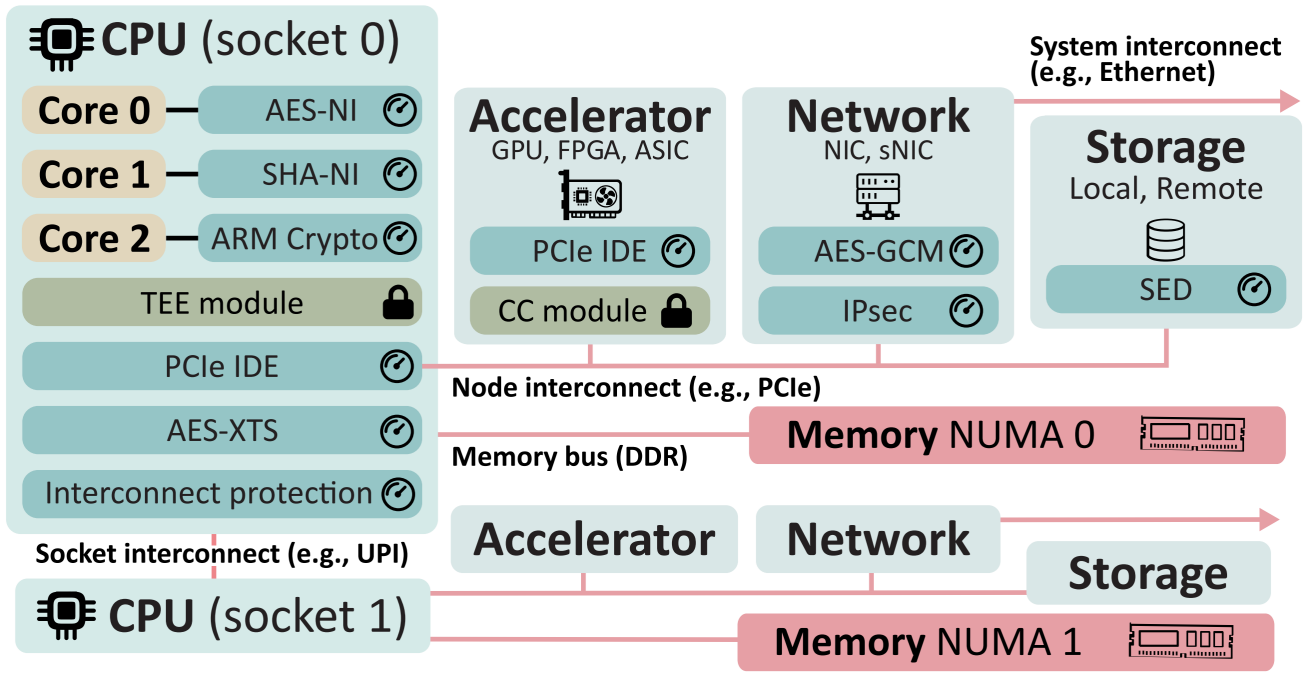
What are the overheads of these systems?

cHPC needs to secure not only compute

How do tradeoffs for these areas look like?

Applications have different security needs

How to characterize them?



Modern Secure Systems from an HPC perspective

Heterogeneous HPC nodes have a lot of secure systems

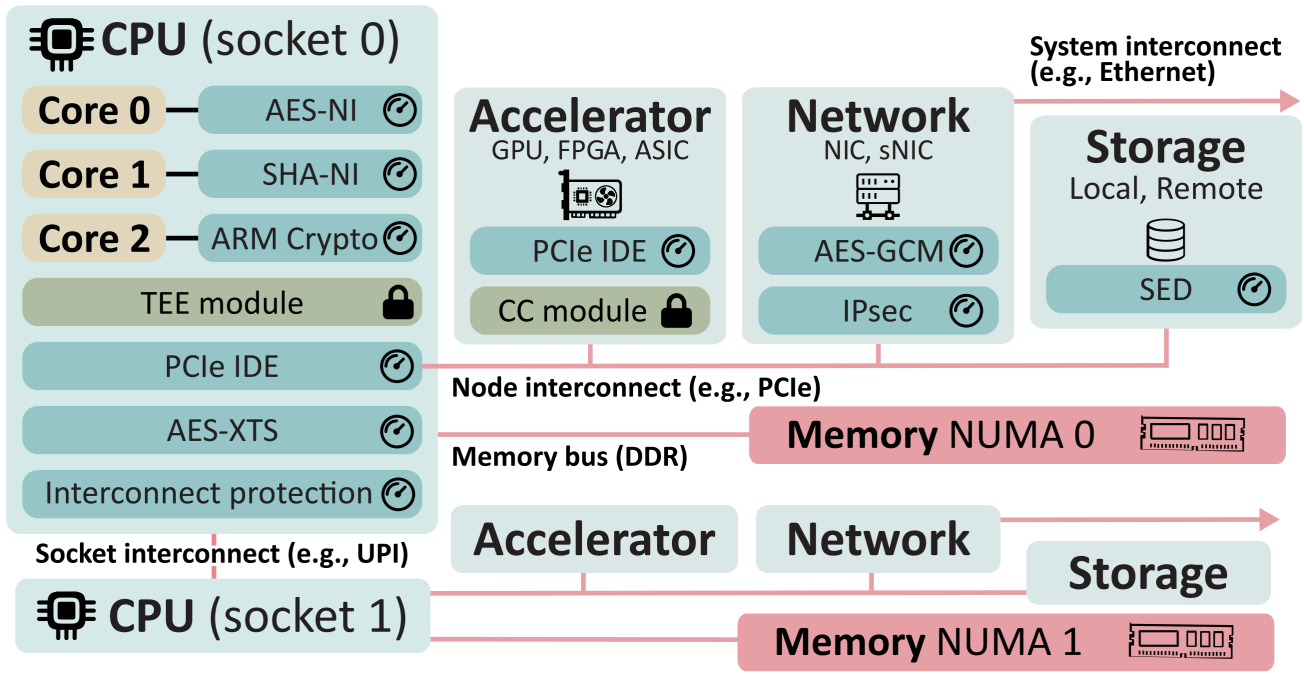
What are the overheads of these systems?

cHPC needs to secure not only compute

How do tradeoffs for these areas look like?

Applications have different security needs

How to characterize them?

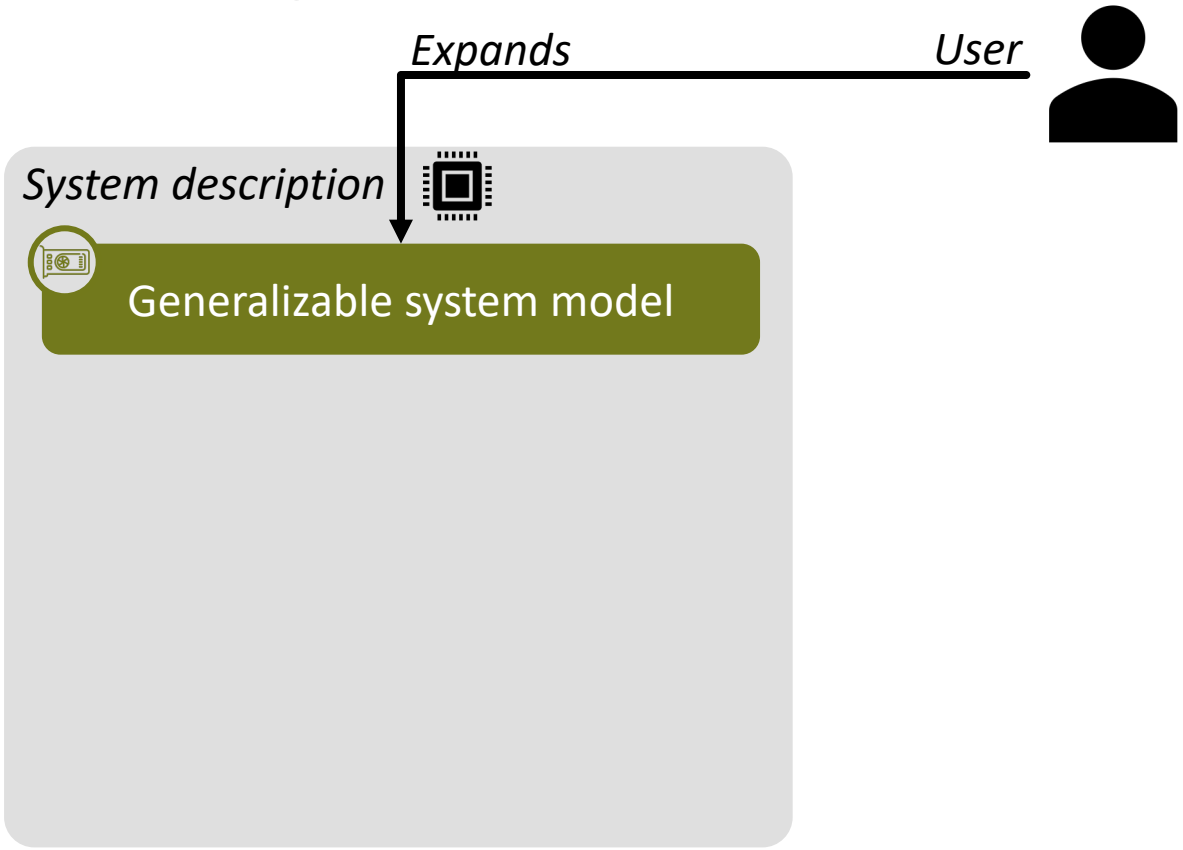


No platform characterizing clearly security and its performance costs

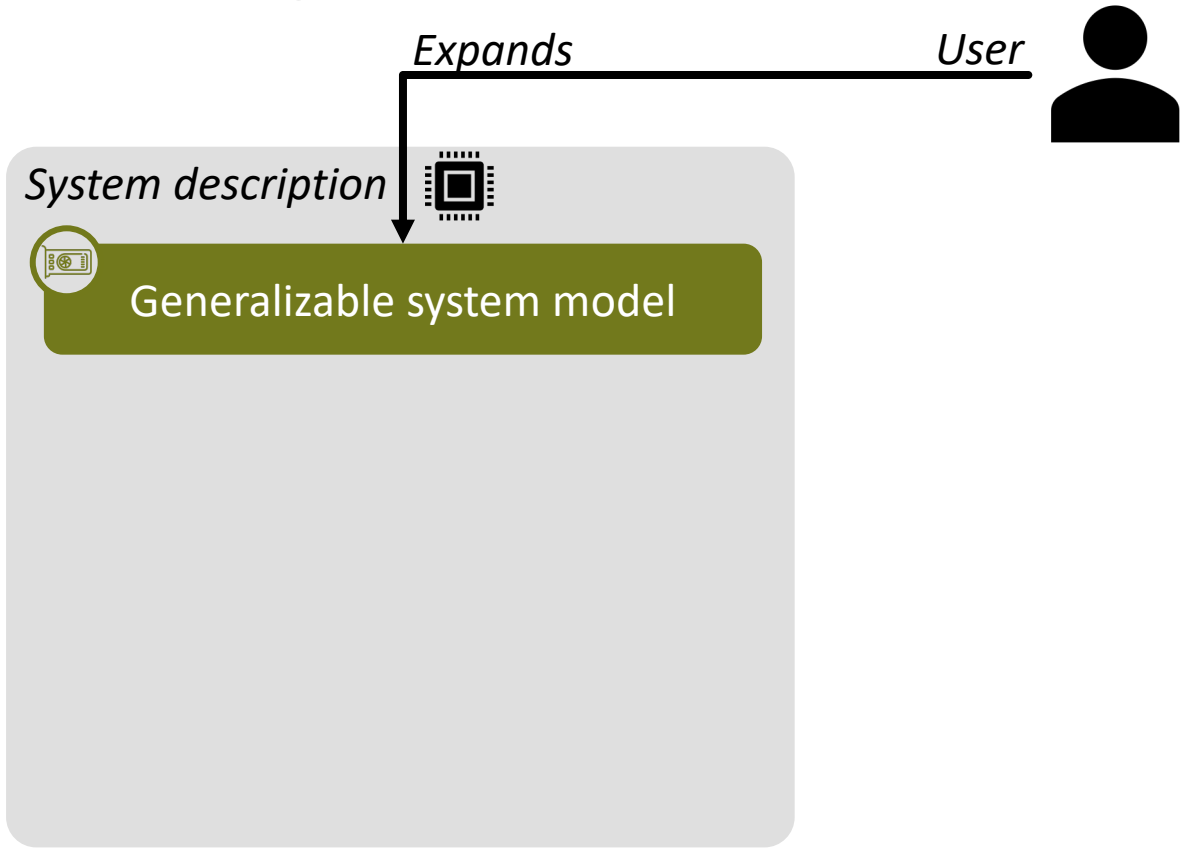
SecPerf design



SecPerf design

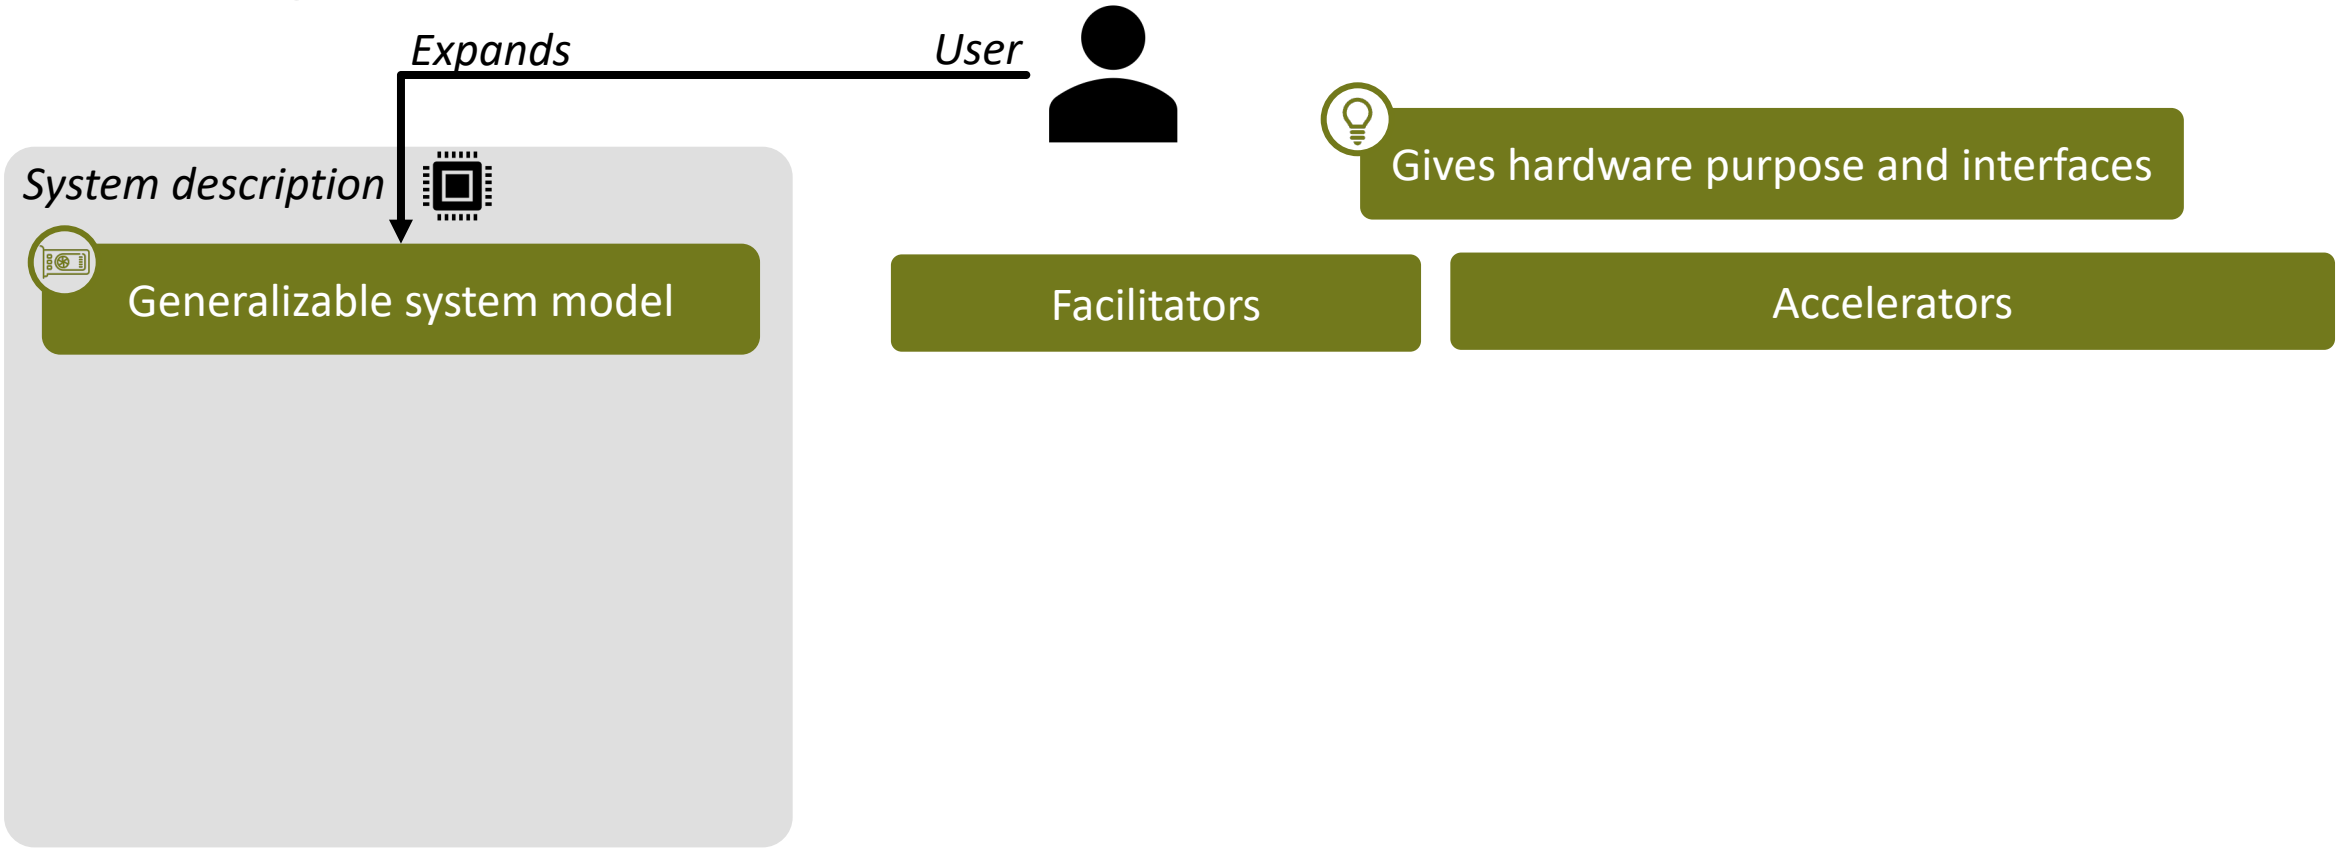


SecPerf design



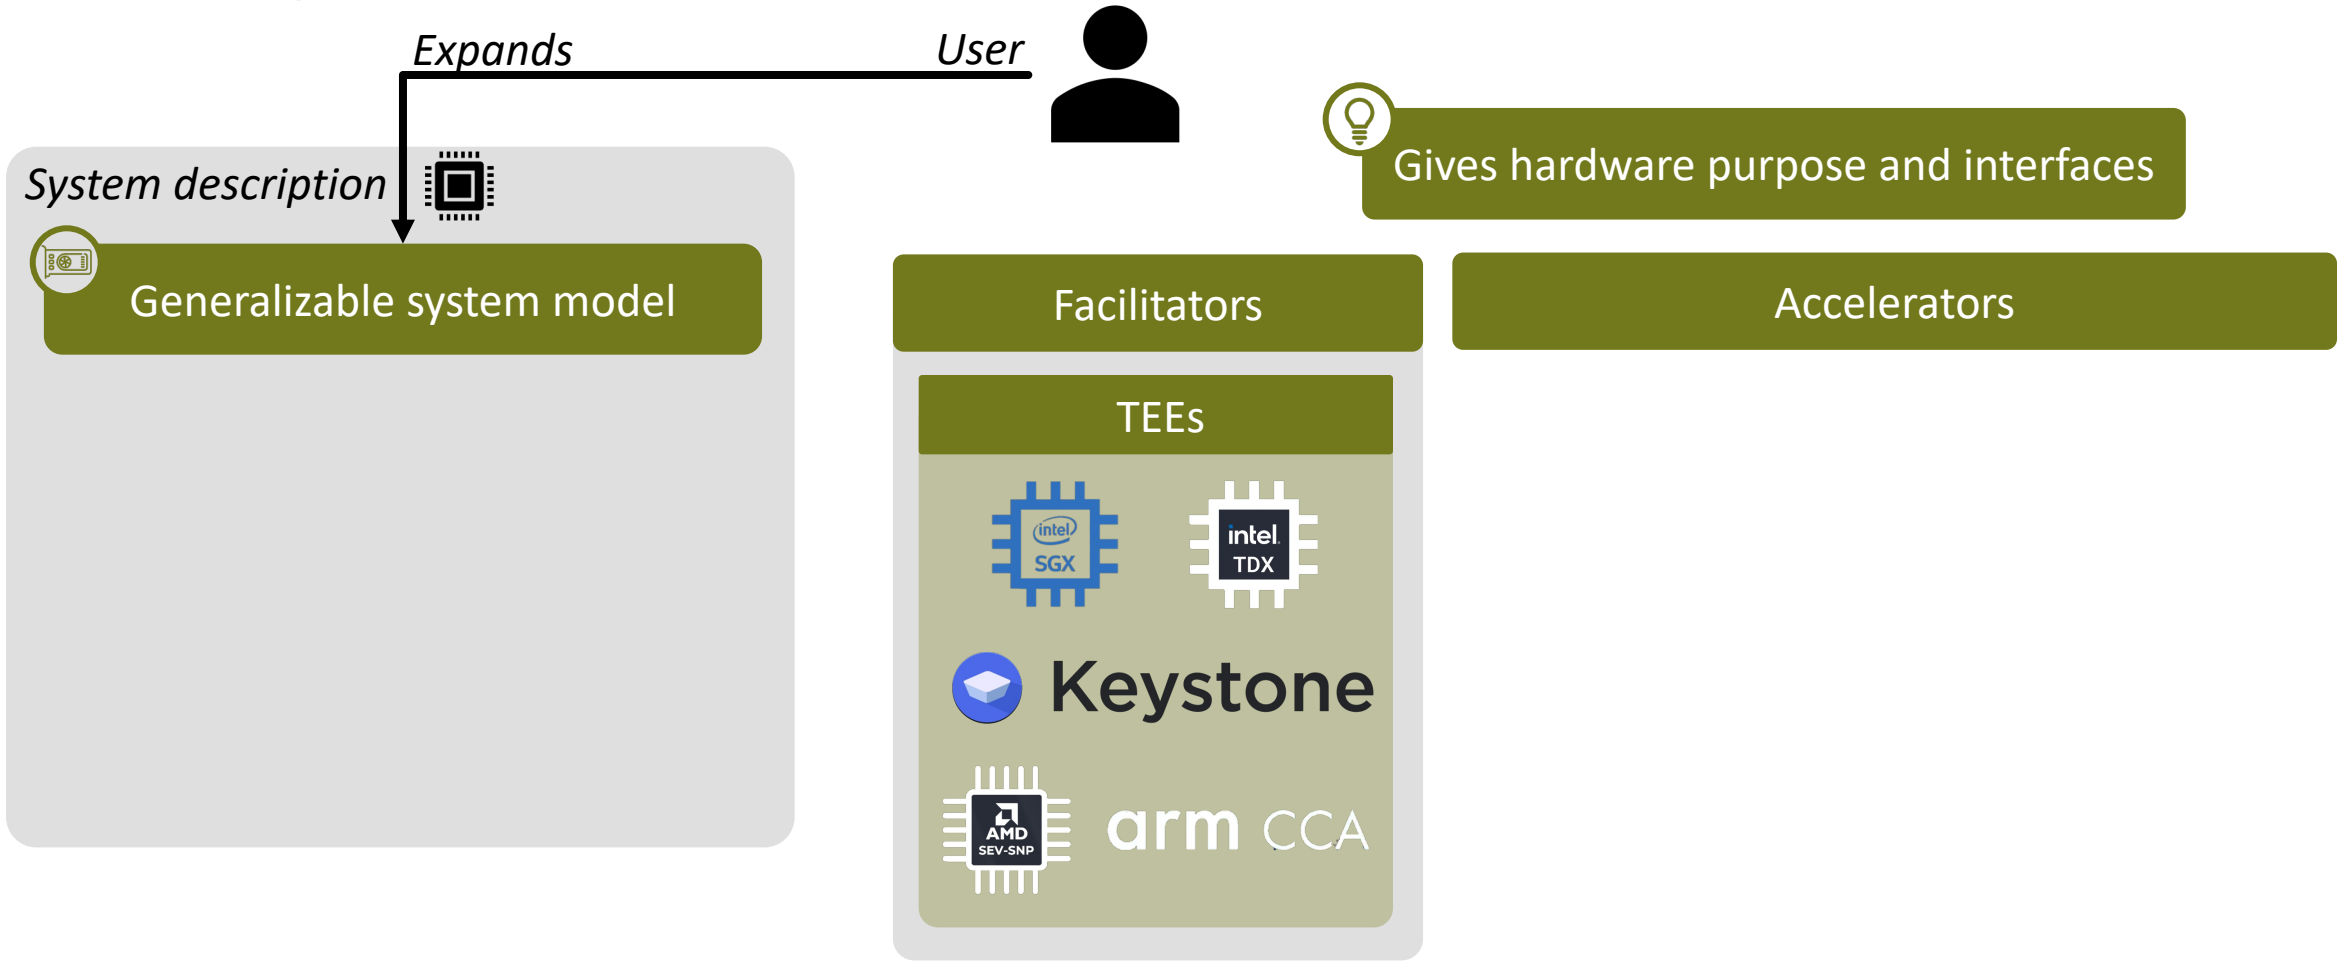
Gives hardware purpose and interfaces

SecPerf design



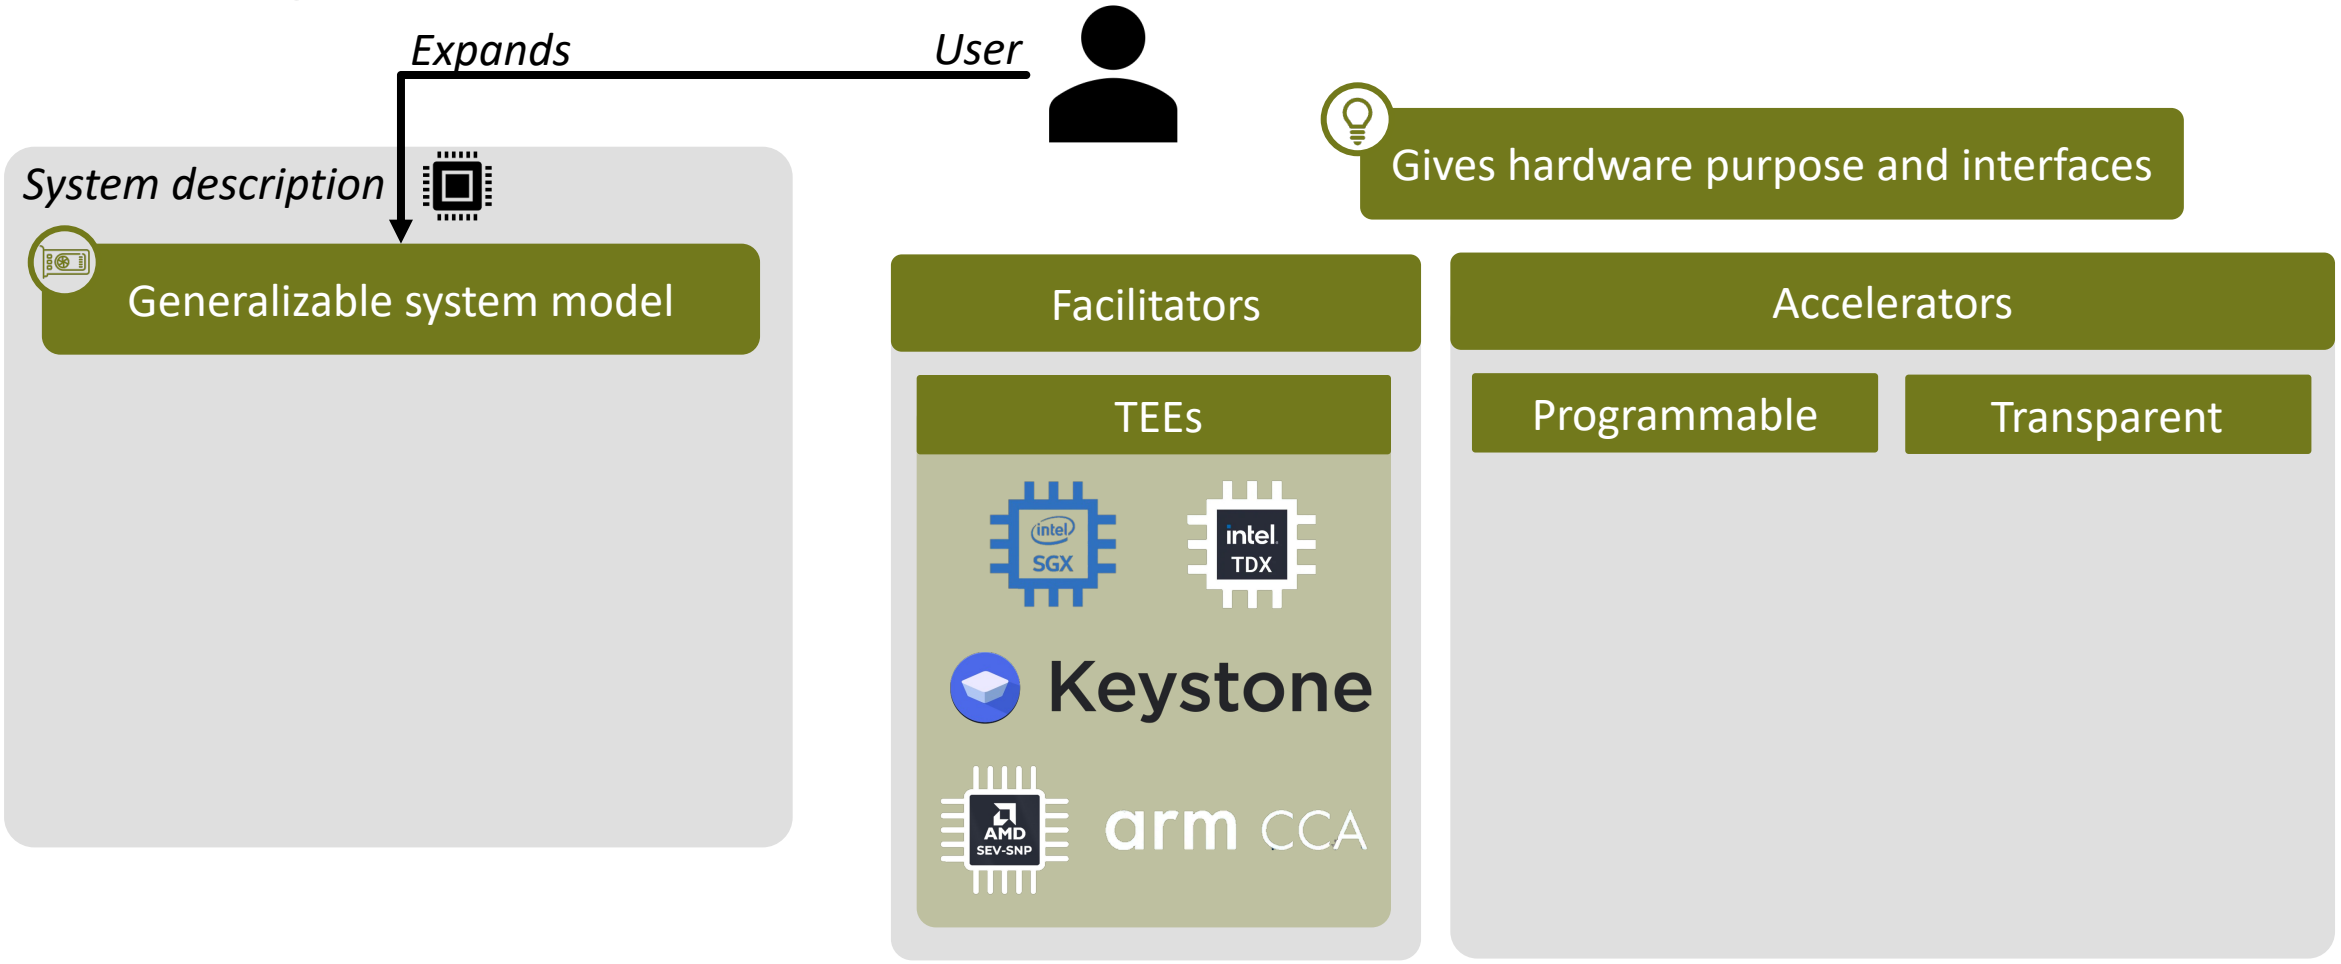
Facilitators provide security while accelerators deal with data-path performance

SecPerf design



Facilitators provide security while accelerators deal with data-path performance

SecPerf design



SecPerf design

Expands *User*

System description

Generalizable system model

Gives hardware purpose and interfaces

Facilitators

TEEs

Keystone

arm CCA

Accelerators

Programmable

AES-NI

arm crypto

GGAs

XILINX bitstream

Transparent

Facilitators provide security while accelerators deal with data-path performance

SecPerf design

Expands — *User*

System description

Generalizable system model

Gives hardware purpose and interfaces

Facilitators

TEEs

Intel SGX

Intel TDX

Keystone

AMD SEV-SNP

arm CCA

Accelerators

Programmable

Intel AES-NI

arm crypto

BlueField-2 DPU — GGAs

XILINX bitstream

Transparent

PCI EXPRESS IDE

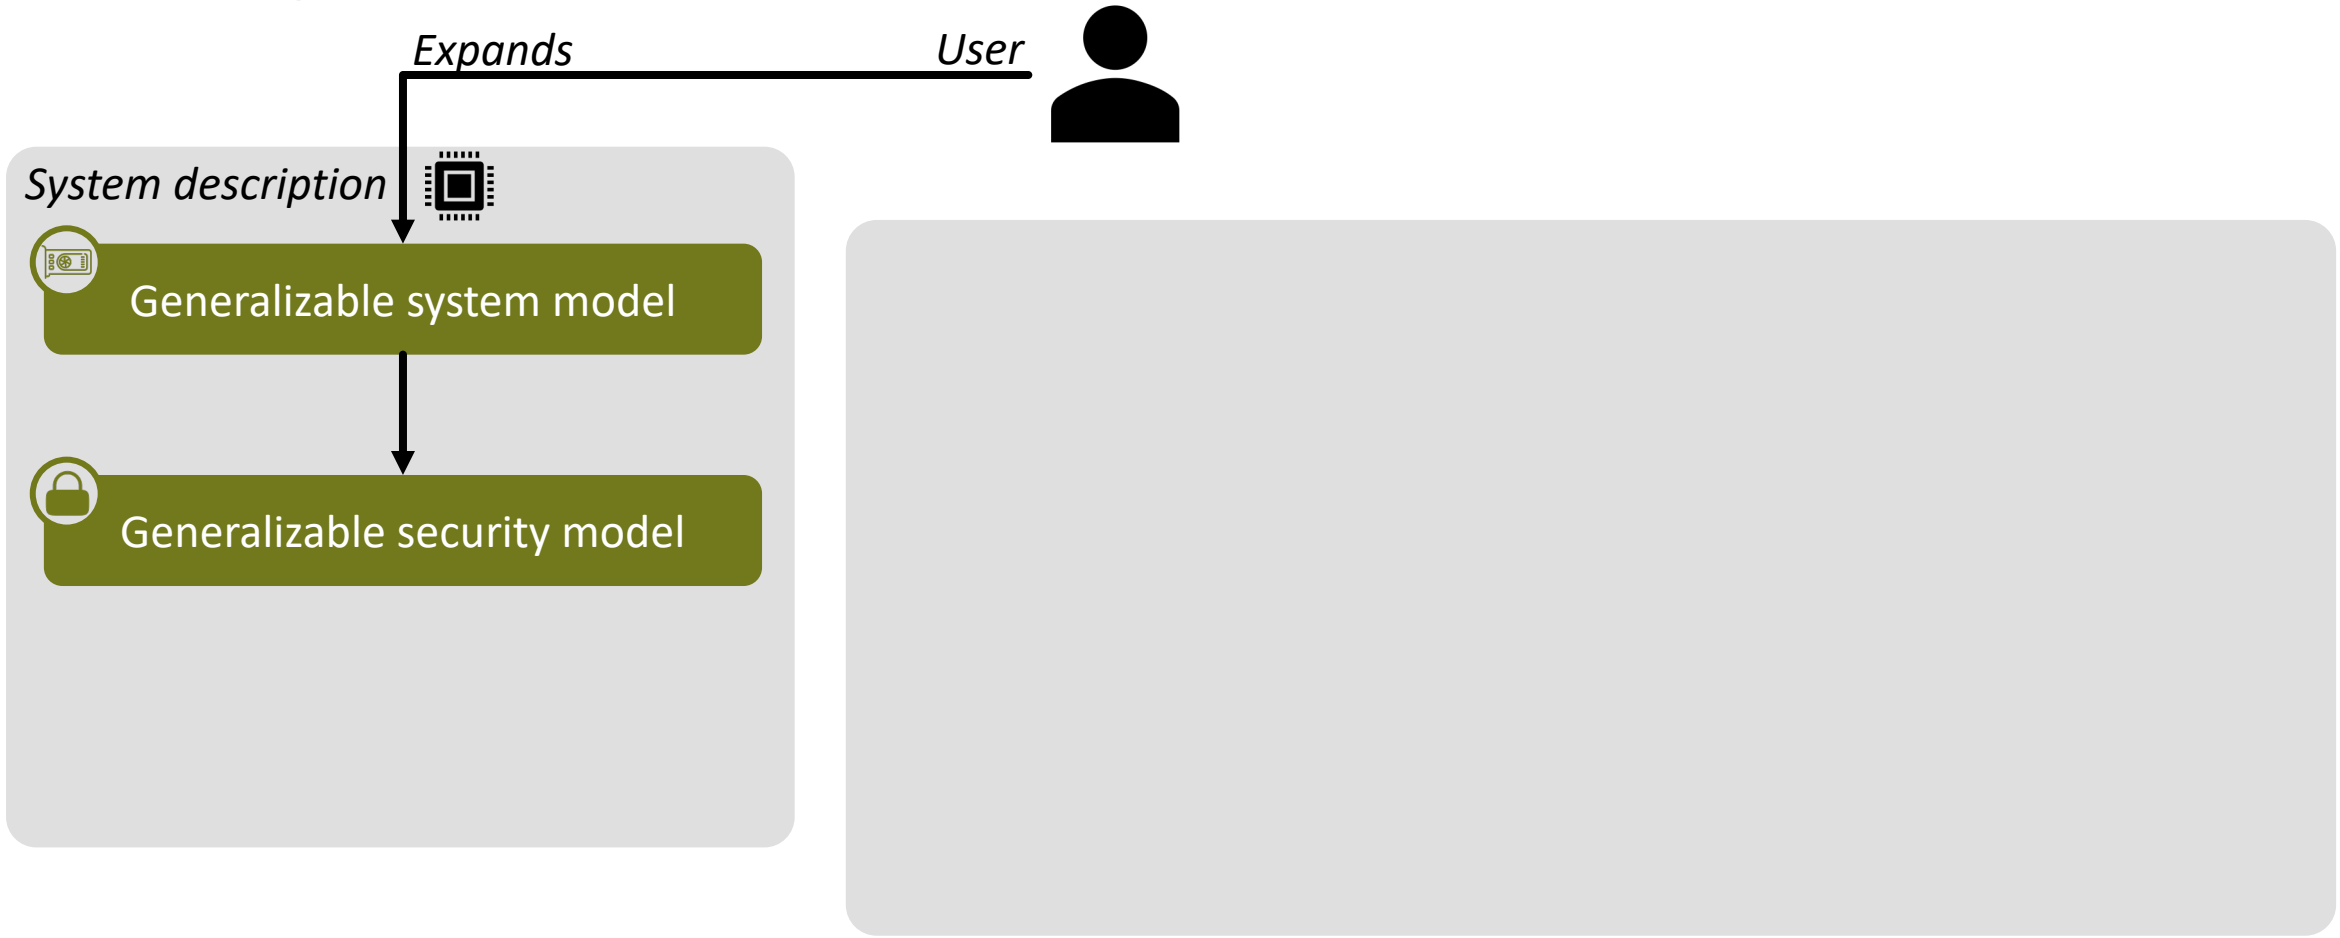
Intel MKTME

NUMA encryption

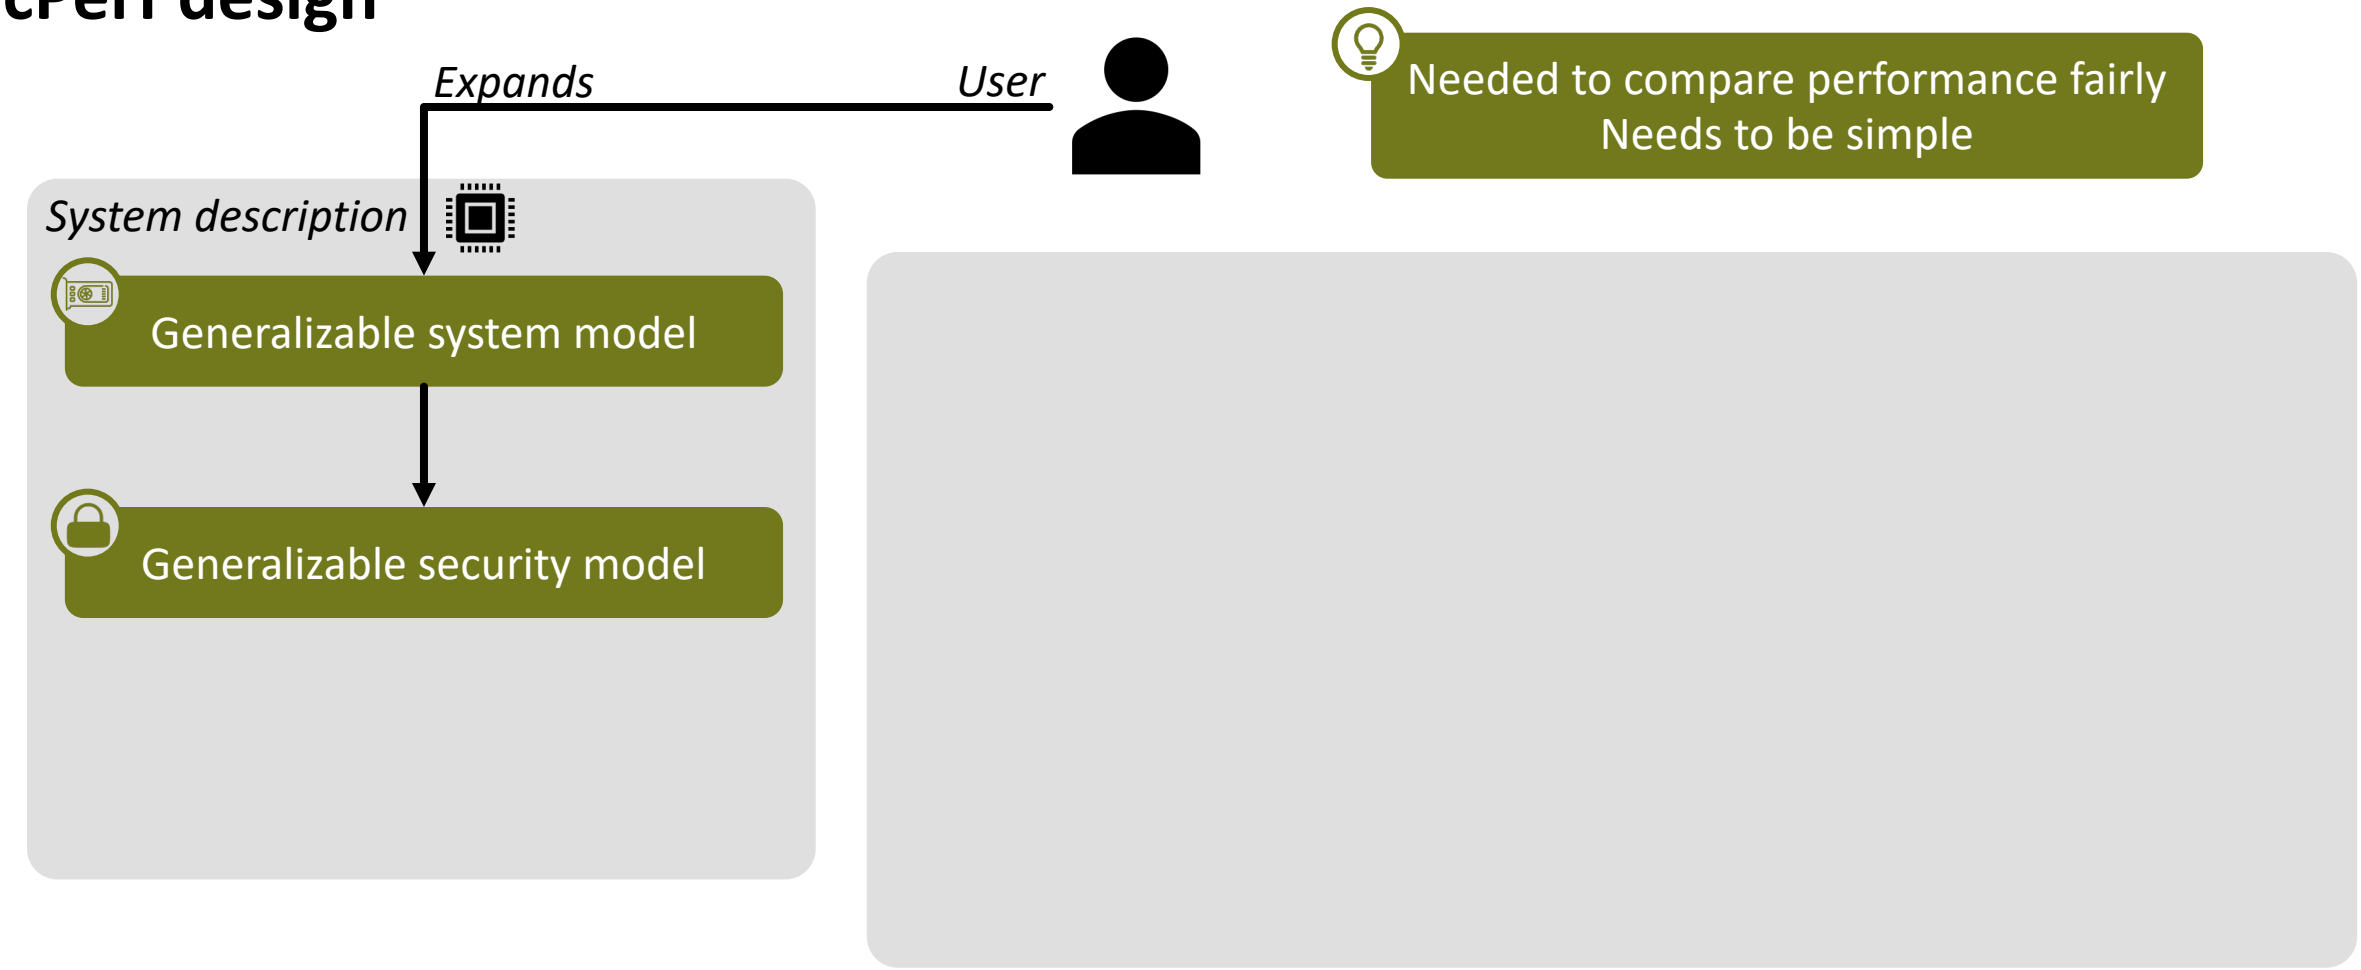
SED offload

Facilitators provide security while accelerators deal with data-path performance

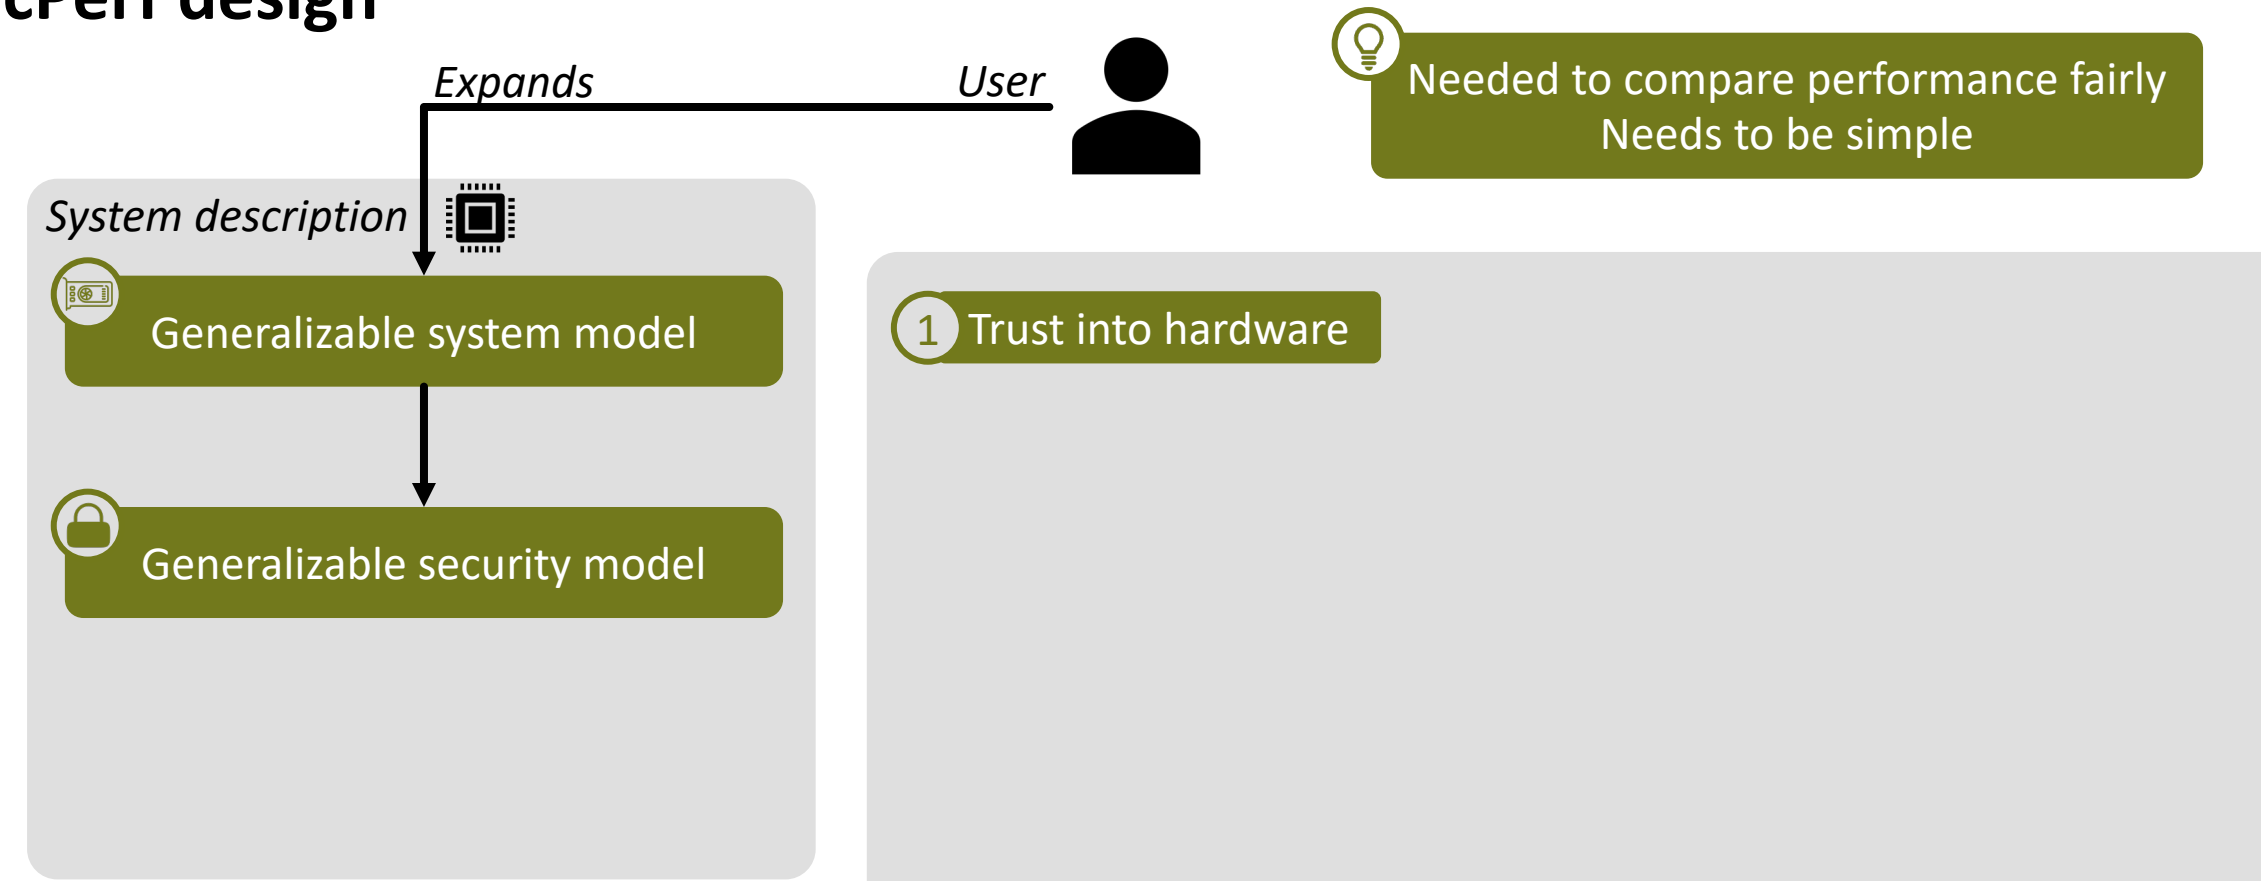
SecPerf design



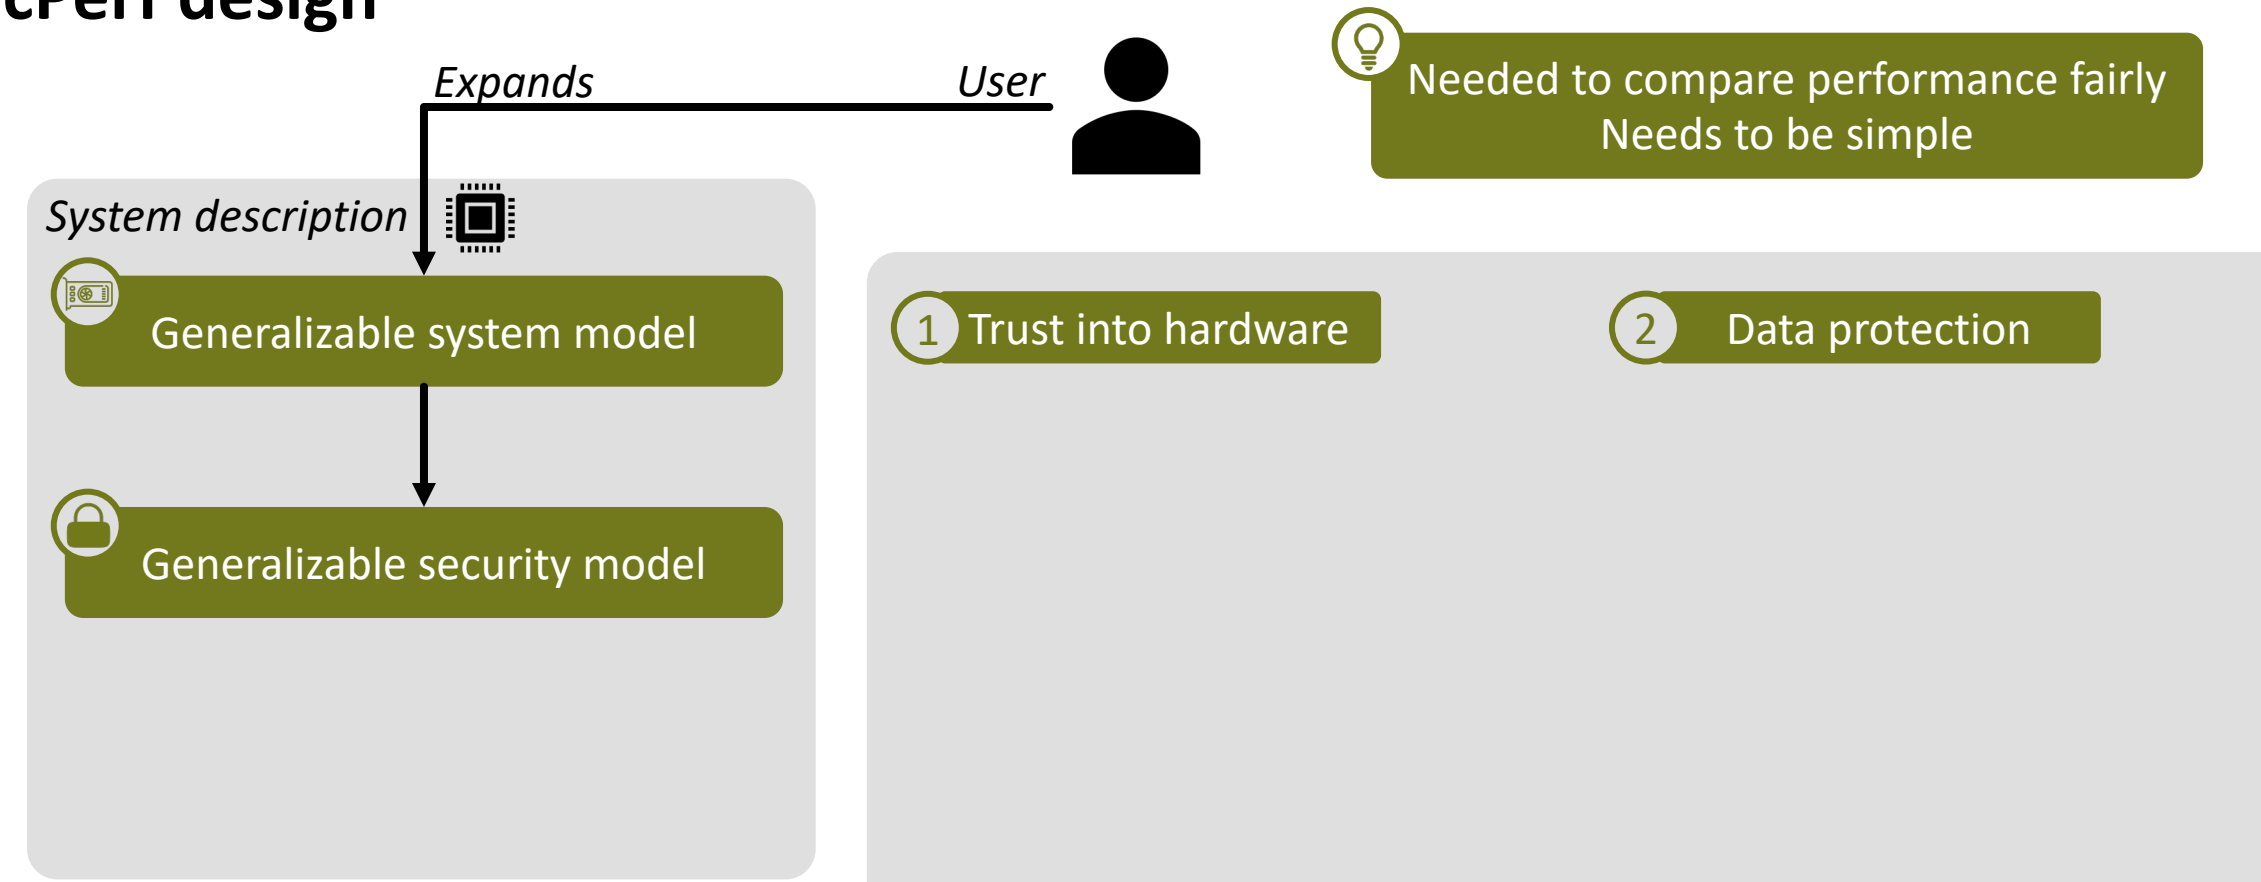
SecPerf design



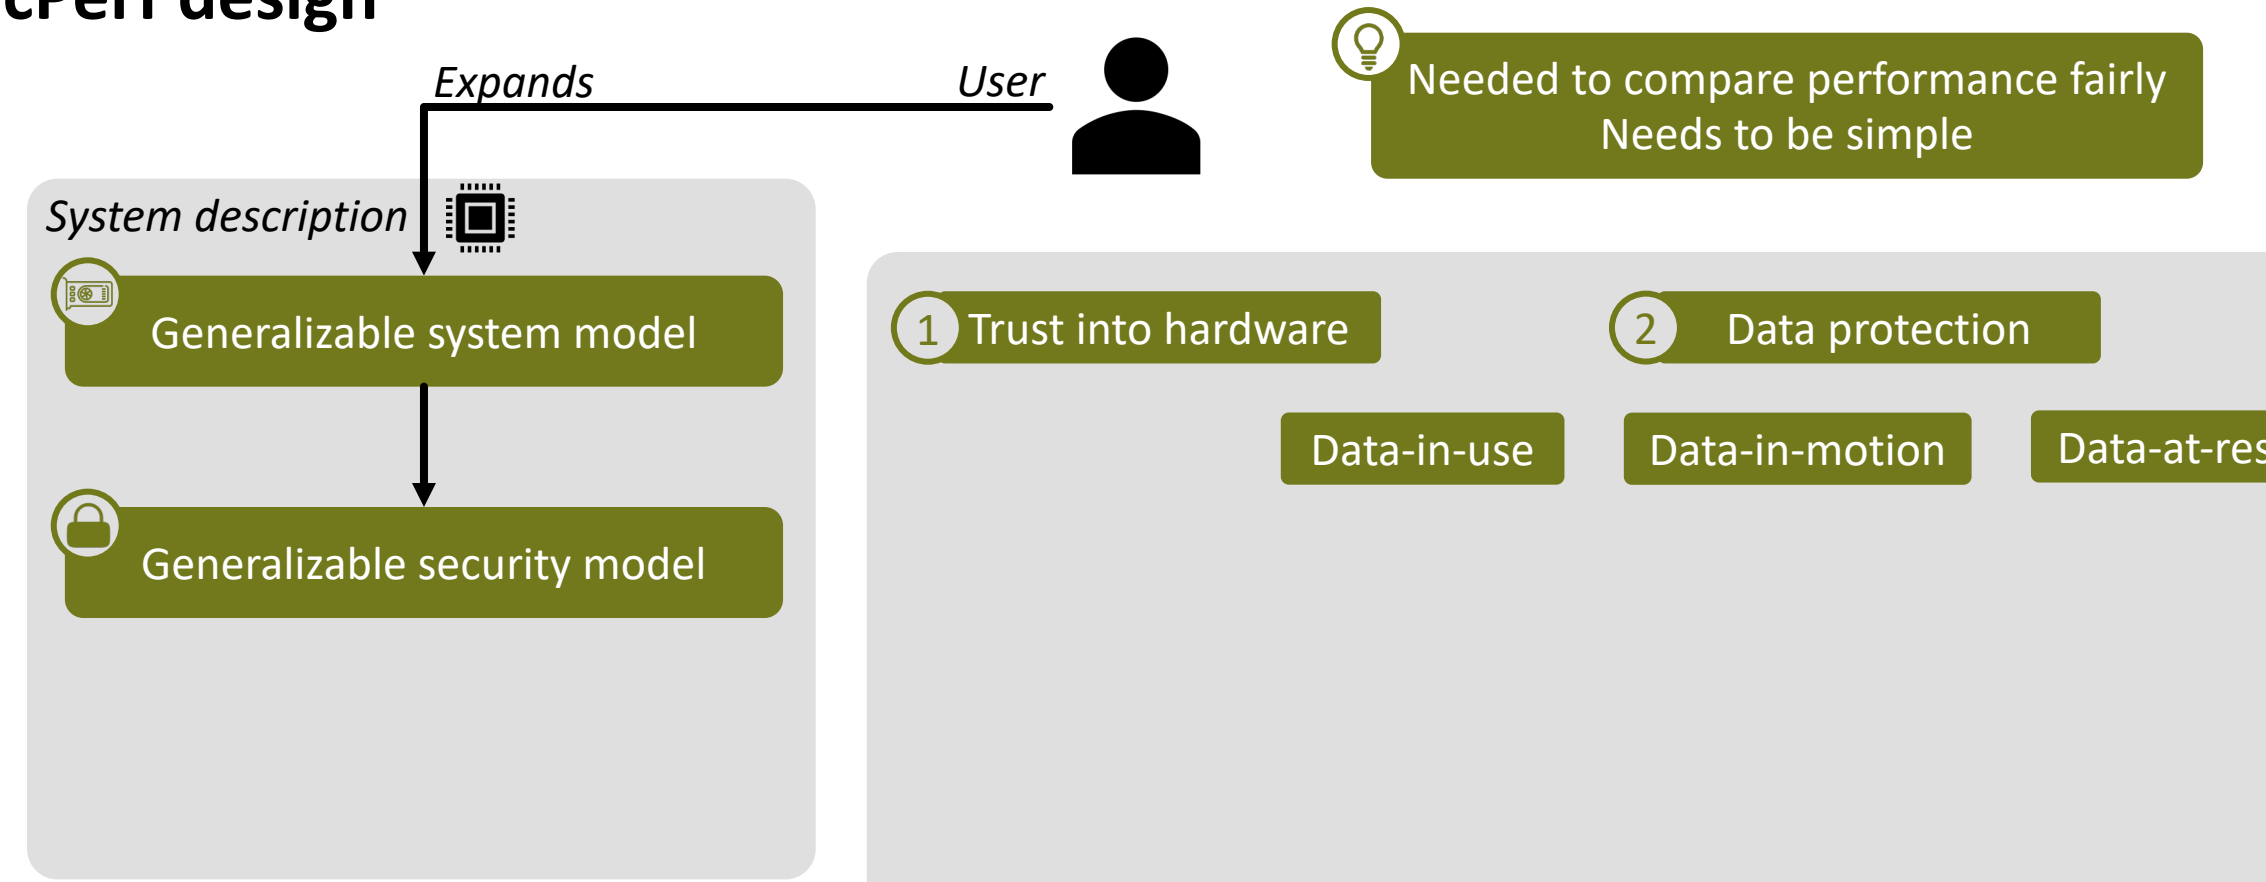
SecPerf design



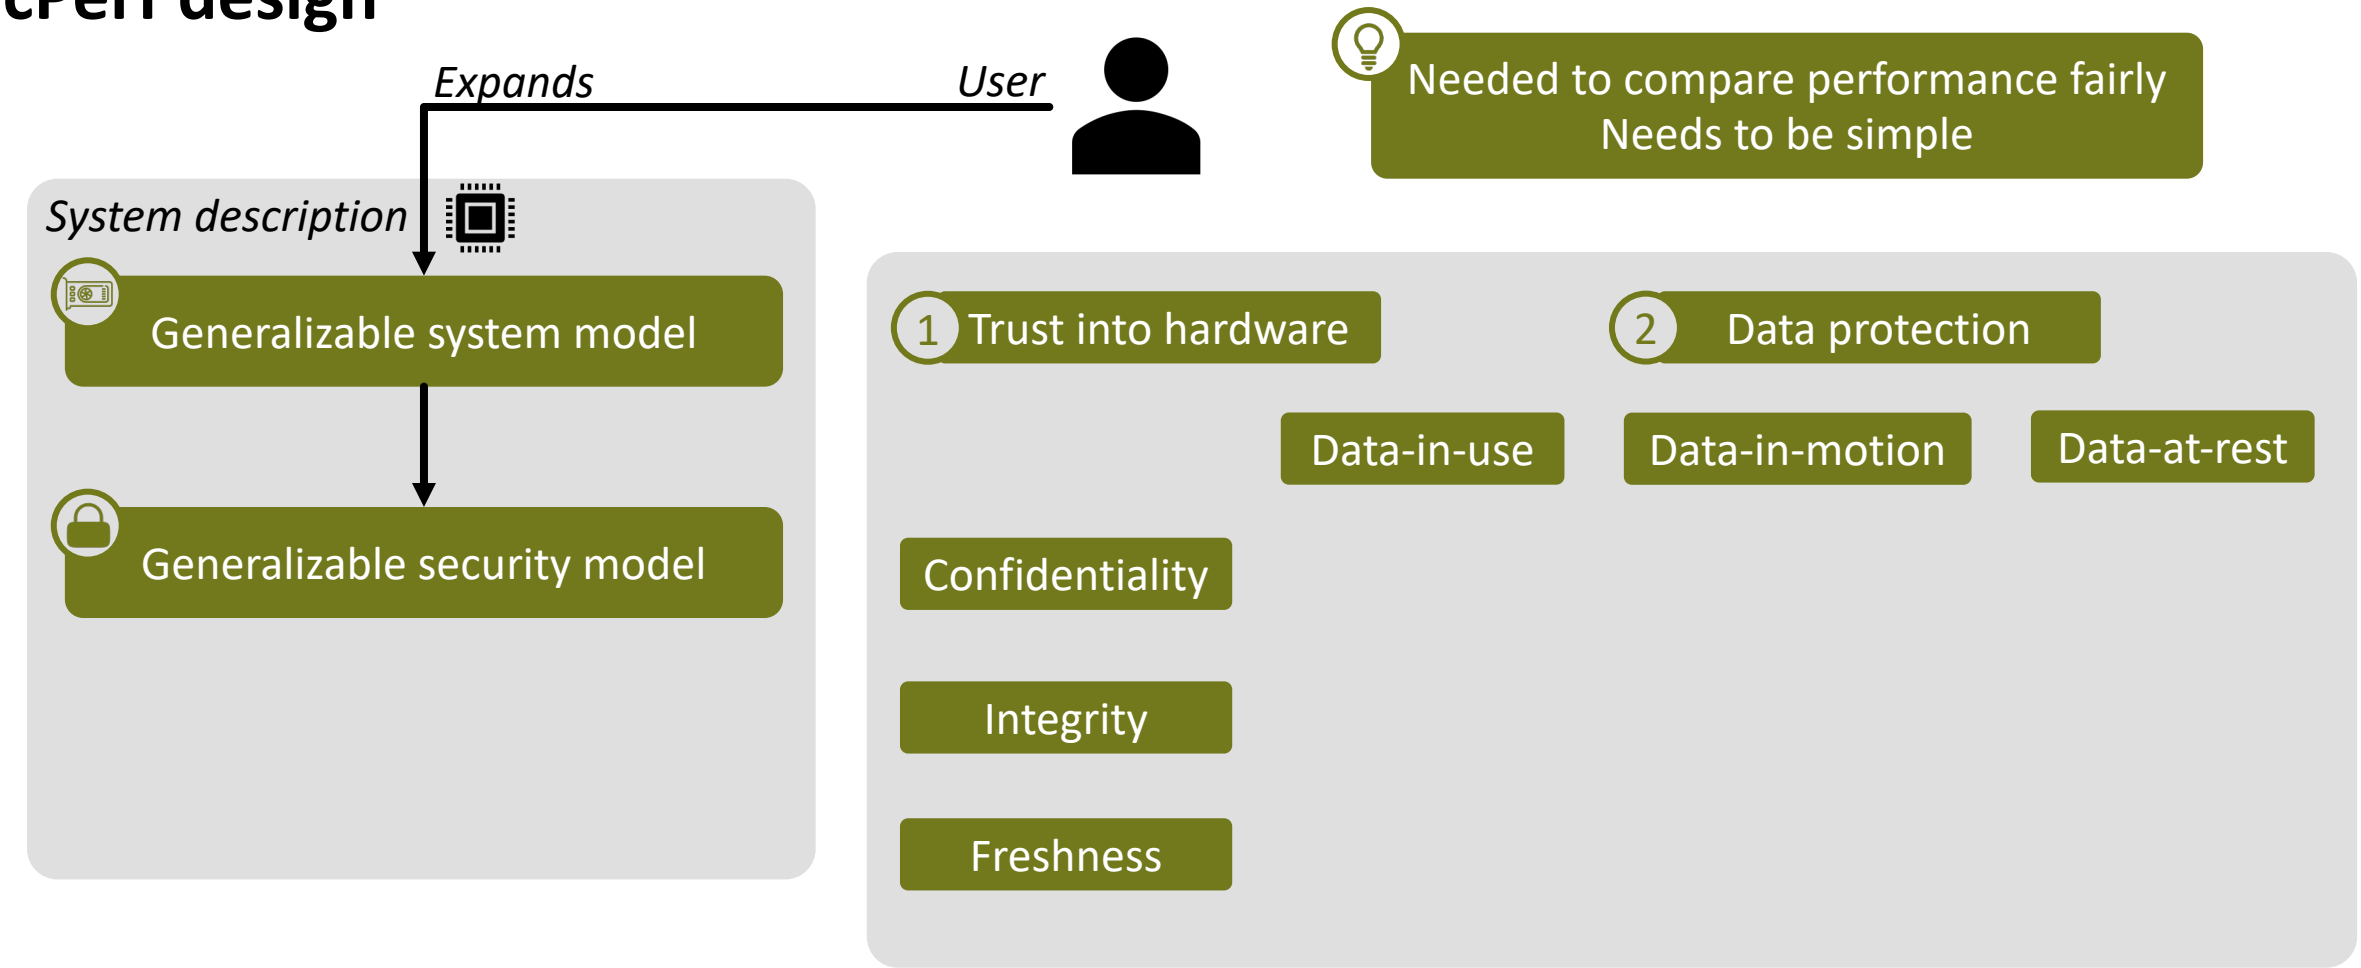
SecPerf design



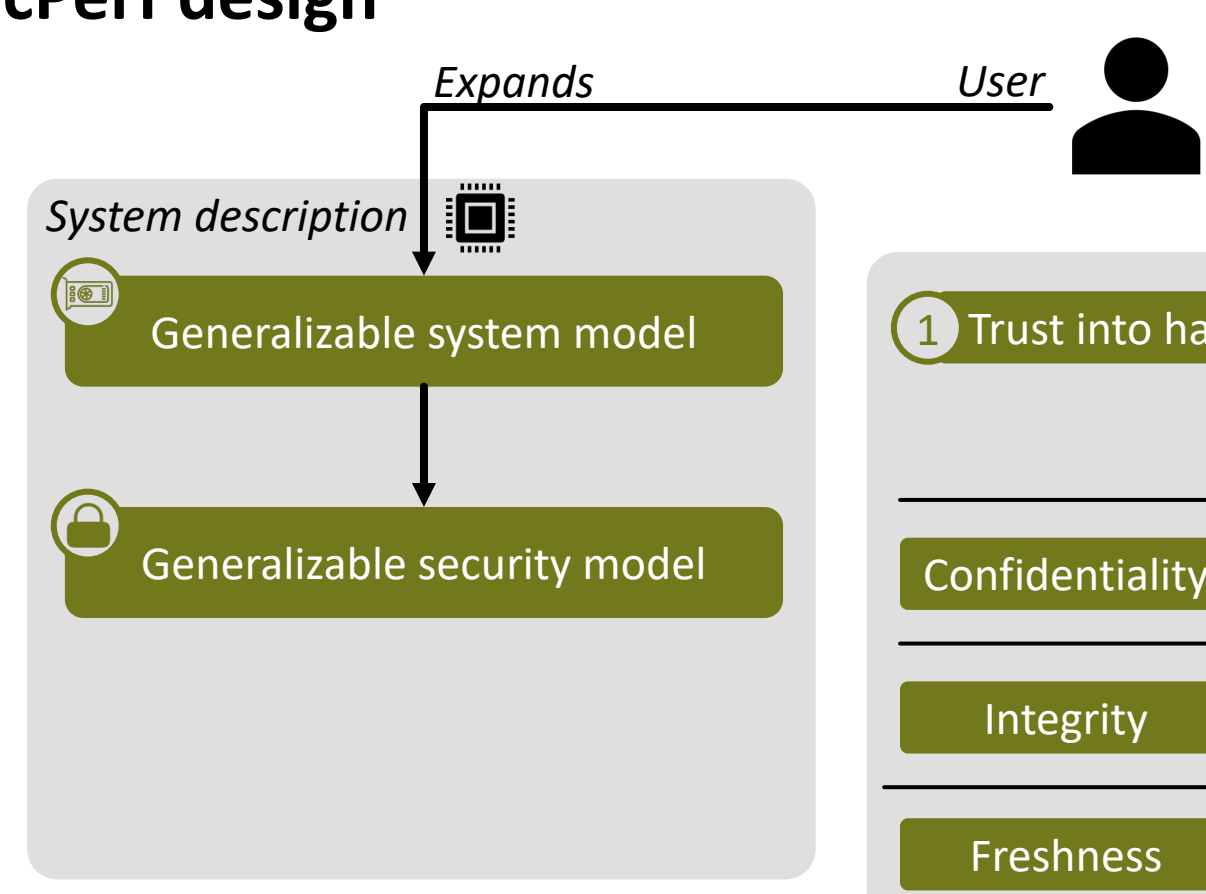
SecPerf design



SecPerf design



SecPerf design



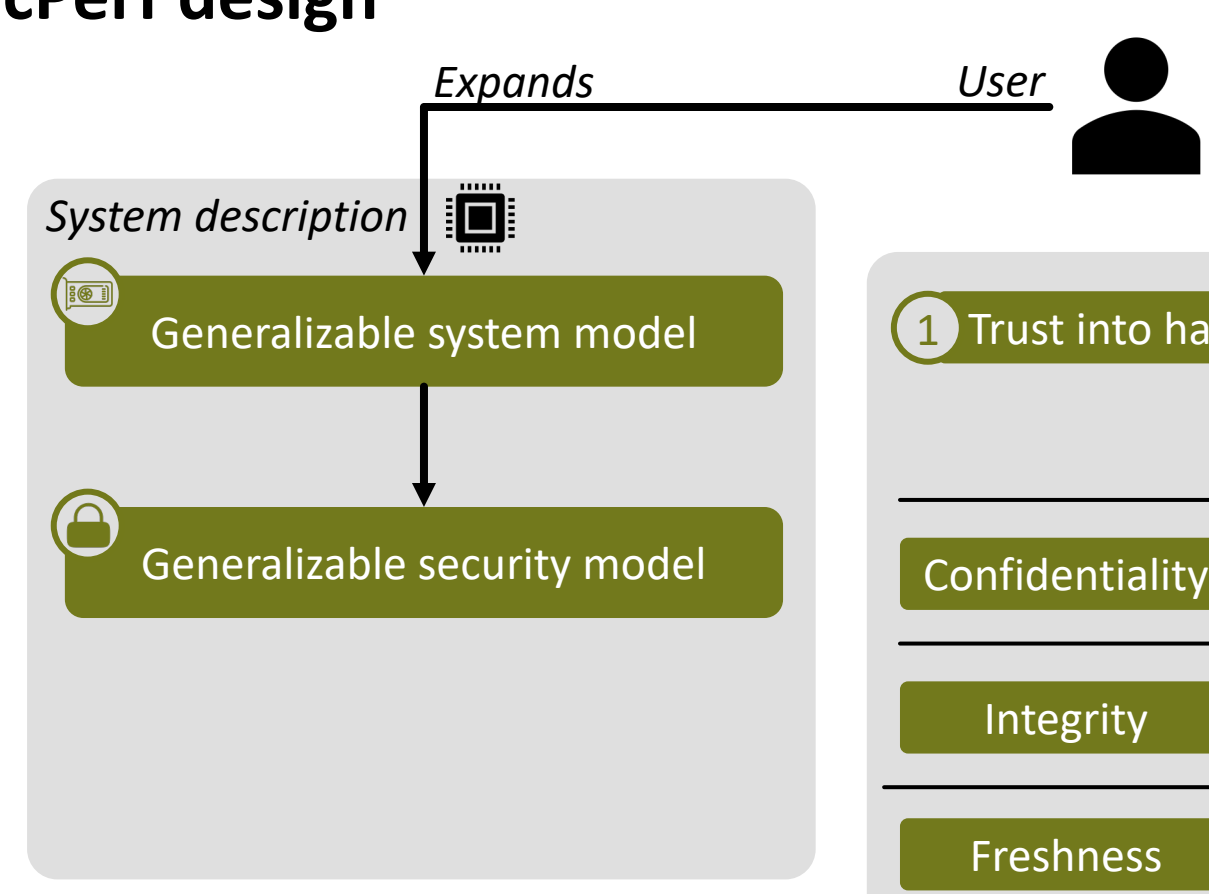
Needed to compare performance fairly
Needs to be simple

1 Trust into hardware

2 Data protection

	Data-in-use	Data-in-motion	Data-at-rest
Confidentiality			
Integrity			
Freshness			

SecPerf design



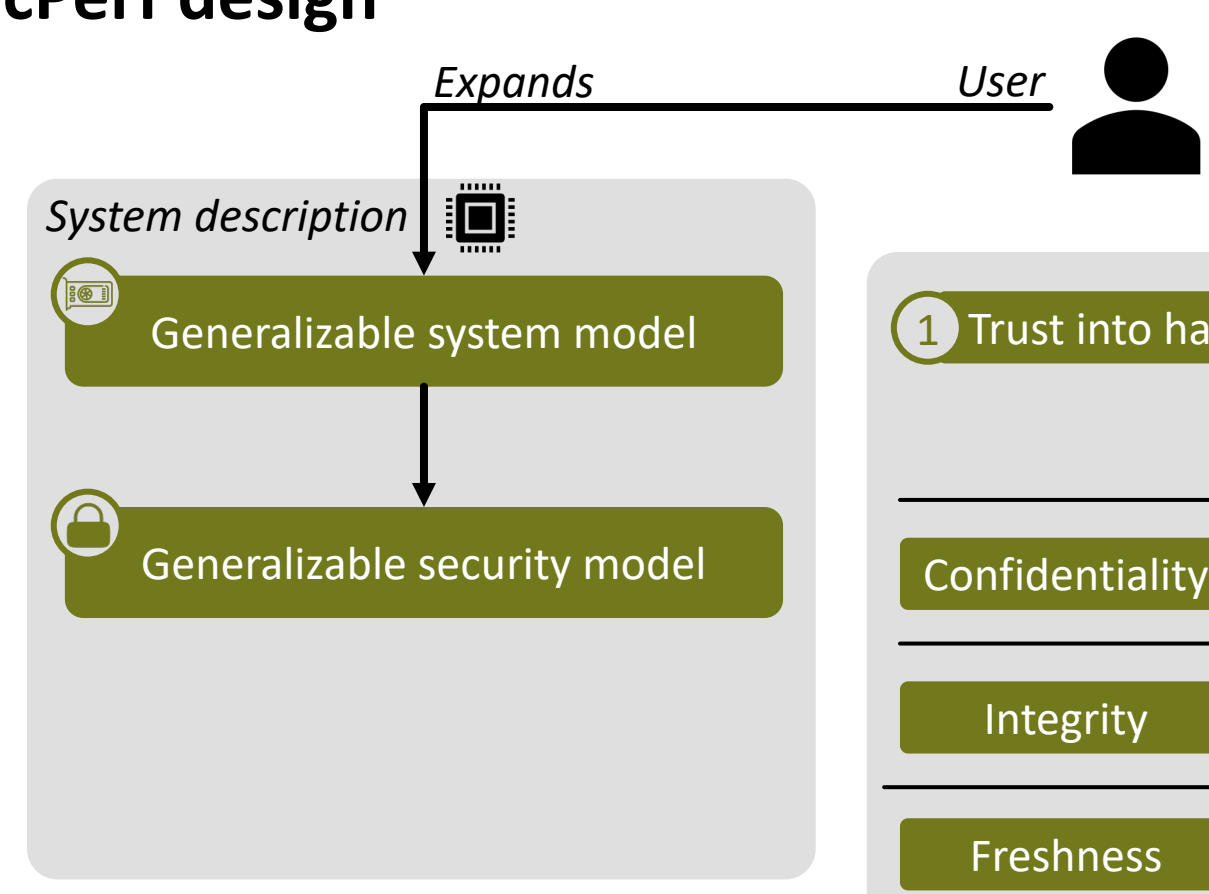
Needed to compare performance fairly
Needs to be simple

1 Trust into hardware

2 Data protection

	Data-in-use	Data-in-motion	Data-at-rest
Confidentiality	Intel TDX *AES-XTS*		
Integrity			
Freshness			

SecPerf design



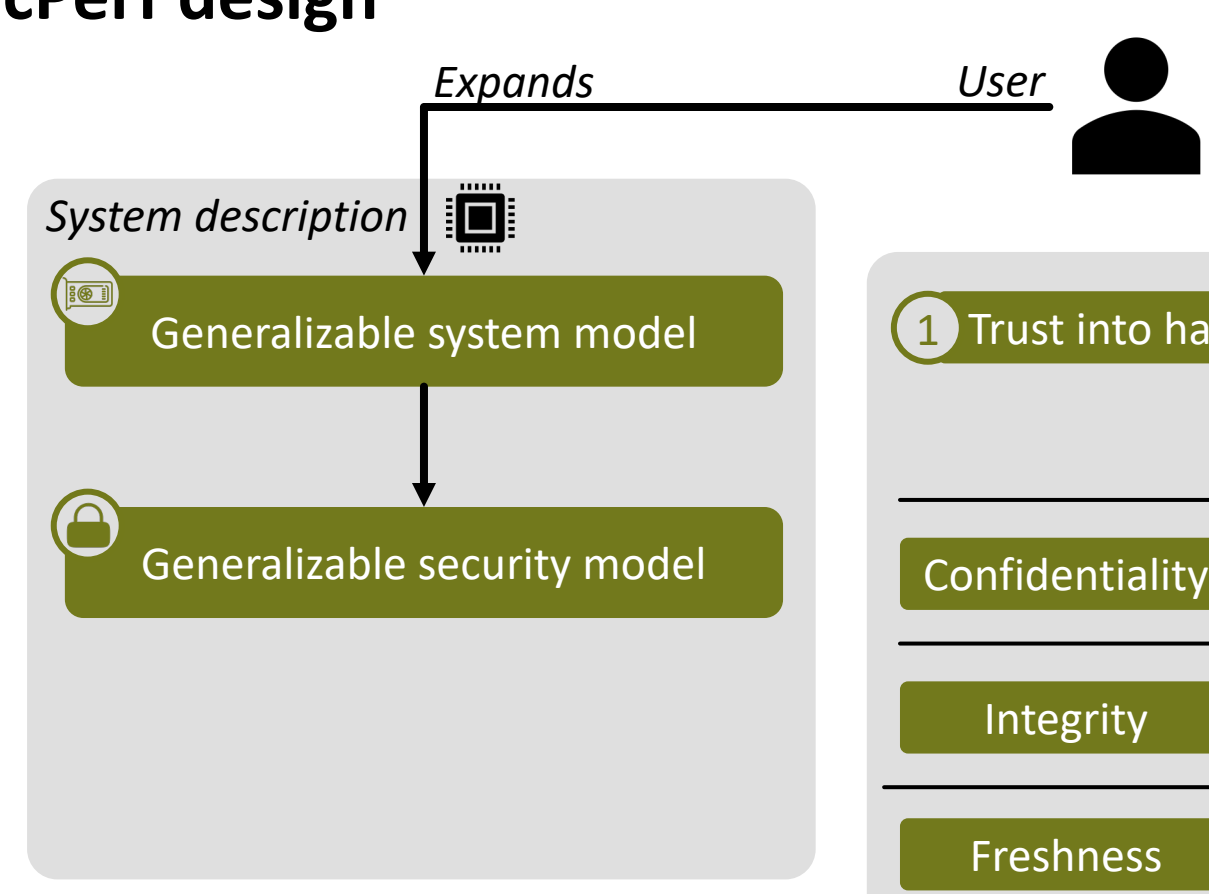
Needed to compare performance fairly
Needs to be simple

1 Trust into hardware

2 Data protection

	Data-in-use	Data-in-motion	Data-at-rest
Confidentiality	Intel TDX *AES-XTS*		
Integrity		Wireguard *Poly1305*	
Freshness			

SecPerf design



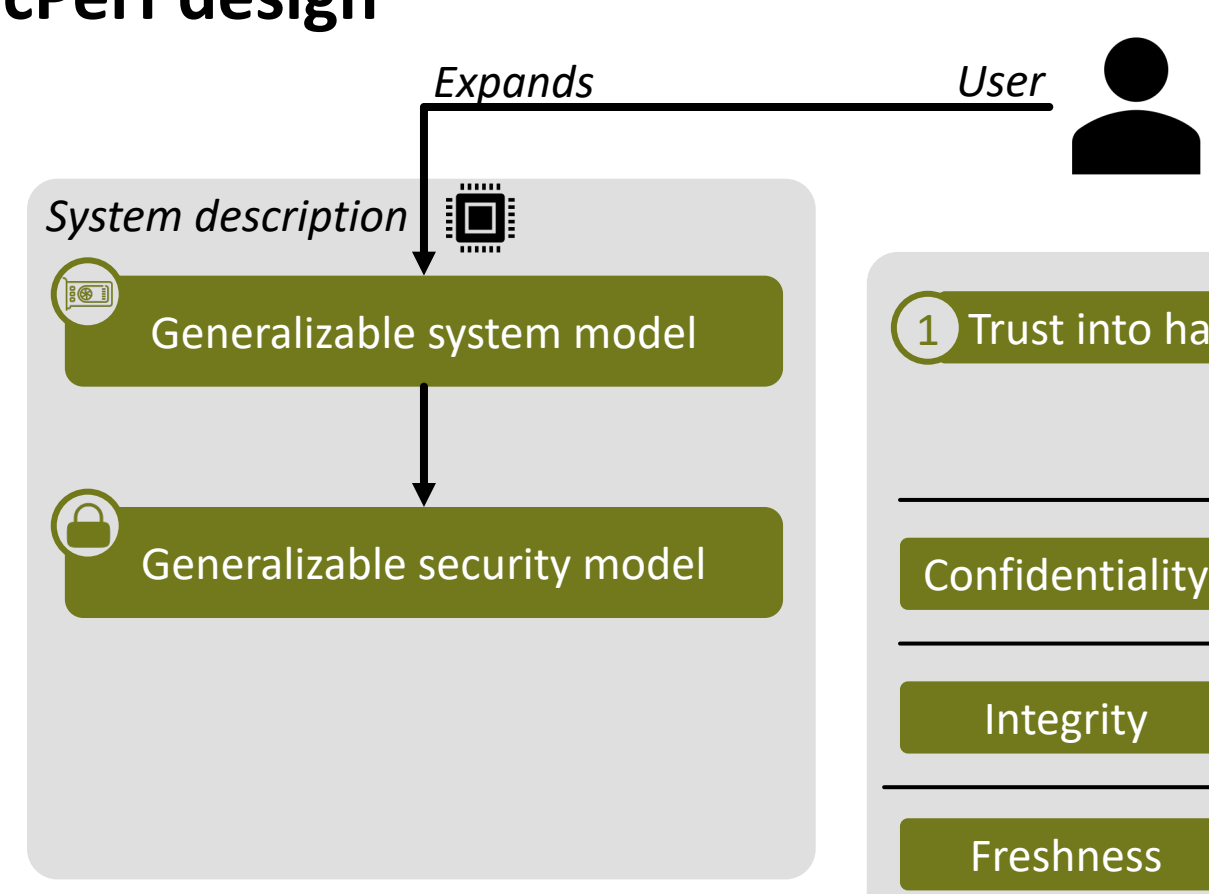
Needed to compare performance fairly
Needs to be simple

1 Trust into hardware

2 Data protection

	Data-in-use	Data-in-motion	Data-at-rest
Confidentiality	Intel TDX *AES-XTS*		
Integrity		Wireguard *Poly1305*	
Freshness			dm-x *Merkle tree*

SecPerf design



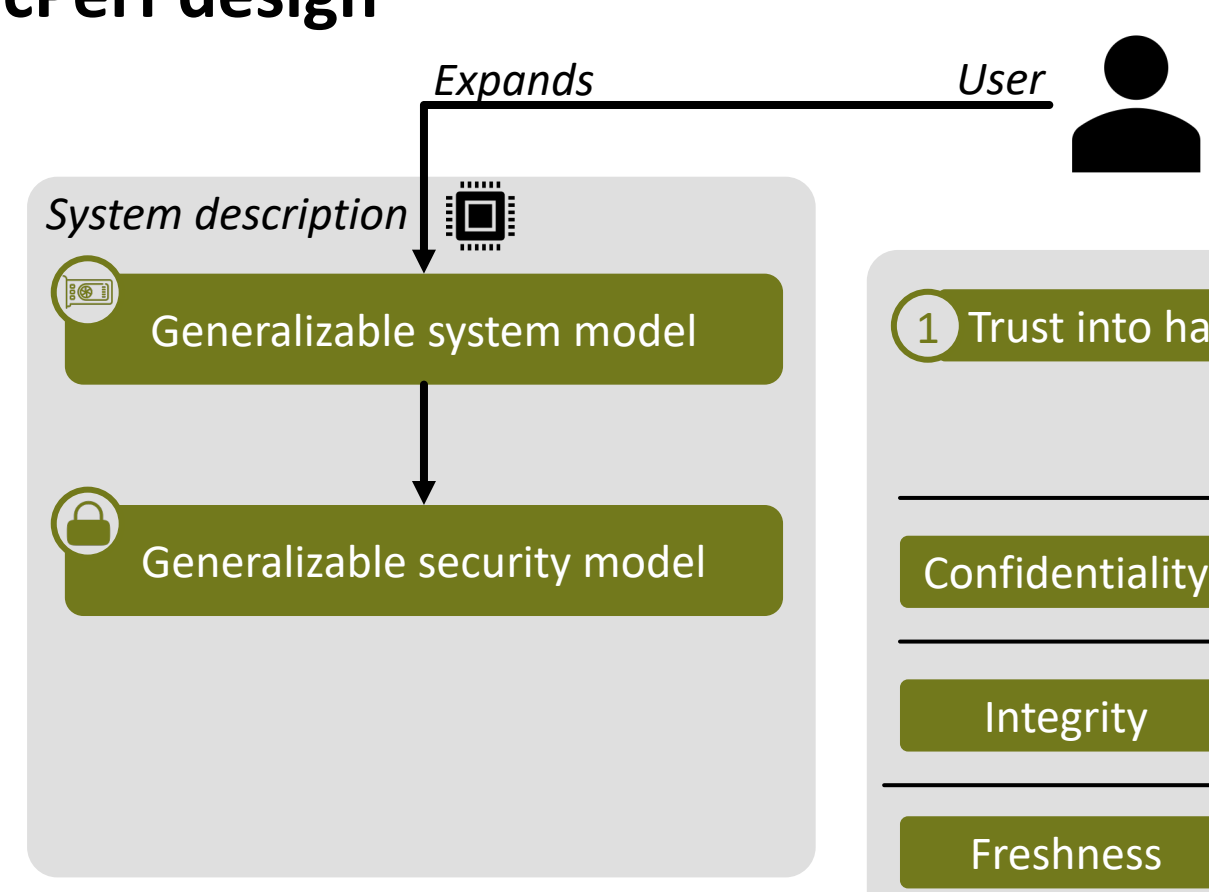
Needed to compare performance fairly
Needs to be simple

1 Trust into hardware

2 Data protection

	Data-in-use	Data-in-motion	Data-at-rest
Confidentiality	Intel TDX *AES-XTS*	TLS *ChaCha20*	dm-crypt *AES-XTS*
Integrity	TDX ECC fields *SHA3*	Wireguard *Poly1305*	dm-integrity *AES-GCM*
Freshness	SGX v1 *Merkle tree*	IPsec *PSN window*	dm-x *Merkle tree*

SecPerf design



Needed to compare performance fairly
Needs to be simple

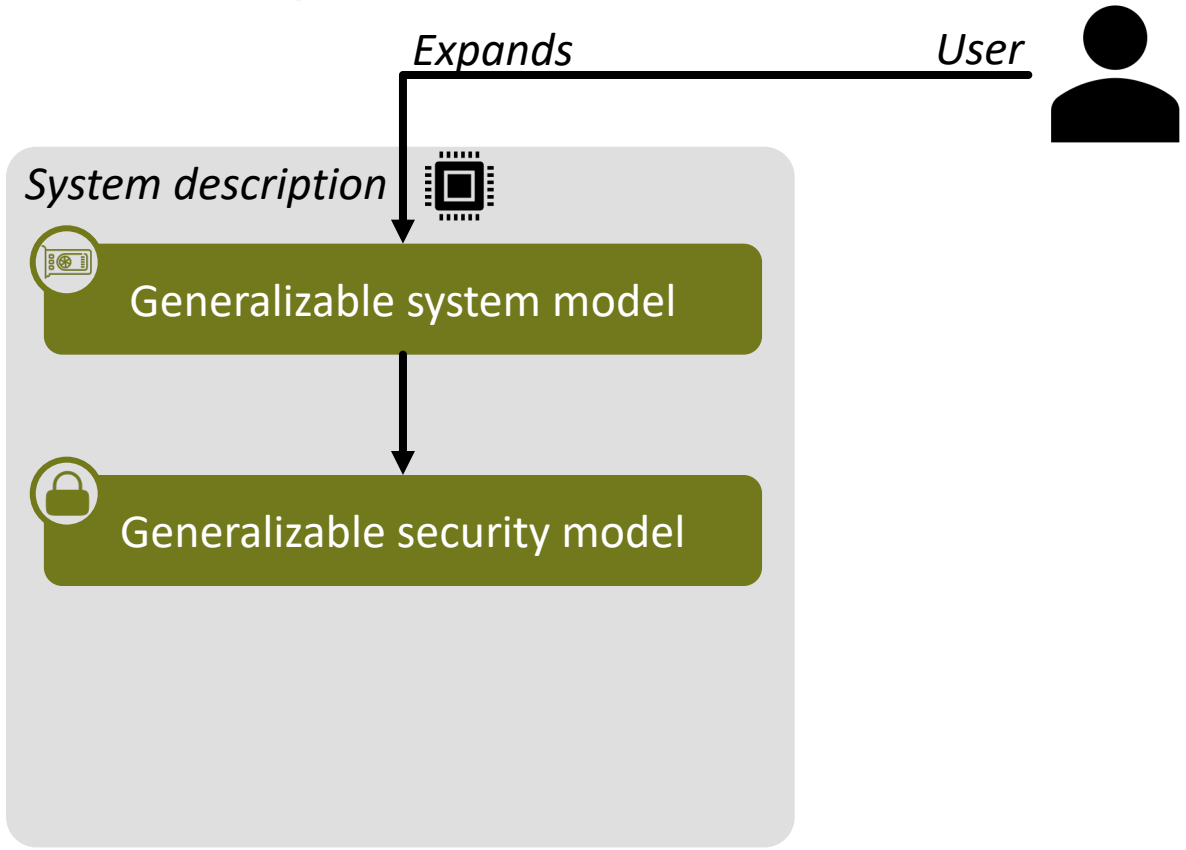
1 Trust into hardware

2 Data protection

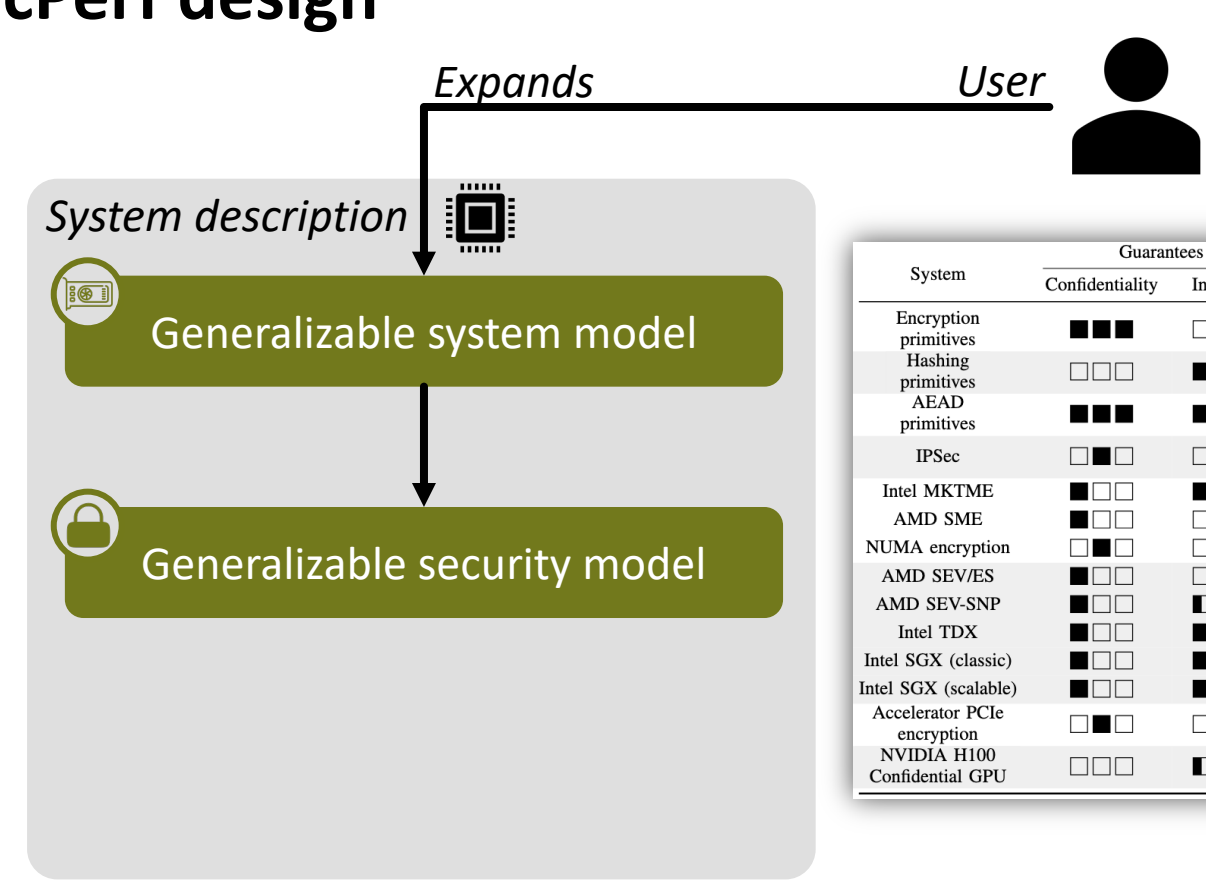
	Data-in-use	Data-in-motion	Data-at-rest
Confidentiality	Intel TDX *AES-XTS*	TLS *ChaCha20*	dm-crypt *AES-XTS*
Integrity	TDX ECC fields *SHA3*	Wireguard *Poly1305*	dm-integrity *AES-GCM*
Freshness	SGX v1 *Merkle tree*	IPsec *PSN window*	dm-x *Merkle tree*

Minimal differentiable and high-level set of security properties

SecPerf design



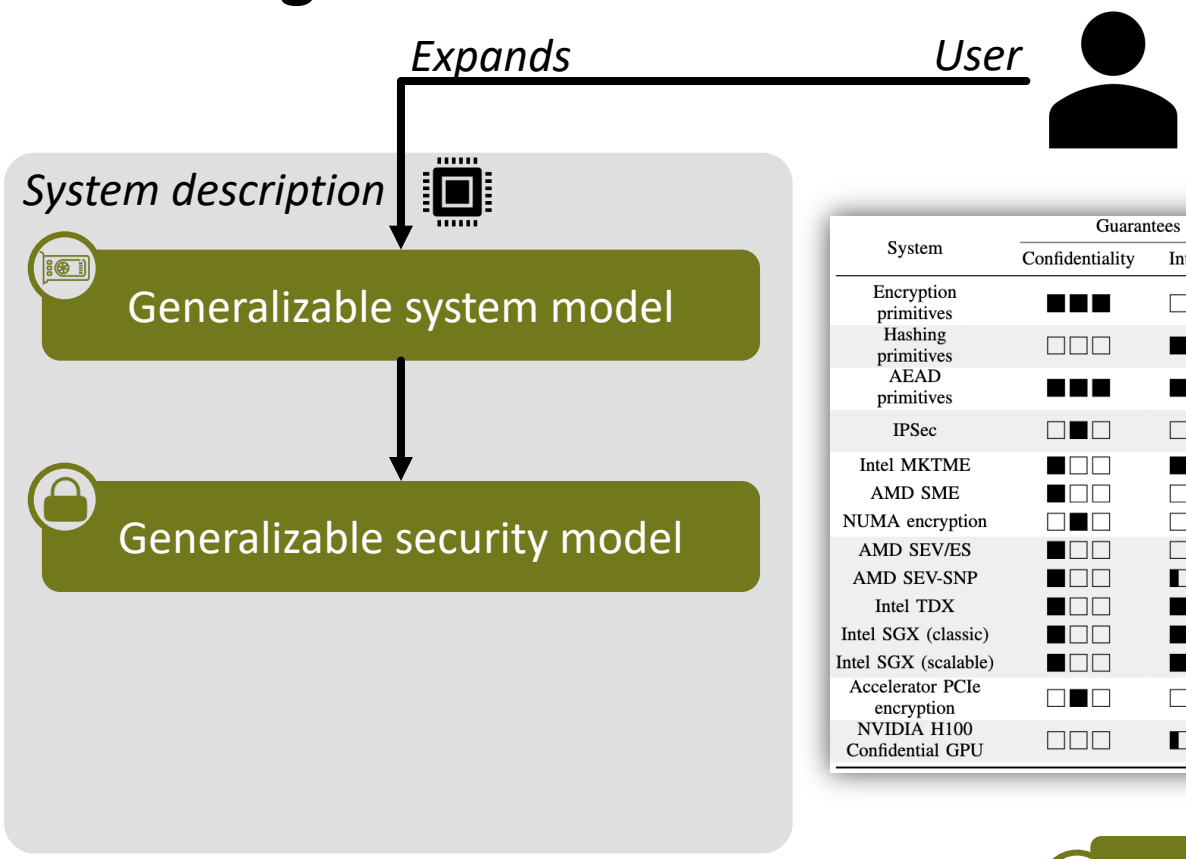
SecPerf design



Combined in a table with hardware, primitives, benchmarks, and frontends

System	Guarantees (use, motion, store)				Type	Hardware	Primitives	Frontends	Benchmark
	Confidentiality	Integrity	Freshness	Trust					
Encryption primitives	■■■	□□□	□□□	□		AES-NI, SIMD	AES CTR256, AES XTS256, ChaCha20	CryptoPP, libsodium, OpenSSL	Own
Hashing primitives	□□□	■■■	□□□	□		SHA-NI, SIMD, BF2 GGA	SHA256, SHA512, Blake2, Blake3	CryptoPP, libsodium, OpenSSL, BLAKE3, DOCA	Own DOCA Bench
AEAD primitives	■■■	■■■	□□□	□		AES-NI, SIMD	AES GCM, AEGIS128L, ChaCha20-Poly1305	CryptoPP, libsodium, OpenSSL	Own
IPSec	□■□	□■□	□■□	□		AES-NI, SIMD BF2 NIC	AES GCM	Wireguard, OpenSwan NIC offload	UCX, OSU, Graph500 LULESH, ML proxy
Intel MKTME	■□□	■□□	□□□	□		MKTME EE	AES XTS, SHA3	–	MLC
AMD SME	■□□	□□□	□□□	□		SME EE	AES XEX	–	
NUMA encryption	□■□	□■□	□□□	□		?	?	–	
AMD SEV/ES	■□□	□□□	□□□	■		AMD SEV/ES	–	–	Black Scholes, GEMM, GROMACS, LULESH, HPCG, HPL, MILC, NPB, SPEC CPU 2017, TPC-C
AMD SEV-SNP	■□□	◧□□	□□□	■		AMD SEV-SNP	–	–	
Intel TDX	■□□	■□□	□□□	■		Intel TDX	–	–	
Intel SGX (classic)	■□□	■□□	■□□	■		Intel SGX	–	–	
Intel SGX (scalable)	■□□	■□□	□□□	■		Intel SGX	–	–	
Accelerator PCIe encryption	□■□	□■□	□□□	□		NVIDIA bounce buffer EE	AES-GCM	–	NVbandwidth
NVIDIA H100 Confidential GPU	□□□	◧□□	□□□	■		NVIDIA bounce buffer EE	–	–	torchbench

SecPerf design

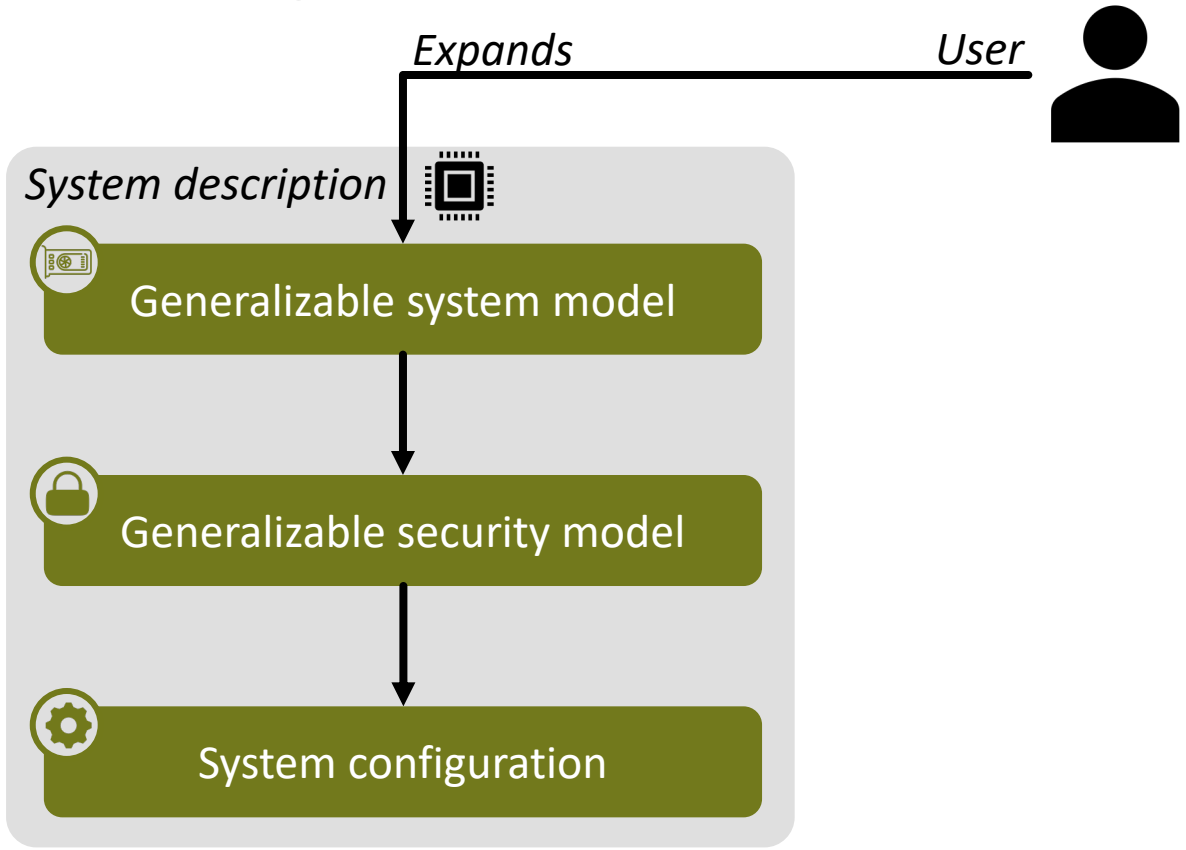


Combined in a table with hardware, primitives, benchmarks, and frontends

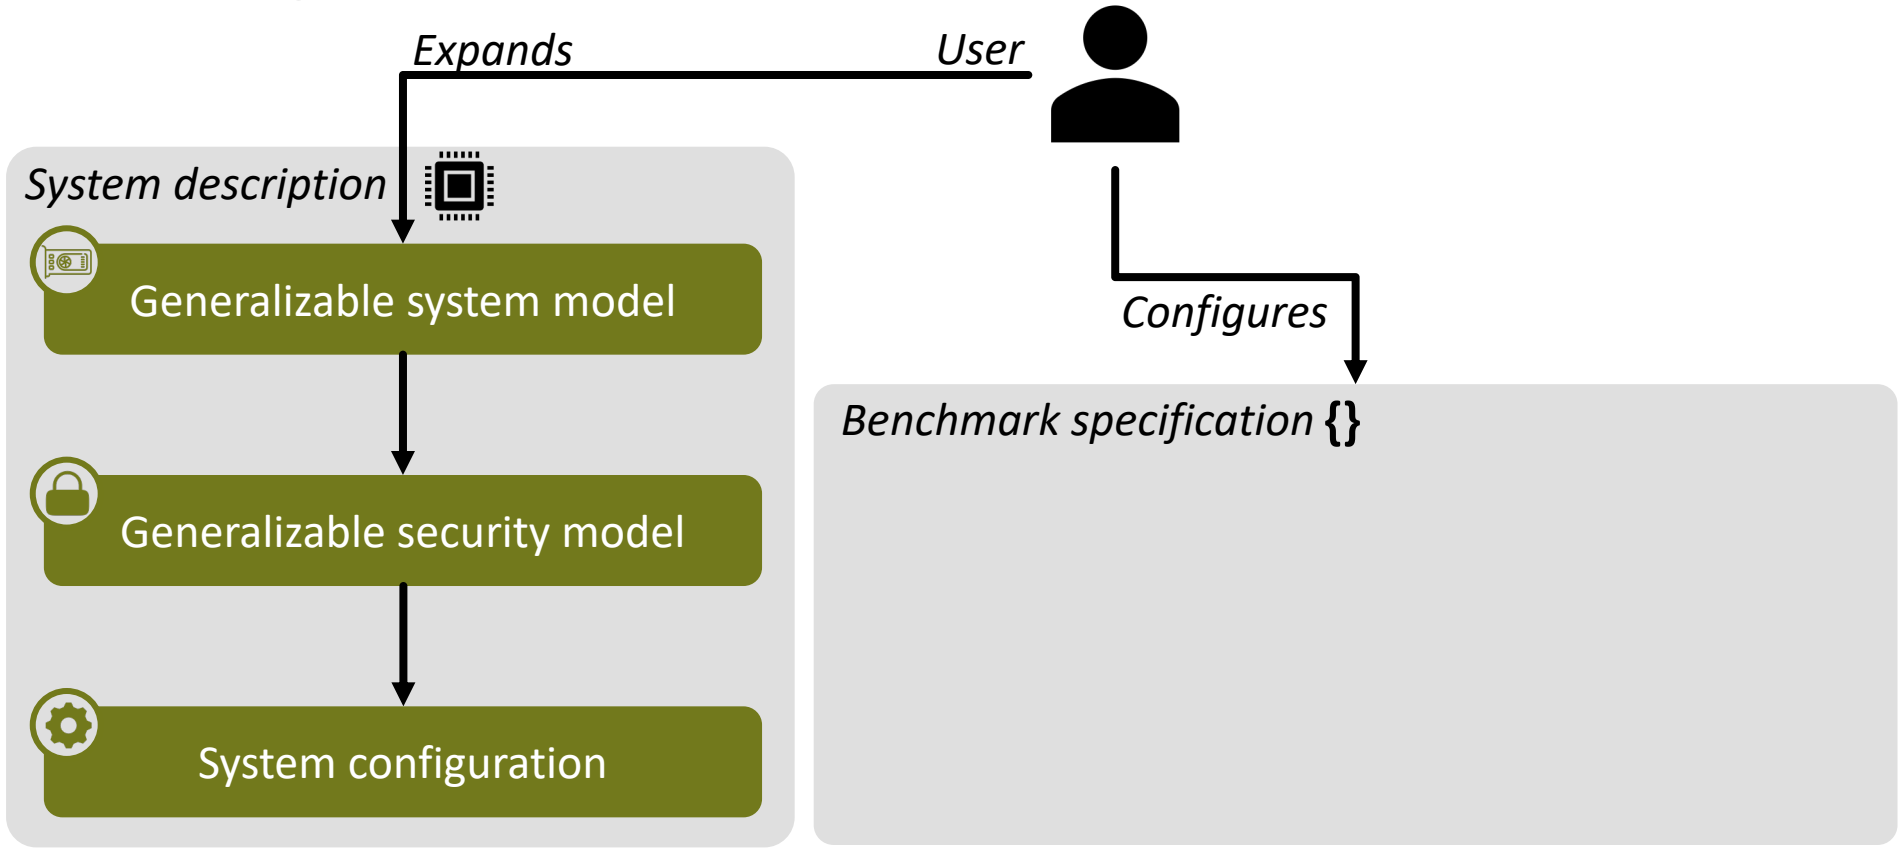
System	Guarantees (use, motion, store)				Type	Hardware	Primitives	Frontends	Benchmark
	Confidentiality	Integrity	Freshness	Trust					
Encryption primitives	■■■	□□□	□□□	□		AES-NI, SIMD	AES CTR256, AES XTS256, ChaCha20	CryptoPP, libsodium, OpenSSL	Own
Hashing primitives	□□□	■■■	□□□	□		SHA-NI, SIMD, BF2 GGA	SHA256, SHA512, Blake2, Blake3	CryptoPP, libsodium, OpenSSL, BLAKE3, DOCA	Own DOCA Bench
AEAD primitives	■■■	■■■	□□□	□		AES-NI, SIMD	AES GCM, AEGIS128L, ChaCha20-Poly1305	CryptoPP, libsodium, OpenSSL	Own
IPSec	□■□	□■□	□■□	□		AES-NI, SIMD BF2 NIC	AES GCM	Wireguard, OpenSwan NIC offload	UCX, OSU, Graph500 LULESH, ML proxy
Intel MKTME	■□□	■□□	□□□	□		MKTME EE	AES XTS, SHA3	–	MLC
AMD SME	■□□	□□□	□□□	□		SME EE	AES XEX	–	
NUMA encryption	□■□	□■□	□□□	□		?	?	–	
AMD SEV/ES	■□□	□□□	□□□	■		AMD SEV/ES	–	–	Black Scholes, GEMM, GROMACS, LULESH, HPCG, HPL, MILC, NPB, SPEC CPU 2017, TPC-C
AMD SEV-SNP	■□□	◧□□	□□□	■		AMD SEV-SNP	–	–	
Intel TDX	■□□	■□□	□□□	■		Intel TDX	–	–	
Intel SGX (classic)	■□□	■□□	■□□	■		Intel SGX	–	–	
Intel SGX (scalable)	■□□	■□□	□□□	■		Intel SGX	–	–	
Accelerator PCIe encryption	□■□	□■□	□□□	□		NVIDIA bounce buffer EE	AES-GCM	–	NVbandwidth
NVIDIA H100 Confidential GPU	□□□	◧□□	□□□	■		NVIDIA bounce buffer EE	–	–	torchbench

Differentiate solutions by guarantees provided by default

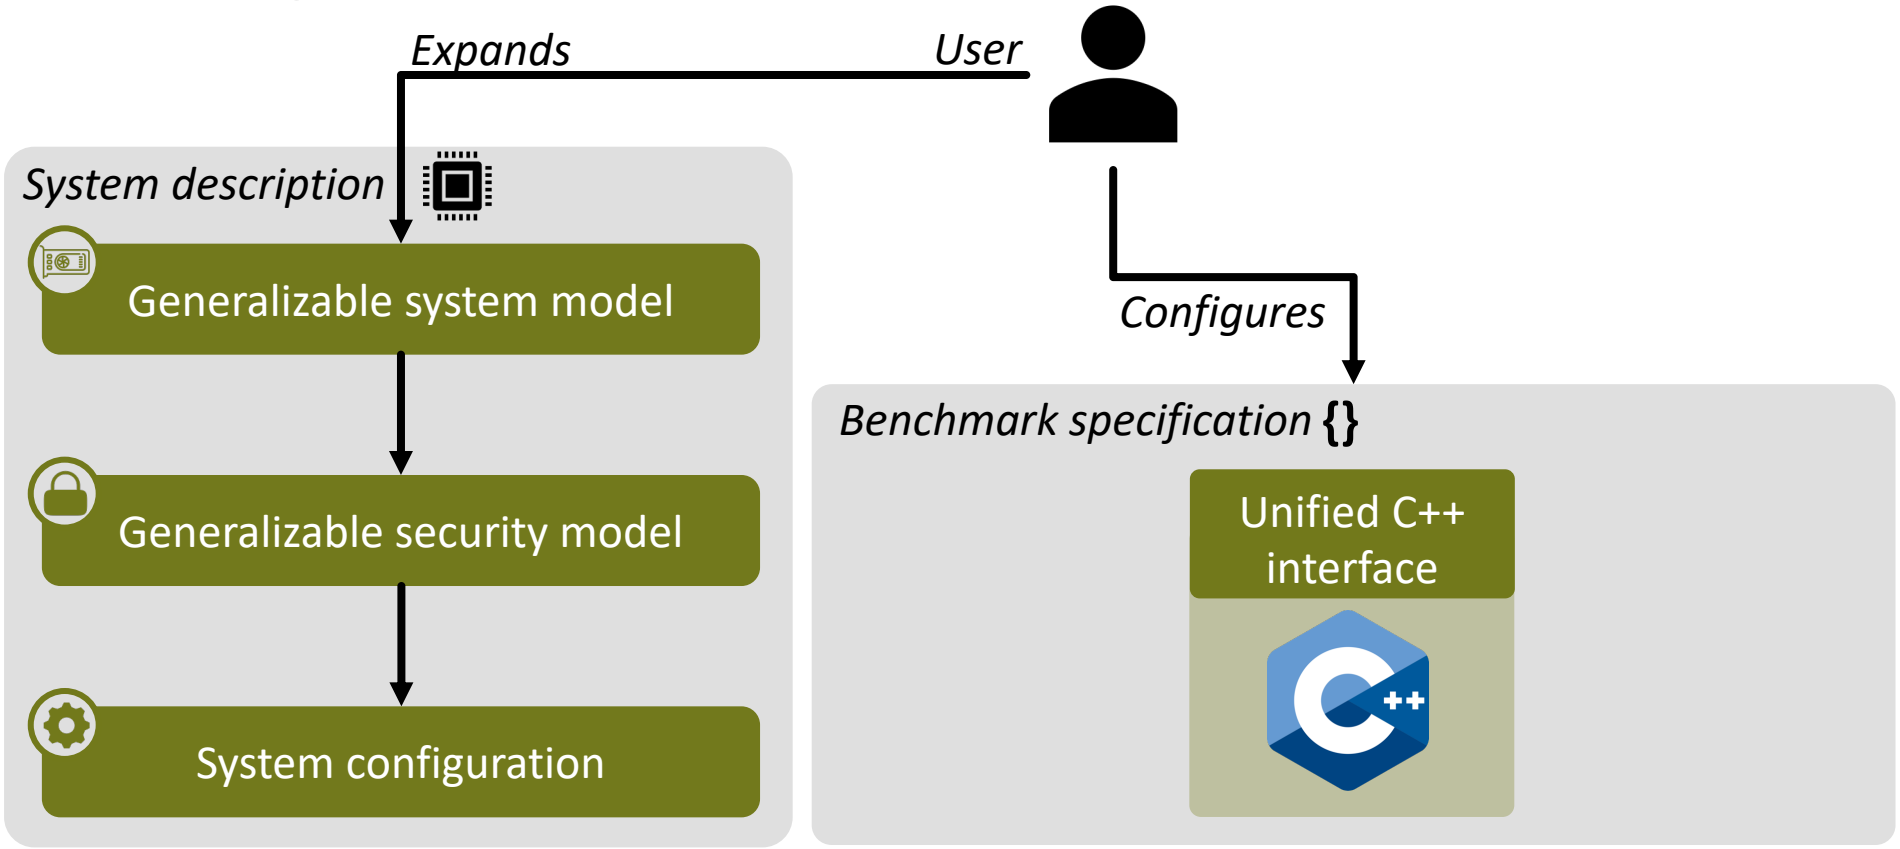
SecPerf design



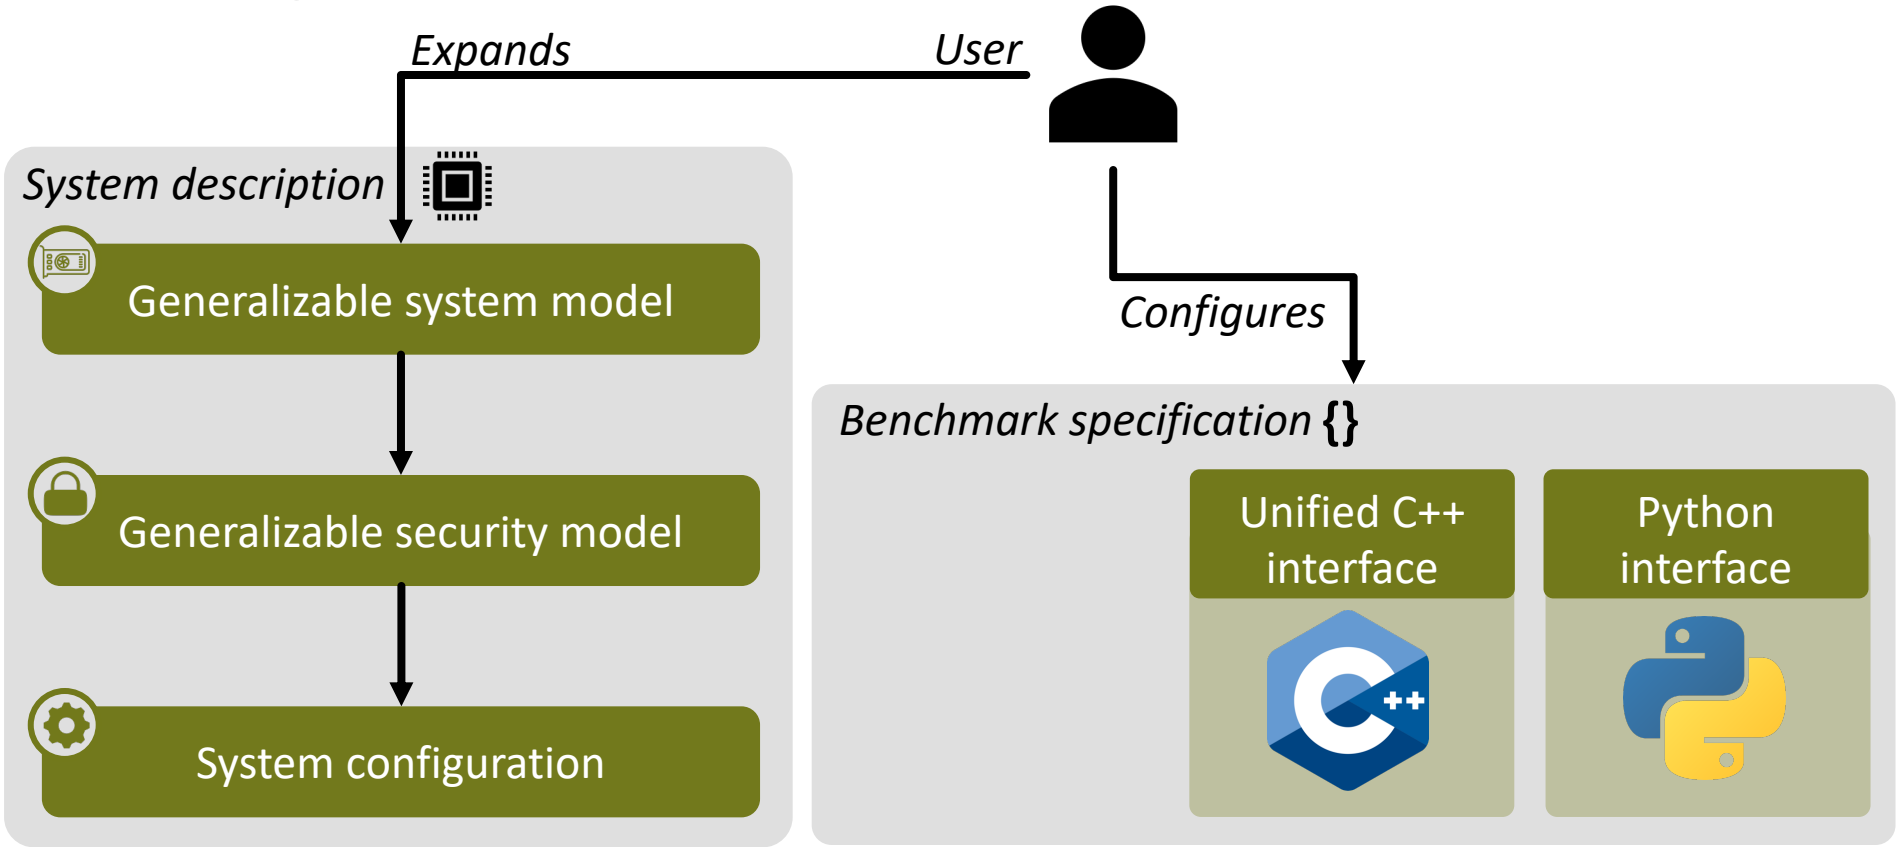
SecPerf design



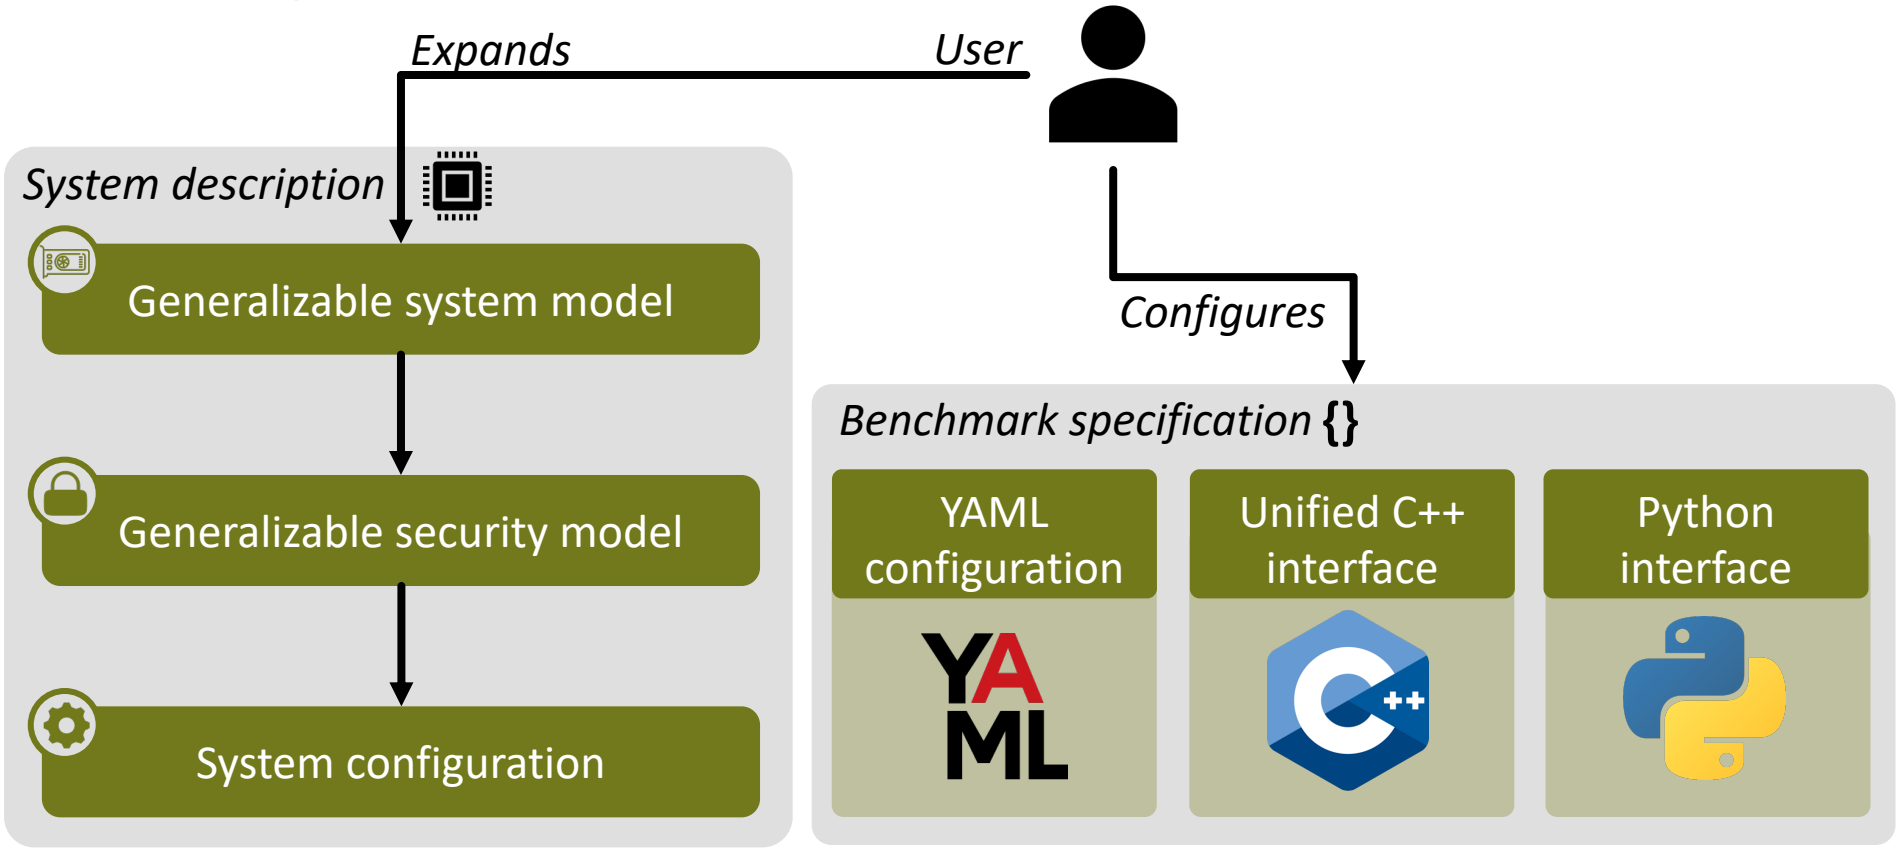
SecPerf design



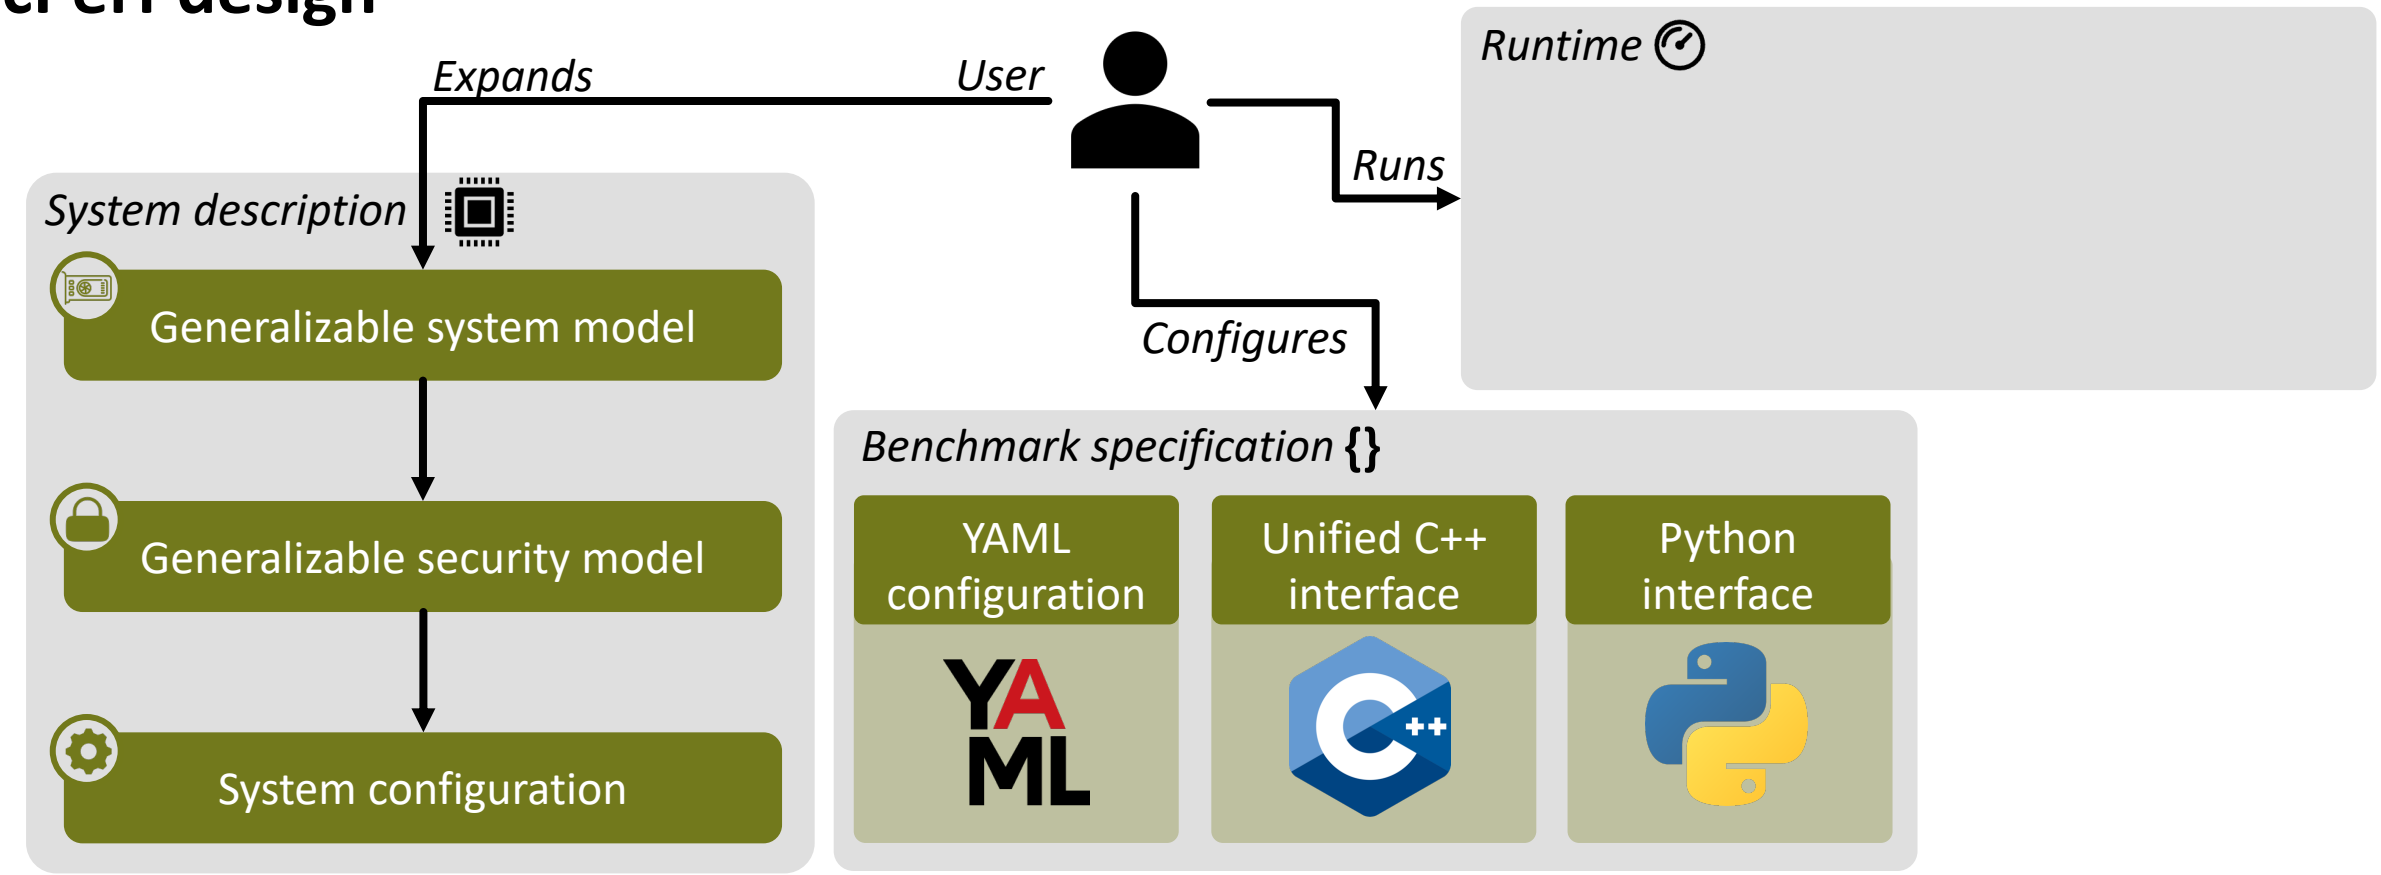
SecPerf design



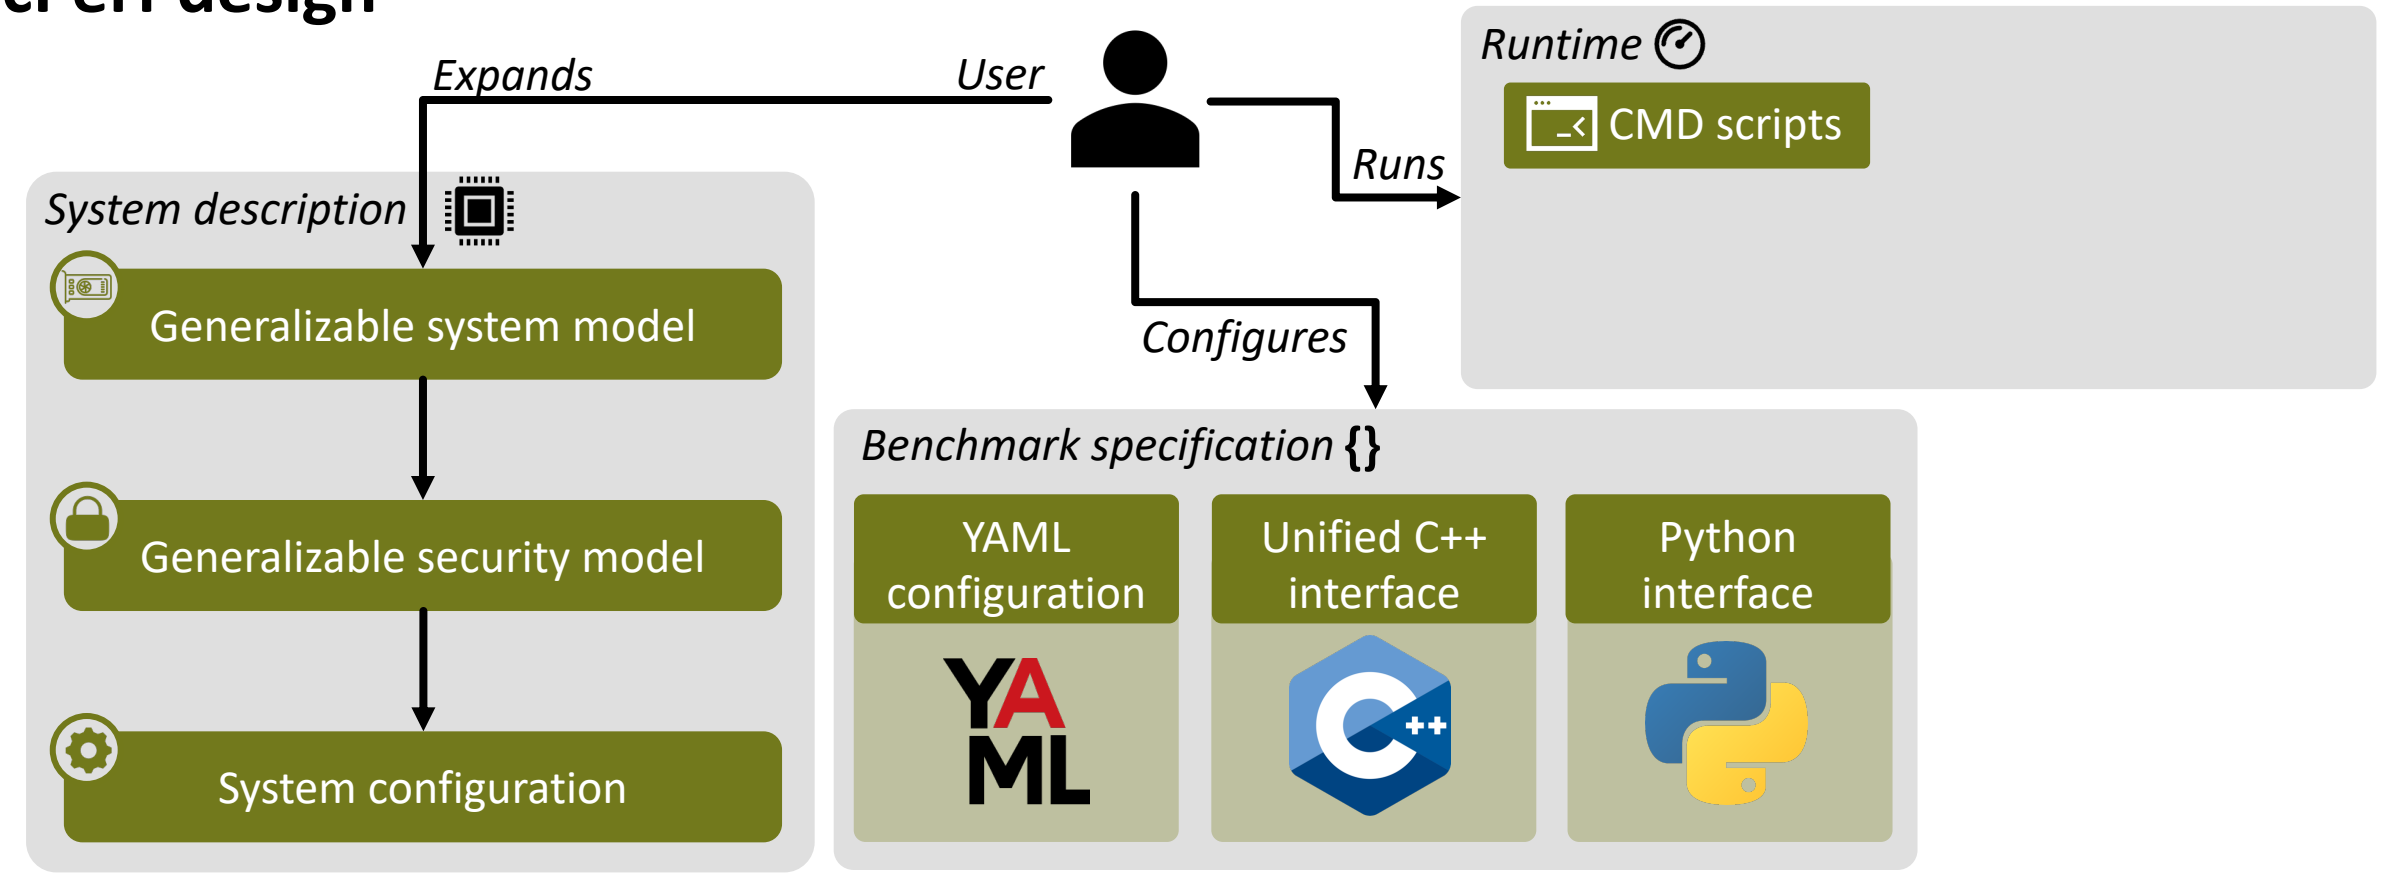
SecPerf design



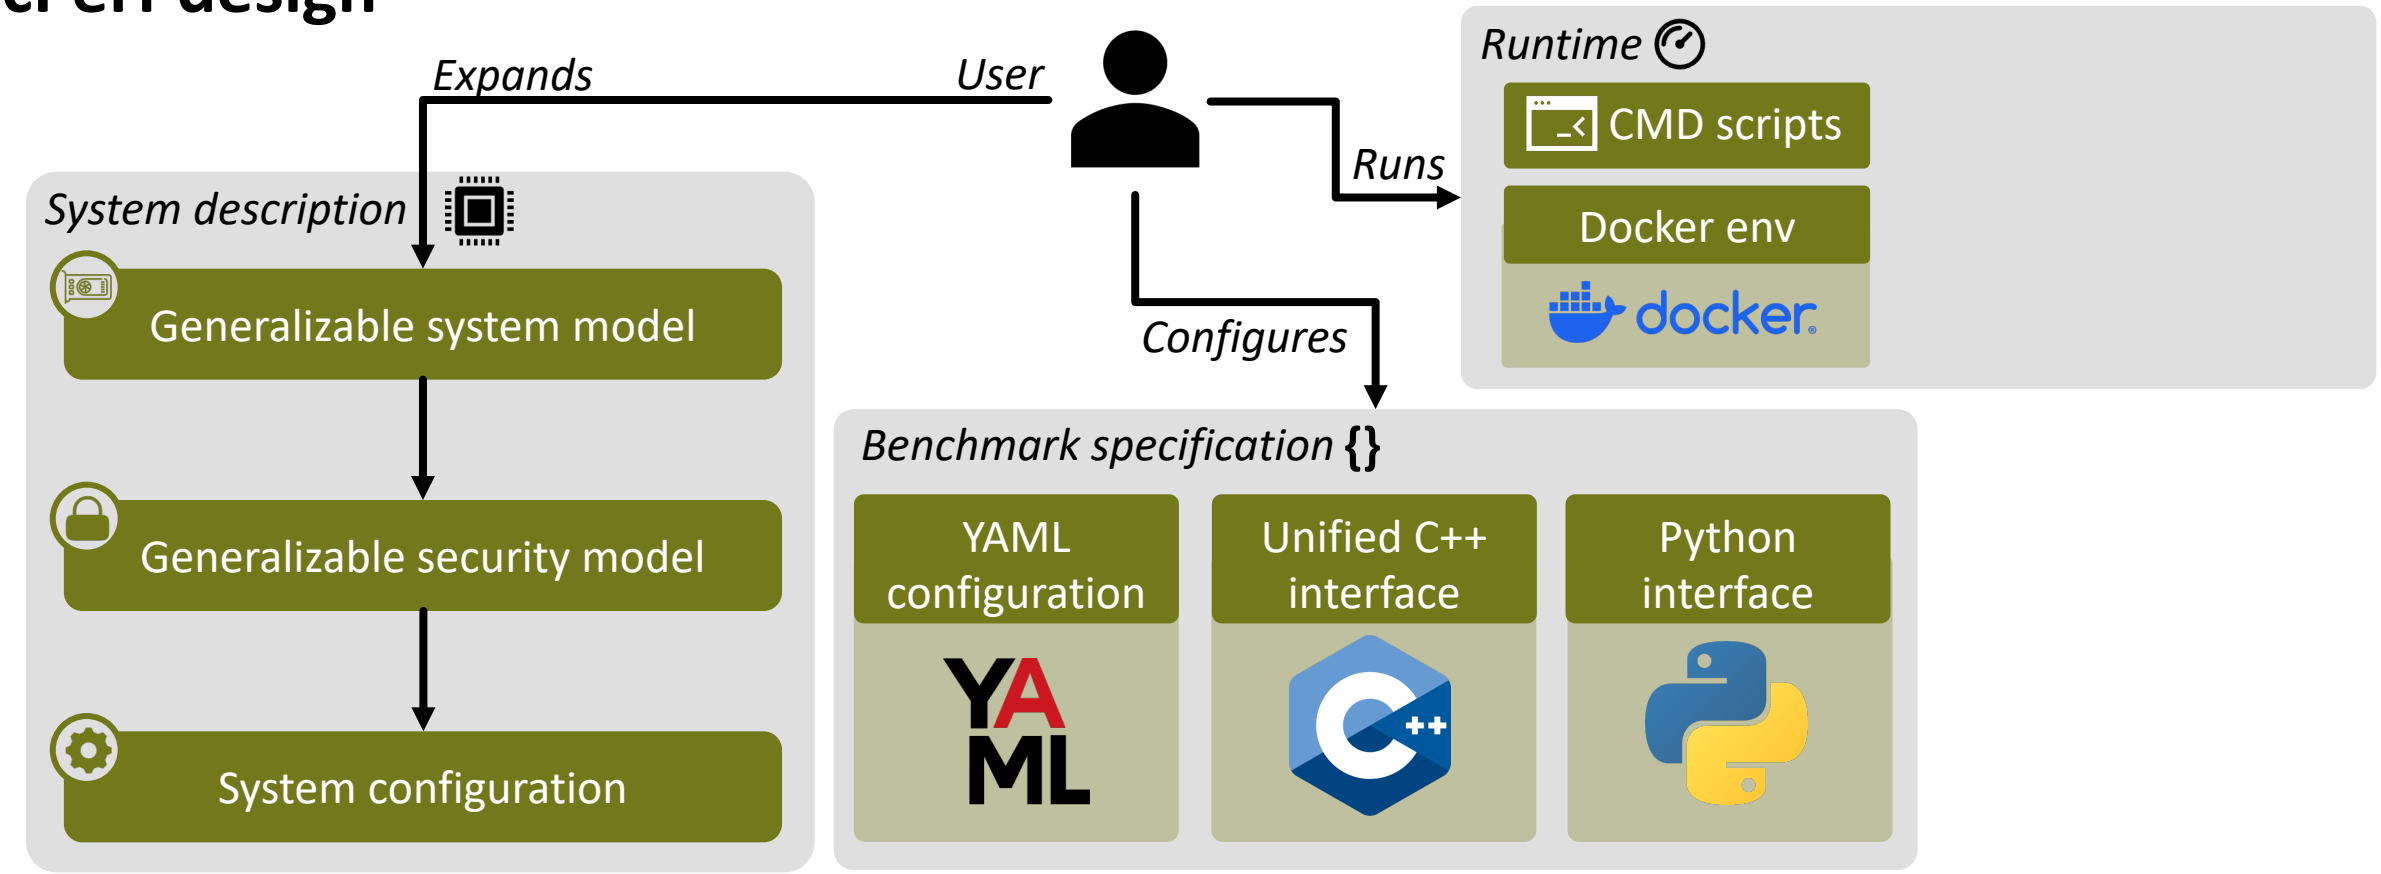
SecPerf design



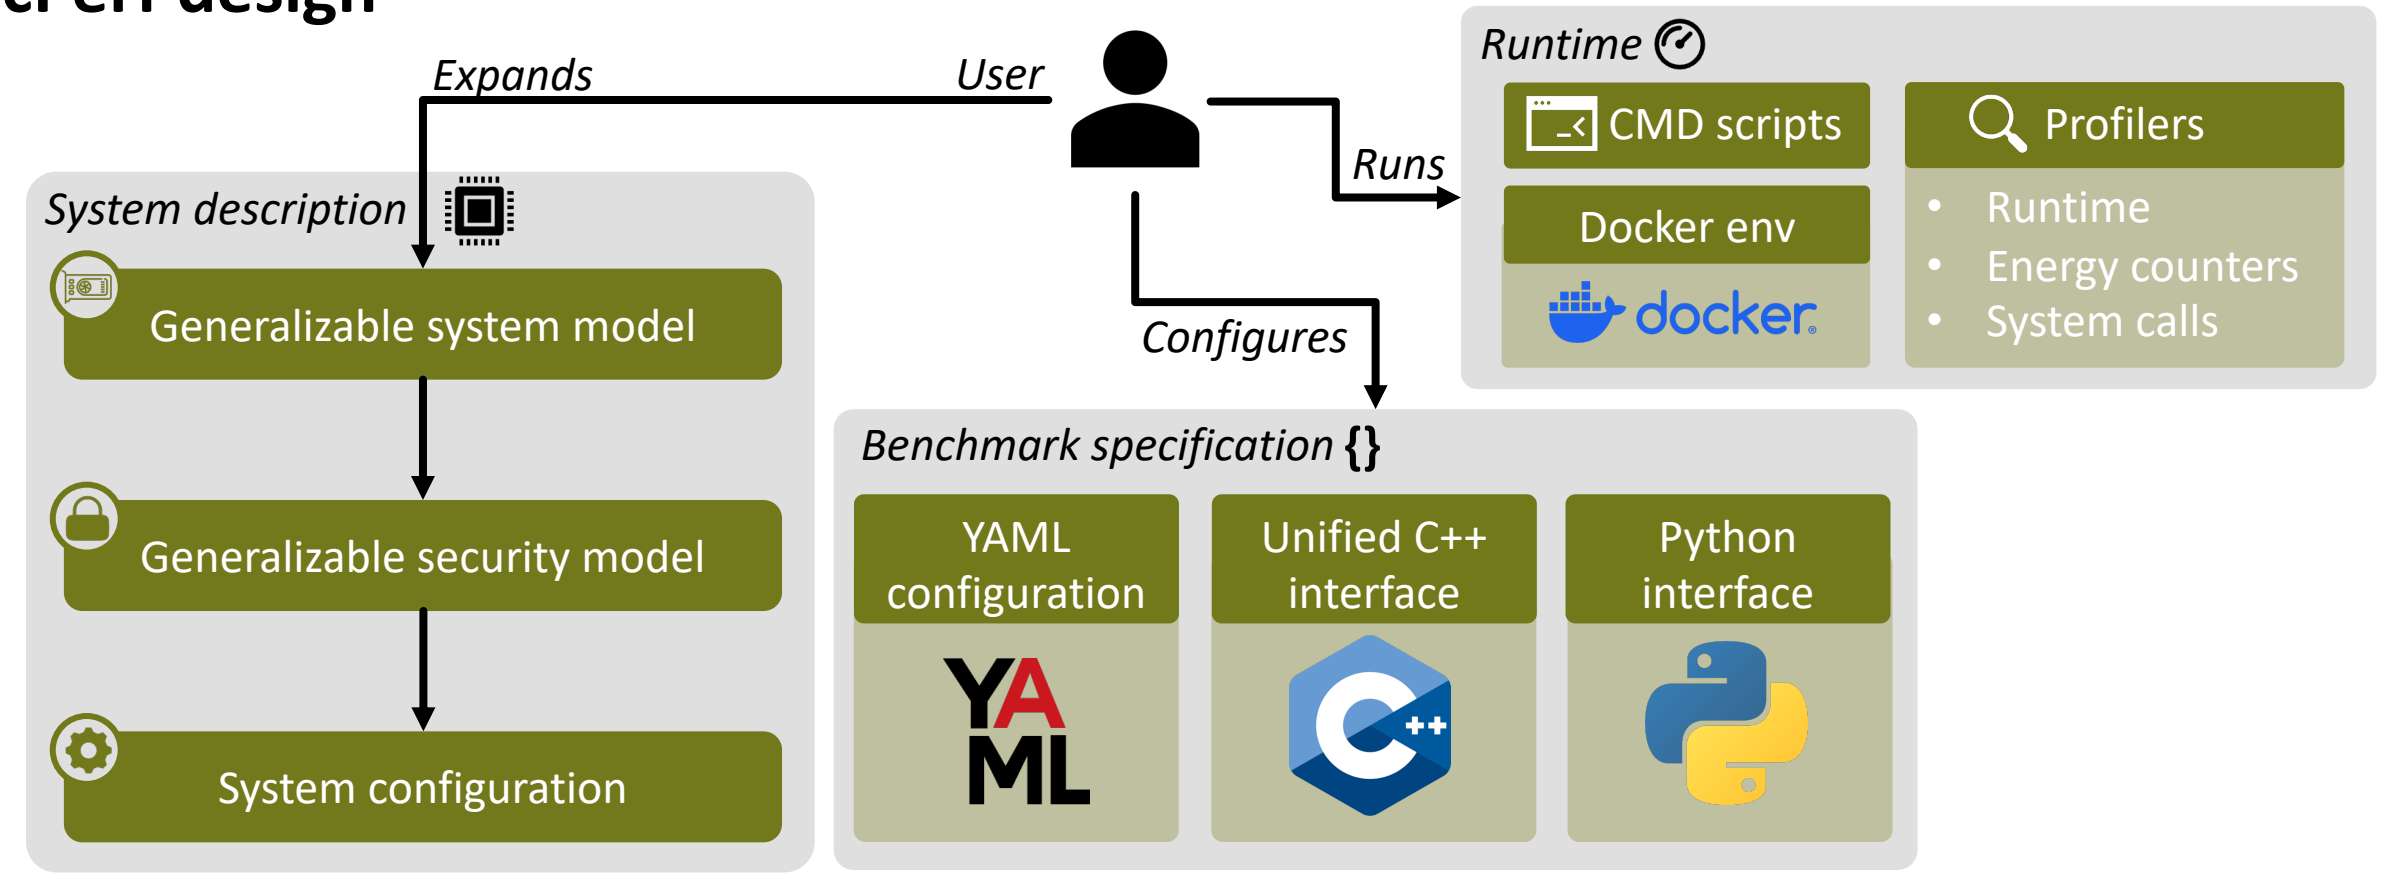
SecPerf design



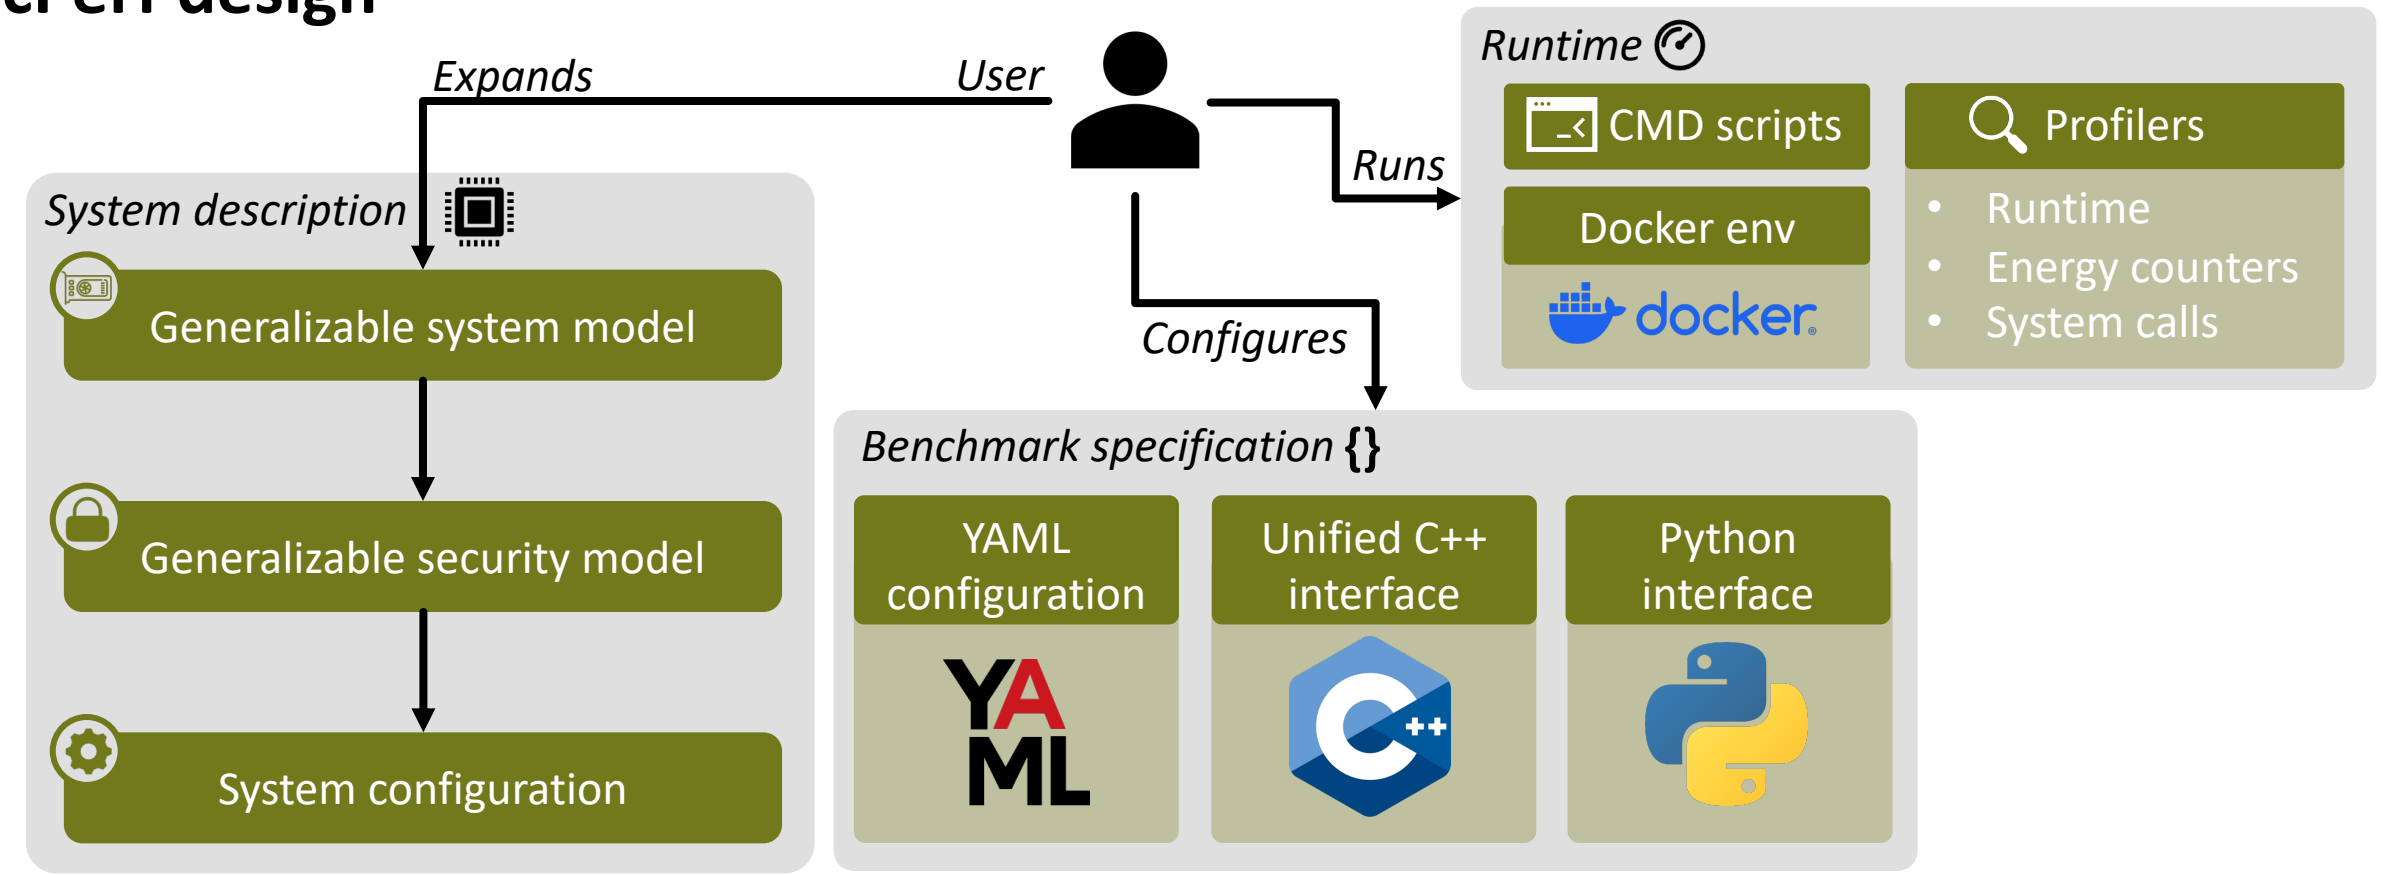
SecPerf design



SecPerf design



SecPerf design



We measure **25 algorithms, 18 network primitives, 29 modern HPC application variants, providing 12 insights and 8 recommendations**
This makes SecPerf largest evaluation to date of secure hardware and systems.

Select systems

Which hashing algorithm is best for secure storage?

Hazel: Secure and Efficient Disaggregated Storage

Marcin Chrapek, Meni Orenbach, Ahmad Atamli, Marcin Copik, Mikhail Khalilov, Fritz Alder, Torsten Hoefler

Select systems

Which hashing algorithm is best for secure storage?

Implemented 12 different interfaces for hashing (C++ interface)

SIMD
SHA-NI
BF2

OpenSSL::SHA256
Cryptopp::SHA256
Libsodium::SHA256
DOCA::SHA256

OpenSSL::SHA512
Cryptopp::SHA512
Libsodium::SHA512
DOCA::SHA512

OpenSSL::BLAKE2
Cryptopp::BLAKE2
Libsodium::BLAKE2

BLAKE3::BLAKE3

Hazel: Secure and Efficient Disaggregated Storage

Marcin Chrapek, Meni Orenbach, Ahmad Atamli, Marcin Copik, Mikhail Khalilov, Fritz Alder, Torsten Hoefler

Select systems

Which hashing algorithm is best for secure storage?

Implemented 12 different interfaces for hashing (C++ interface)

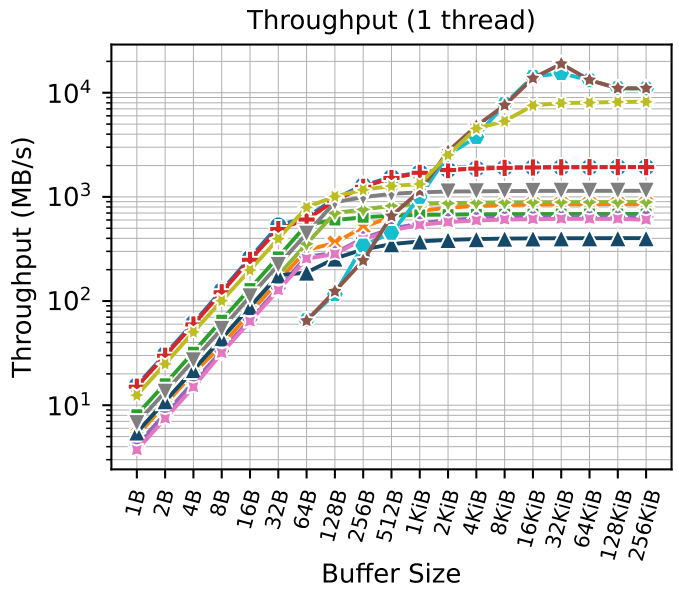
SIMD
SHA-NI
BF2

OpenSSL::SHA256
Cryptopp::SHA256
Libsodium::SHA256
DOCA::SHA256

OpenSSL::SHA512
Cryptopp::SHA512
Libsodium::SHA512
DOCA::SHA512

OpenSSL::BLAKE2
Cryptopp::BLAKE2
Libsodium::BLAKE2

BLAKE3::BLAKE3



Hazel: Secure and Efficient Disaggregated Storage

Marcin Chrapek, Meni Orenbach, Ahmad Atamli, Marcin Copik, Mikhail Khalilov, Fritz Alder, Torsten Hoefler

Select systems

Which hashing algorithm is best for secure storage?

Implemented 12 different interfaces for hashing (C++ interface)

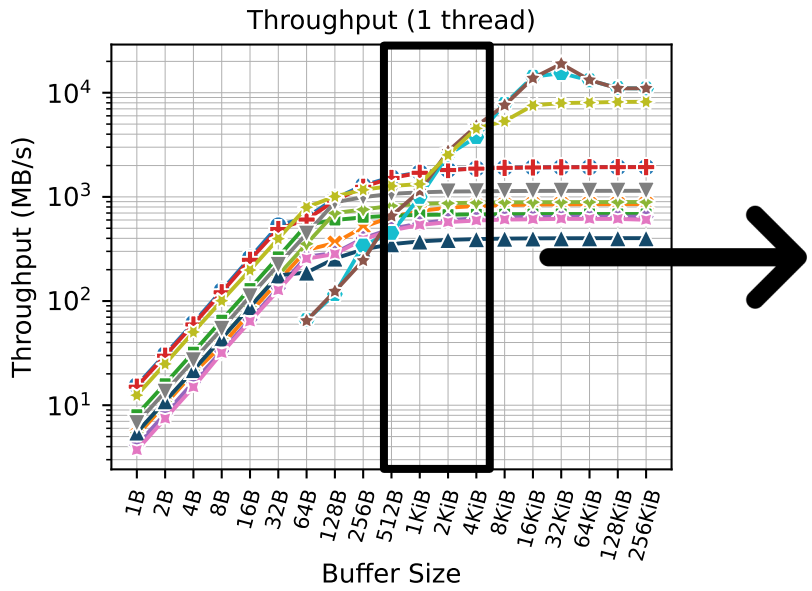
SIMD
SHA-NI
BF2

OpenSSL::SHA256
Cryptopp::SHA256
Libsodium::SHA256
DOCA::SHA256

OpenSSL::SHA512
Cryptopp::SHA512
Libsodium::SHA512
DOCA::SHA512

OpenSSL::BLAKE2
Cryptopp::BLAKE2
Libsodium::BLAKE2

BLAKE3::BLAKE3





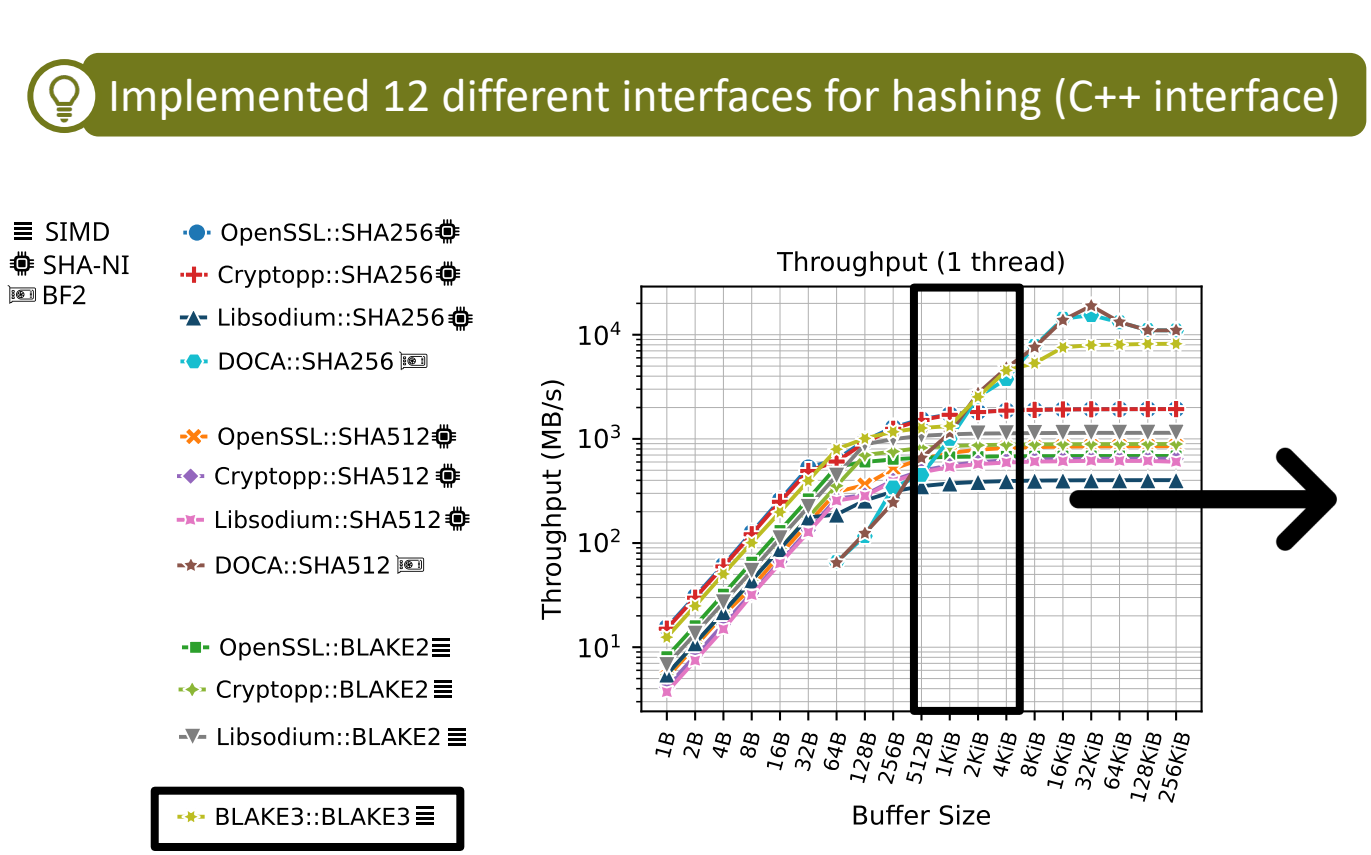
Hazel: Secure and Efficient Disaggregated Storage

Marcin Chrapek, Meni Orenbach, Ahmad Atamli, Marcin Copik, Mikhail Khalilov, Fritz Alder, Torsten Hoefler

Select systems

Which hashing algorithm is best for secure storage?

Implemented 12 different interfaces for hashing (C++ interface)





BLAKE3 performs well for 1-4 KiB and requires less trust than offloading

Hazel: Secure and Efficient Disaggregated Storage

Marcin Chrapek, Meni Orenbach, Ahmad Atamli, Marcin Copik, Mikhail Khalilov, Fritz Alder, Torsten Hoefler

Select systems

Which hashing algorithm is best for secure storage?

Implemented 12 different interfaces for hashing (C++ interface)

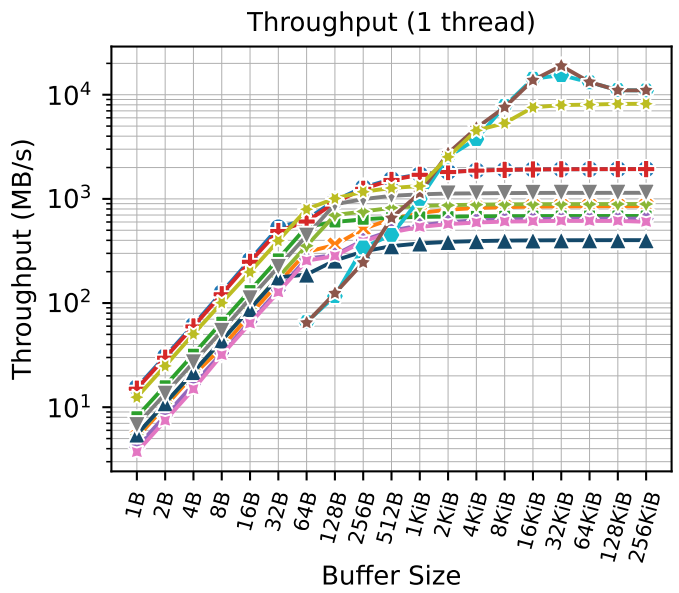
SIMD
SHA-NI
BF2

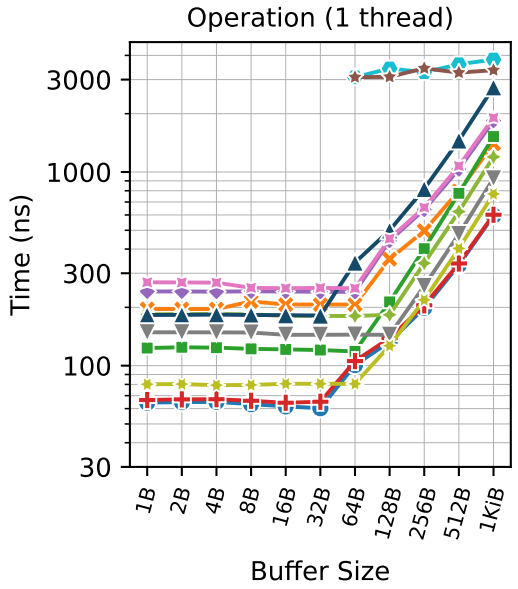
OpenSSL::SHA256
Cryptopp::SHA256
Libsodium::SHA256
DOCA::SHA256

OpenSSL::SHA512
Cryptopp::SHA512
Libsodium::SHA512
DOCA::SHA512

OpenSSL::BLAKE2
Cryptopp::BLAKE2
Libsodium::BLAKE2

BLAKE3::BLAKE3





BLAKE3 performs well for 1-4 KiB and requires less trust than offloading

Hazel: Secure and Efficient Disaggregated Storage

Marcin Chrapek, Meni Orenbach, Ahmad Atamli, Marcin Copik, Mikhail Khalilov, Fritz Alder, Torsten Hoefler

Select systems

Which hashing algorithm is best for secure storage?

Implemented 12 different interfaces for hashing (C++ interface)

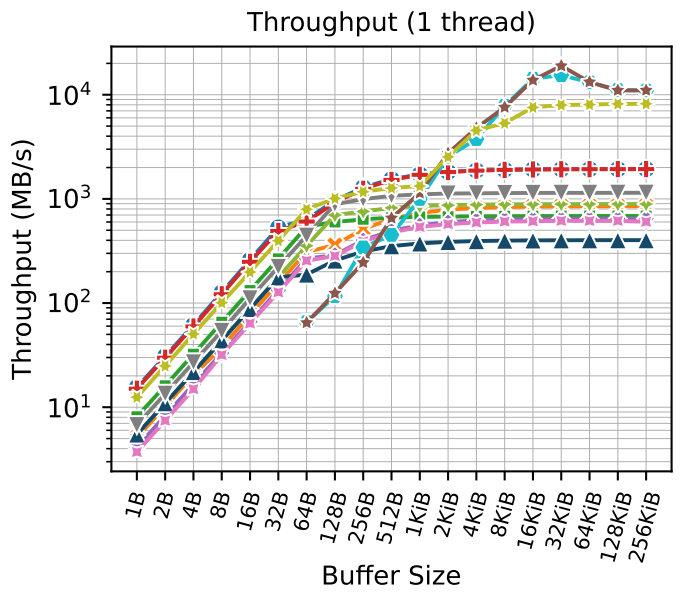
SIMD
SHA-NI
BF2

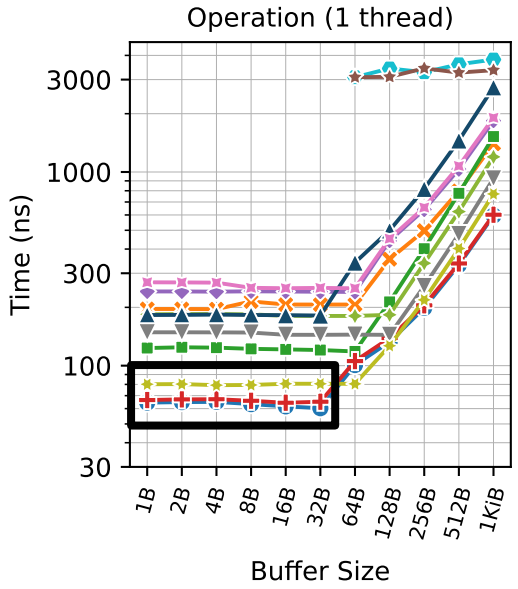
OpenSSL::SHA256
Cryptopp::SHA256
Libsodium::SHA256
DOCA::SHA256

OpenSSL::SHA512
Cryptopp::SHA512
Libsodium::SHA512
DOCA::SHA512

OpenSSL::BLAKE2
Cryptopp::BLAKE2
Libsodium::BLAKE2

BLAKE3::BLAKE3





BLAKE3 performs well for 1-4 KiB and requires less trust than offloading

However, BLAKE3 is not strictly better for all cases!

Hazel: Secure and Efficient Disaggregated Storage

Marcin Chrapek, Meni Orenbach, Ahmad Atamli, Marcin Copik, Mikhail Khalilov, Fritz Alder, Torsten Hoefler

Identify bottlenecks

What is a major bottleneck in common HPC workloads?

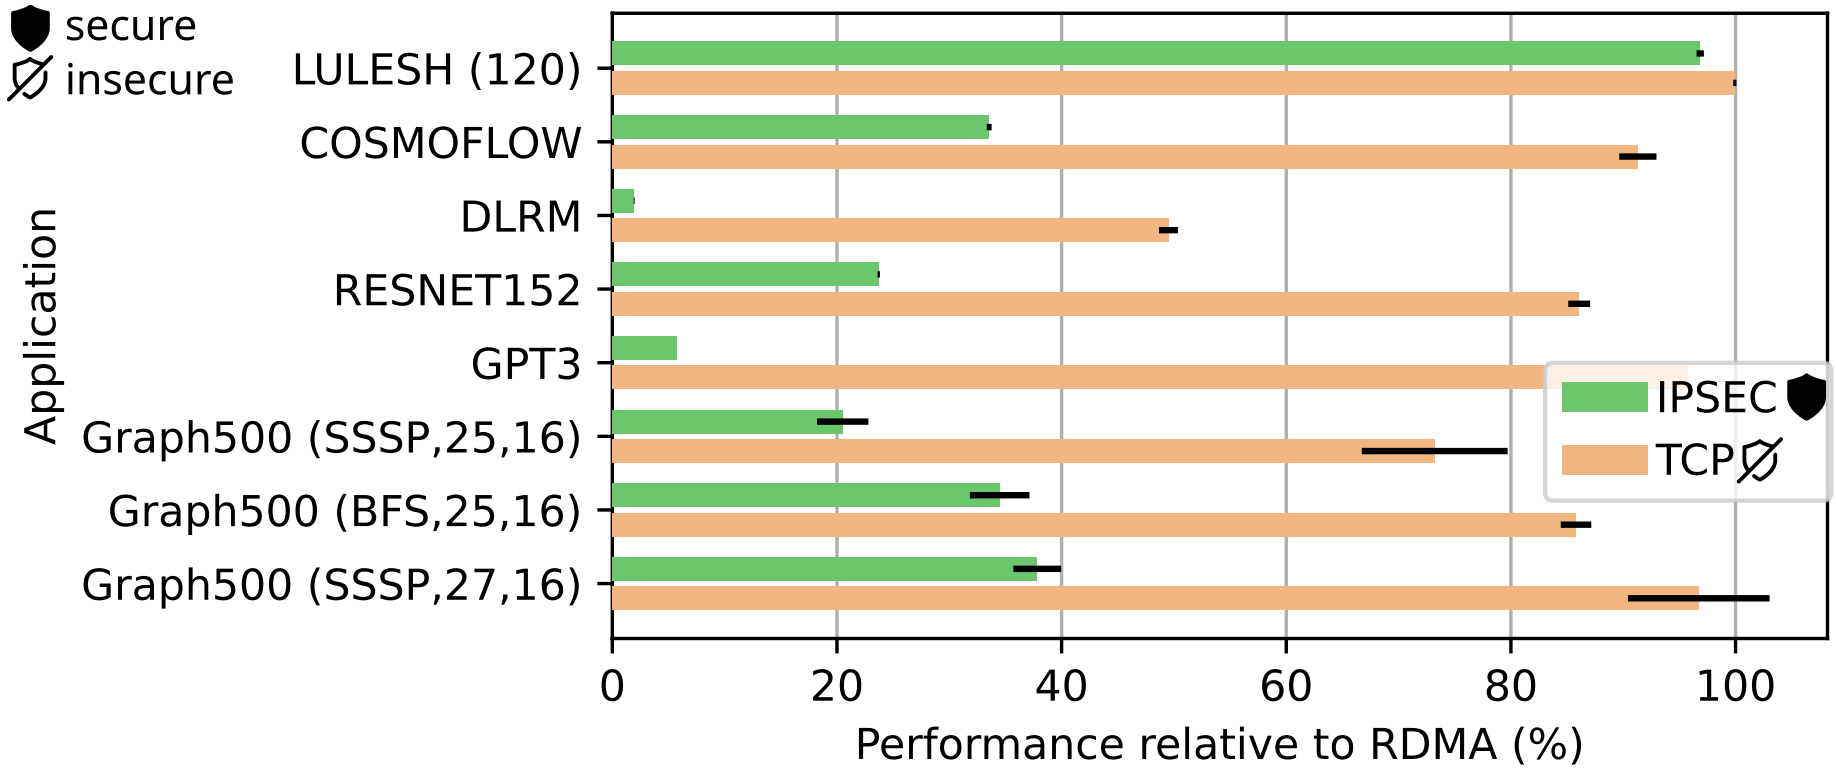
Identify bottlenecks

What is a major bottleneck in common HPC workloads?

Networking requires a protection scheme such as IPsec.

Identify bottlenecks

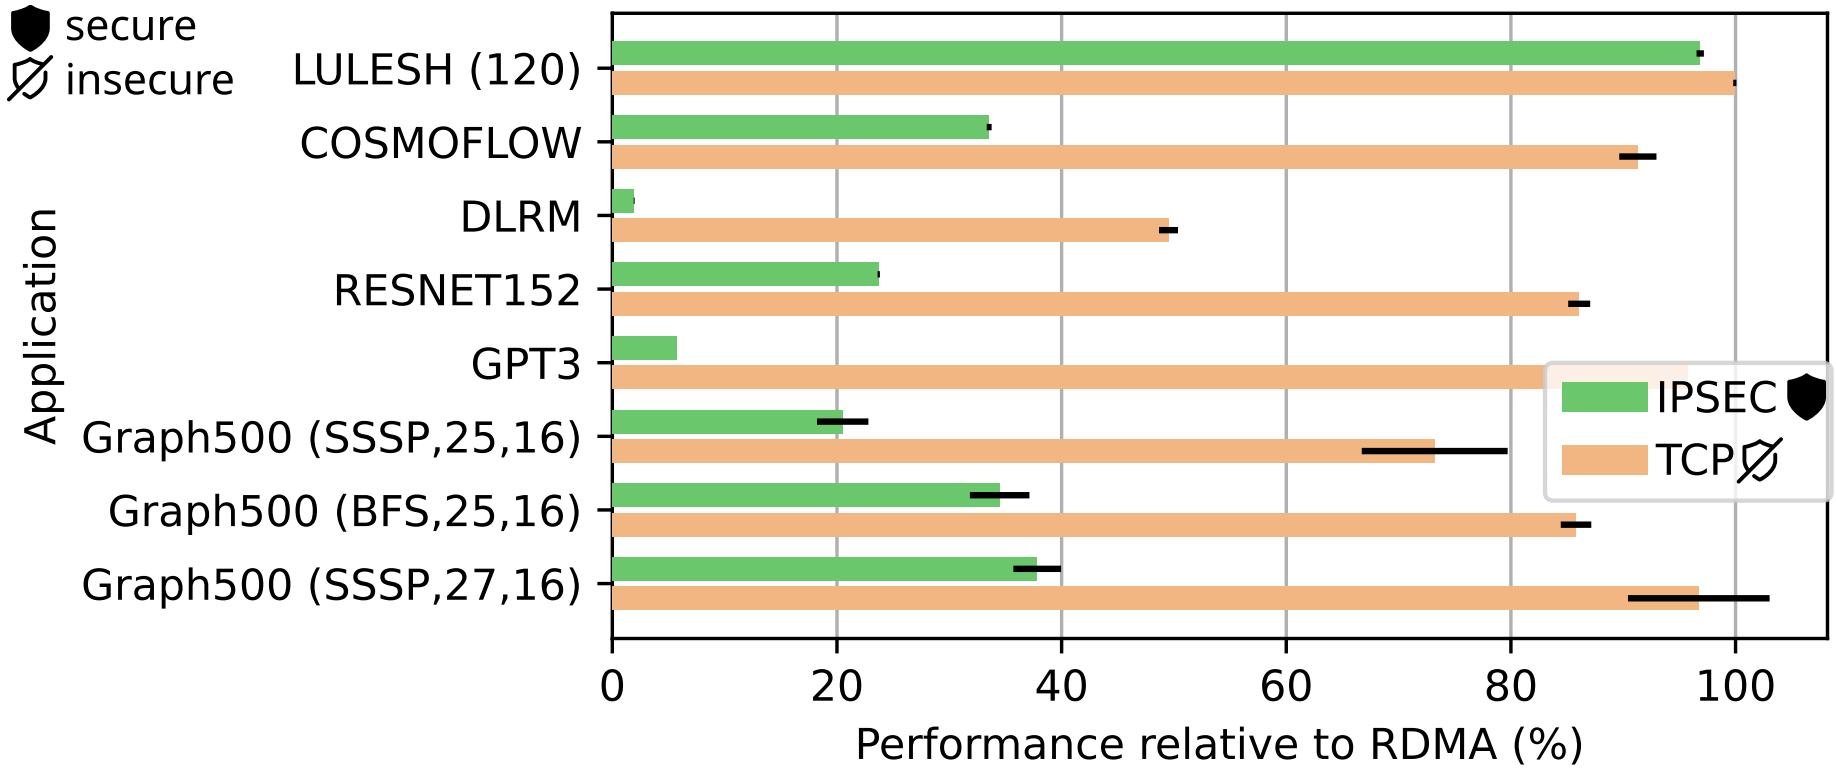
What is a major bottleneck in common HPC workloads?



Networking requires a protection scheme such as IPsec.

Identify bottlenecks

What is a major bottleneck in common HPC workloads?

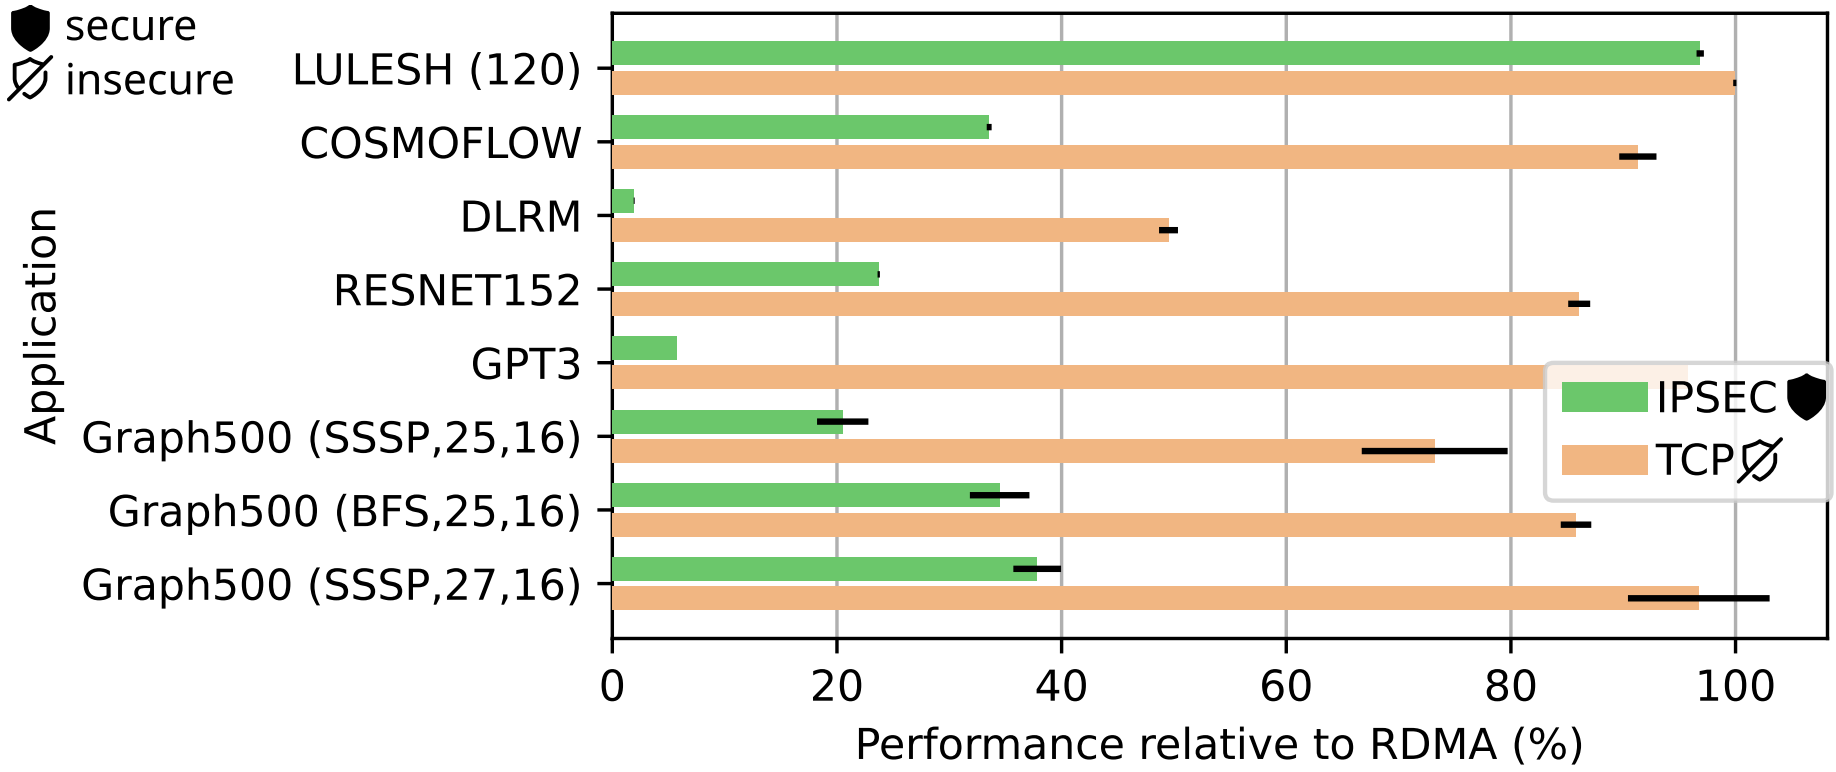


Networking requires a protection scheme such as IPsec.

TCP performs reasonably well
IPsec lowers performance by up to 90%.

Identify bottlenecks

What is a major bottleneck in common HPC workloads?



Networking requires a protection scheme such as IPsec.

TCP performs reasonably well
IPsec lowers performance by up to 90%.

TEEs are not the only source of overheads for HPC setups!

Predict performance

Can we estimate the performance of applications?

Predict performance

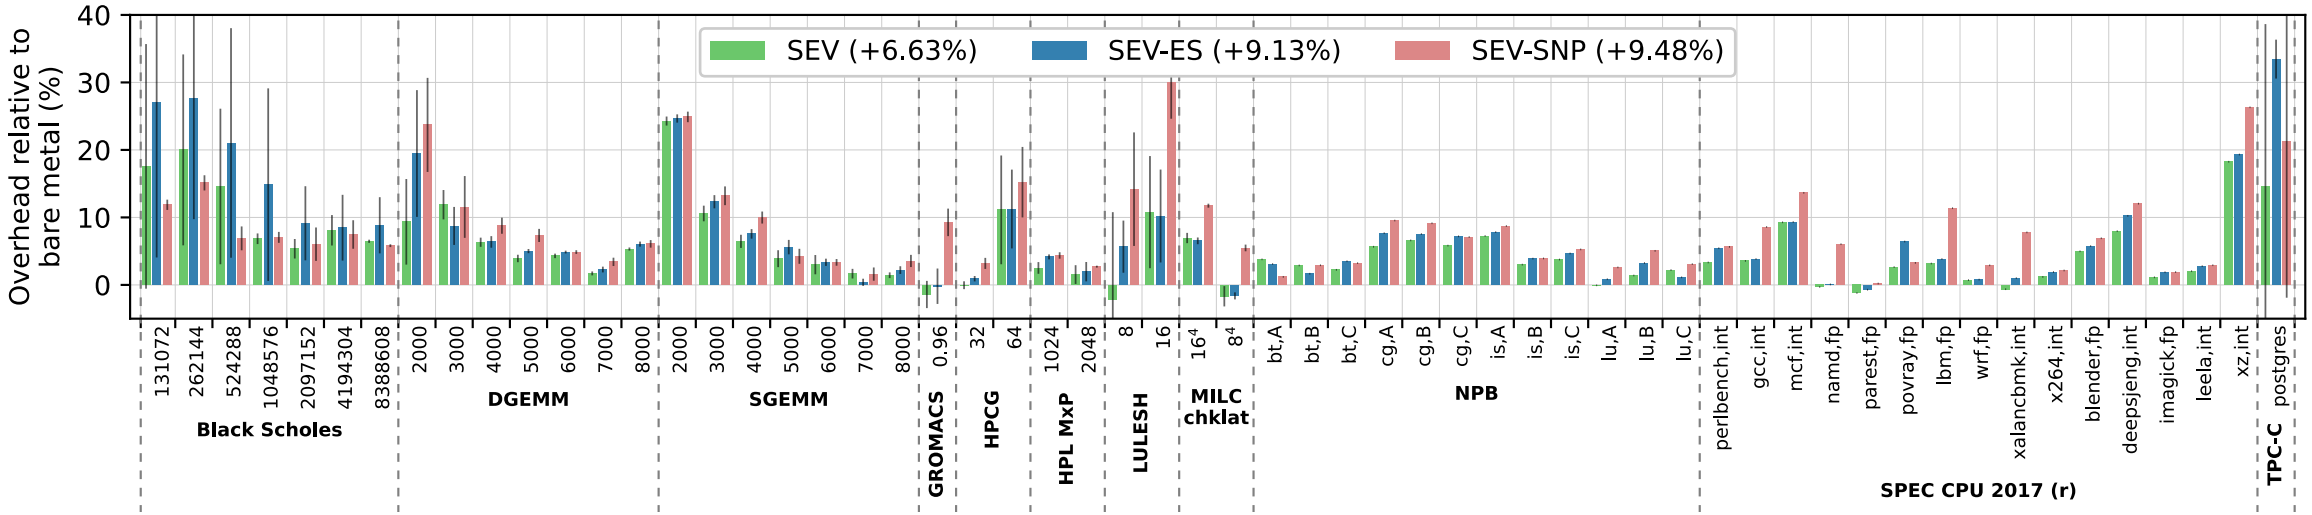
Can we estimate the performance of applications?

11 different applications with subset of runtime parameters (Python interface)

Predict performance

Can we estimate the performance of applications?

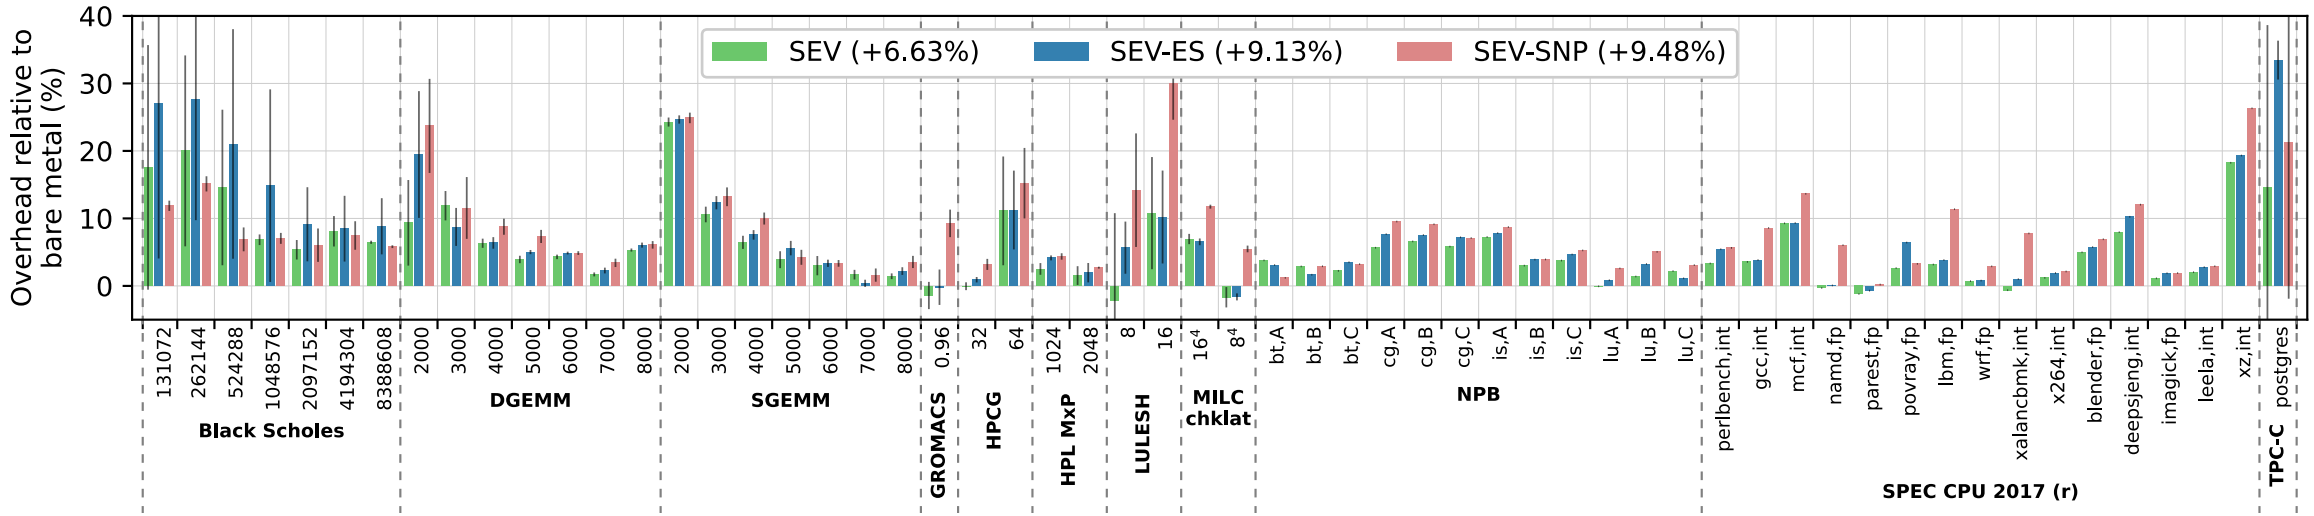
11 different applications with subset of runtime parameters (Python interface)



Predict performance

Can we estimate the performance of applications?

11 different applications with subset of runtime parameters (Python interface)

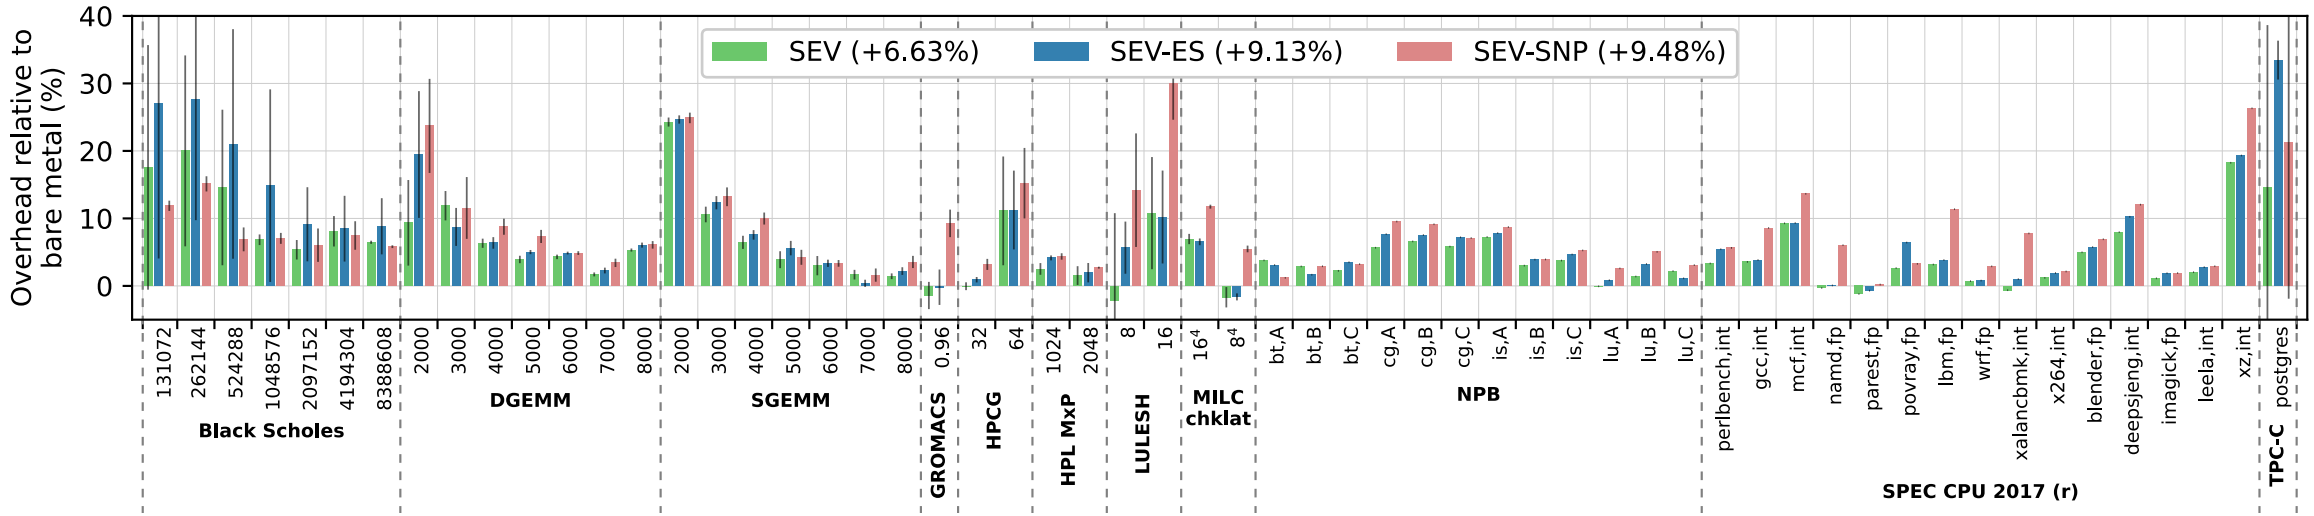


More security = more overhead (see our taxonomy)

Predict performance

Can we estimate the performance of applications?

11 different applications with subset of runtime parameters (Python interface)



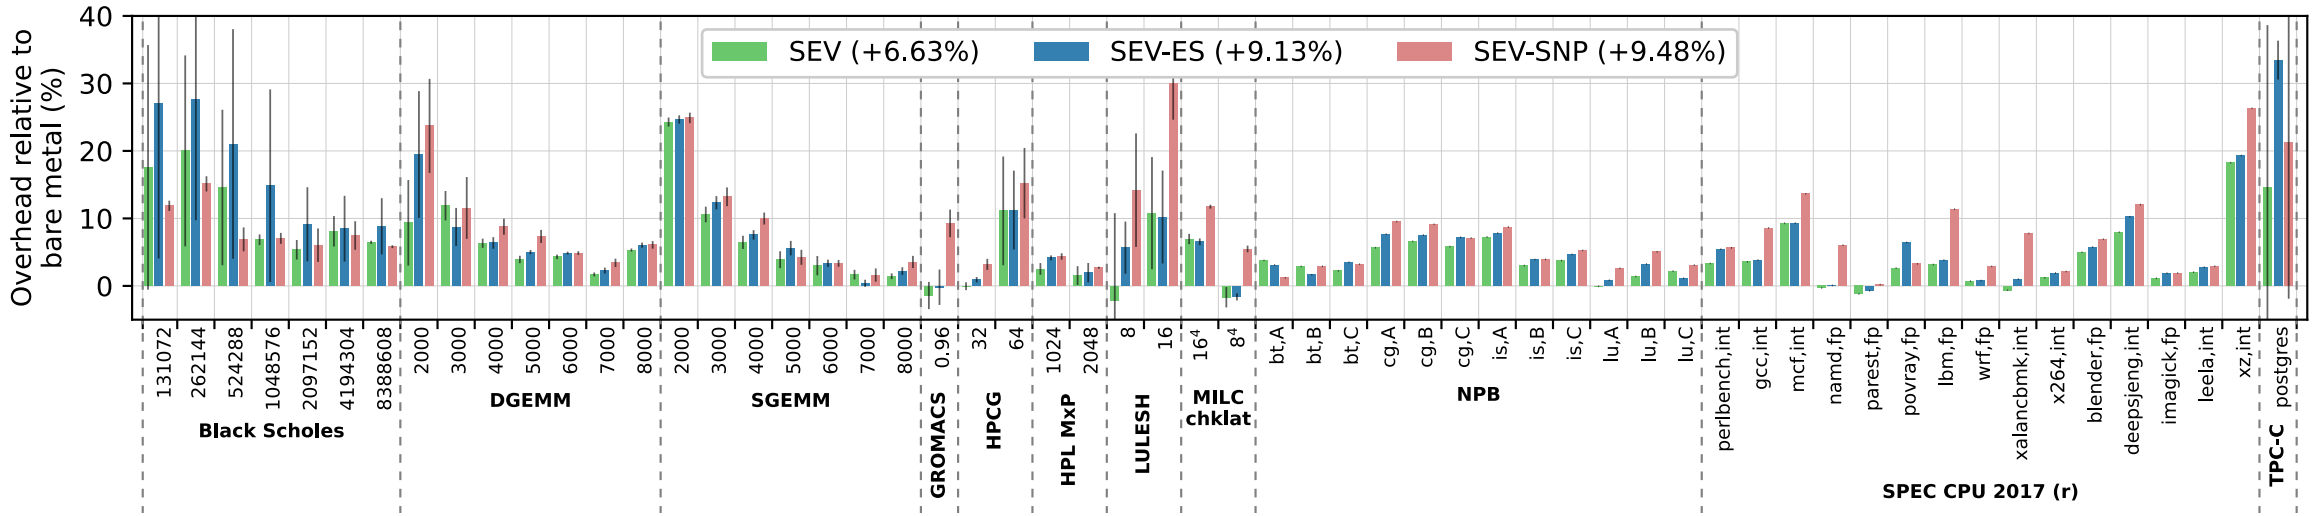
More security = more overhead (see our taxonomy)

More compute heavy = lower overheads

Predict performance

Can we estimate the performance of applications?

11 different applications with subset of runtime parameters (Python interface)



More security = more overhead (see our taxonomy)

More compute heavy = lower overheads

Similar overheads to example applications

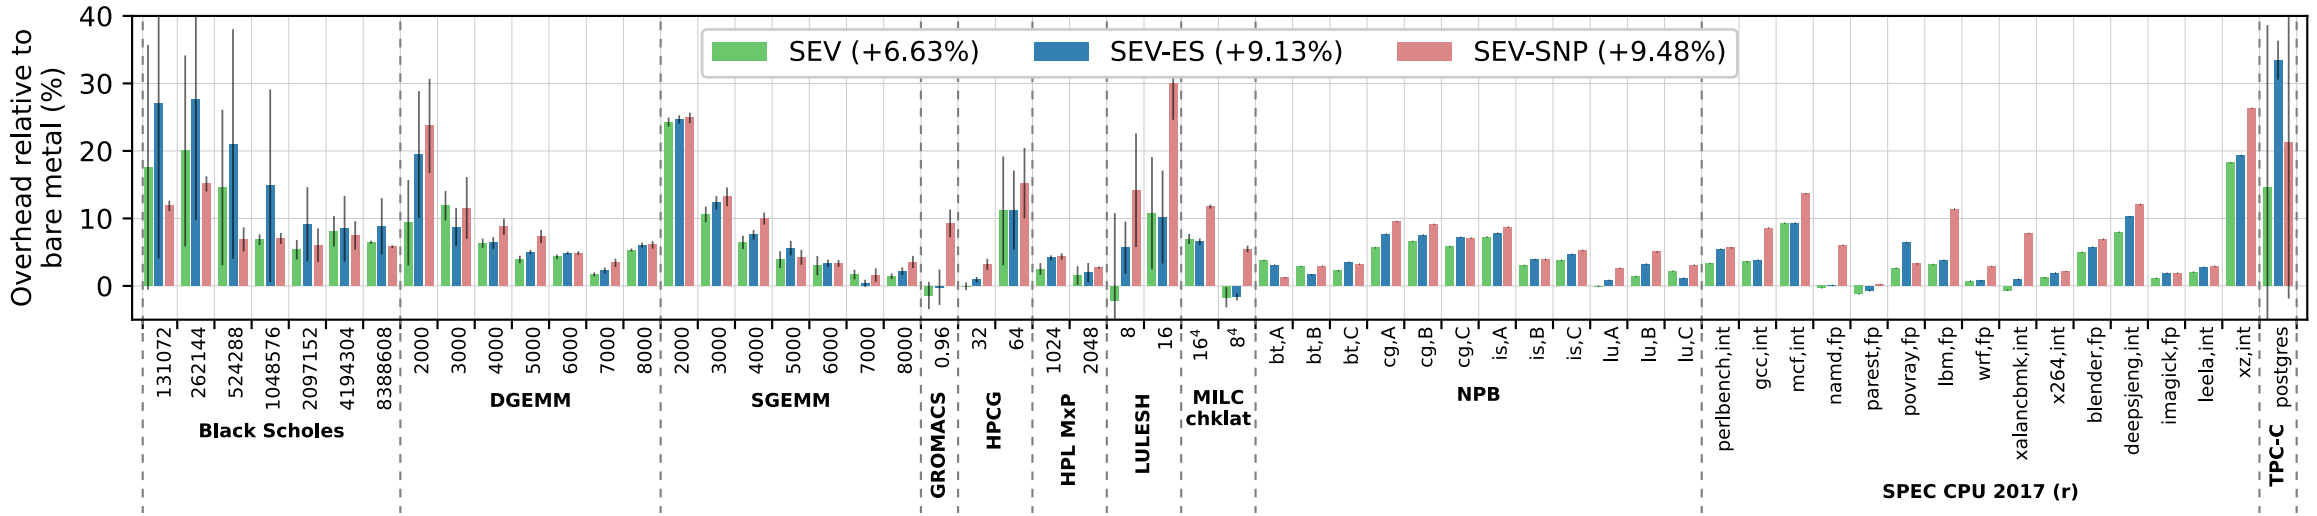
Confidential LLM Inference: Performance and Cost Across CPU and GPU TEEs

Marcin Chrapek, Marcin Copik, Etienne Mettaz, Torsten Hoefler

Predict performance

Can we estimate the performance of applications?

11 different applications with subset of runtime parameters (Python interface)



More security = more overhead (see our taxonomy)

More compute heavy = lower overheads

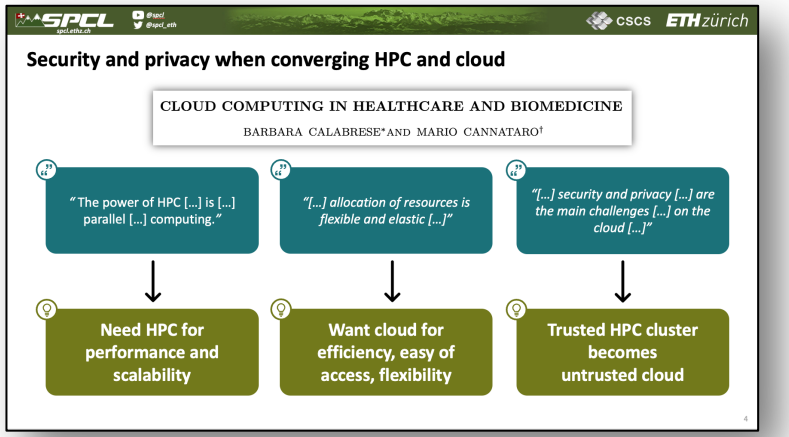
Similar overheads to example applications

We also investigated GPU overheads

Confidential LLM Inference: Performance and Cost Across CPU and GPU TEEs

Marcin Chrapek, Marcin Copik, Etienne Mettaz, Torsten Hoefler

Conclusions



More of SPCL's research:

youtube.com/@spcl — 210+ Talks

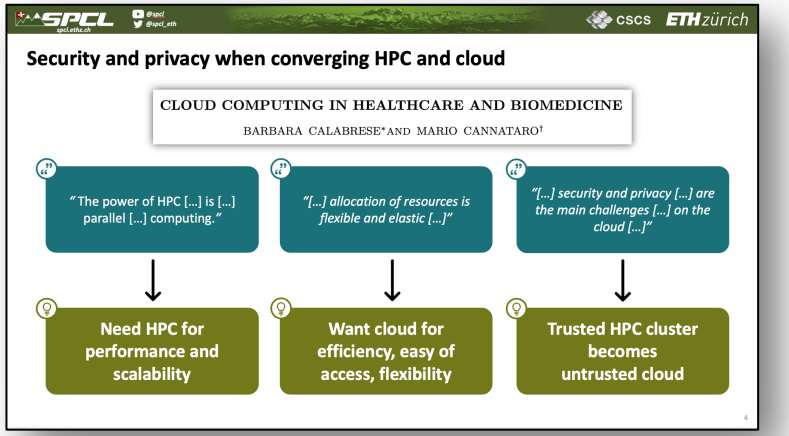
twitter.com/spcl_eth — 1.6K+ Followers

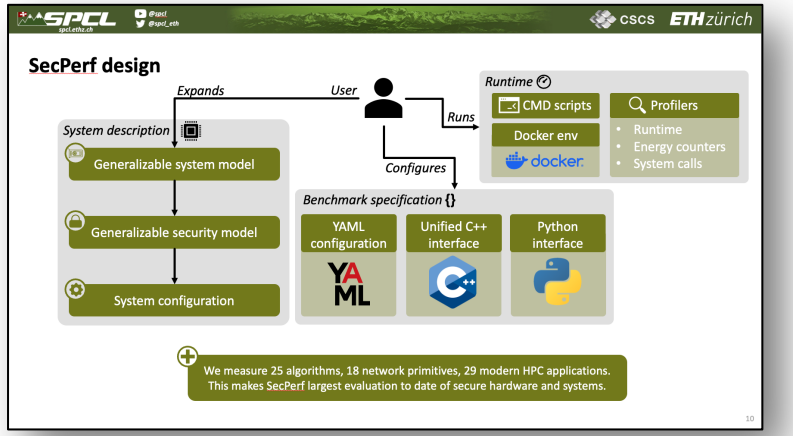
github.com/spcl — 5.6K+ Stars

… or spcl.ethz.ch

github.com/spcl/secperf

Conclusions





More of SPCL's research:

youtube.com/@spcl — 210+ Talks

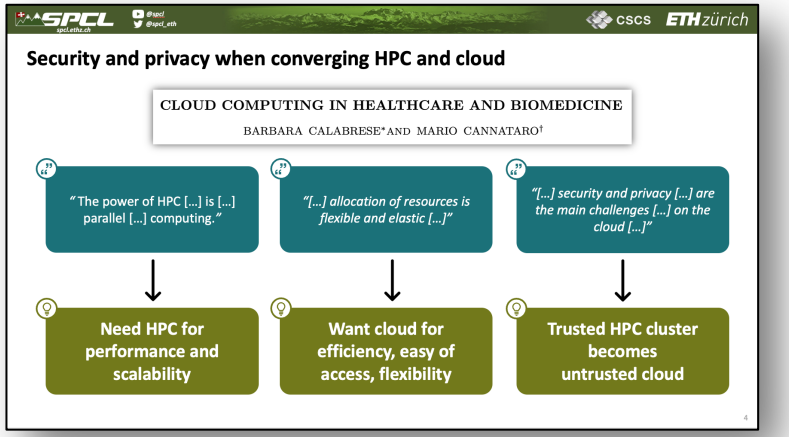
twitter.com/spcl_eth — 1.6K+ Followers

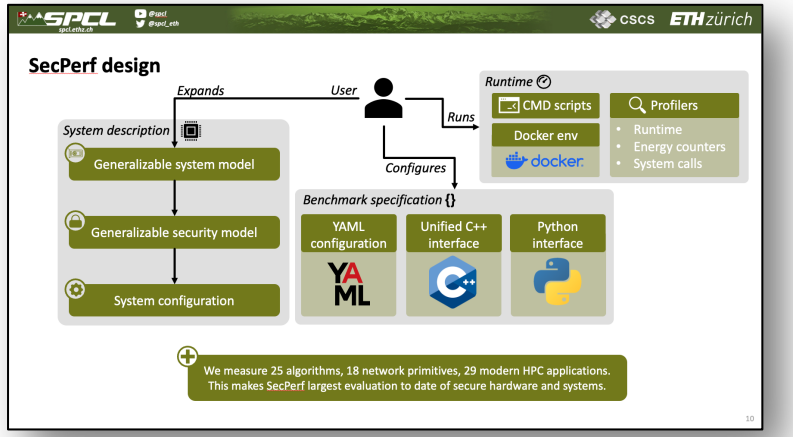
github.com/spcl — 5.6K+ Stars

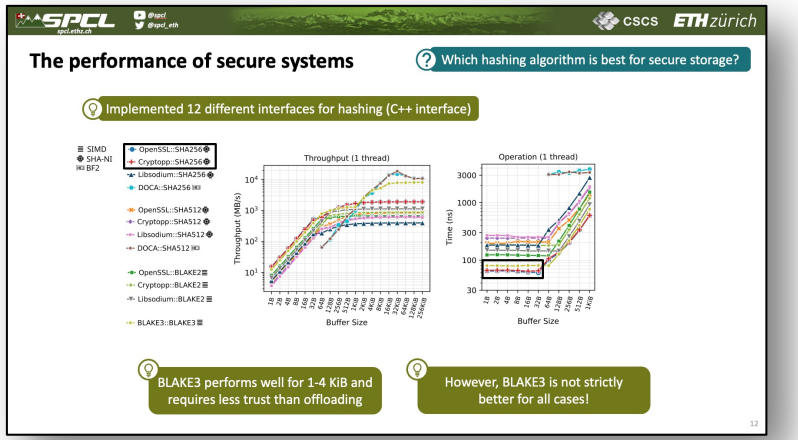
… or spcl.ethz.ch

github.com/spcl/secperf

Conclusions







More of SPCL's research:

youtube.com/@spcl — 210+ Talks

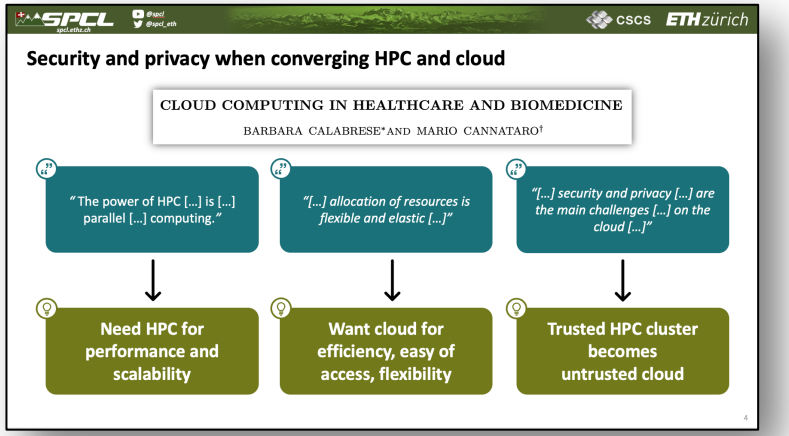
twitter.com/spcl_eth — 1.6K+ Followers

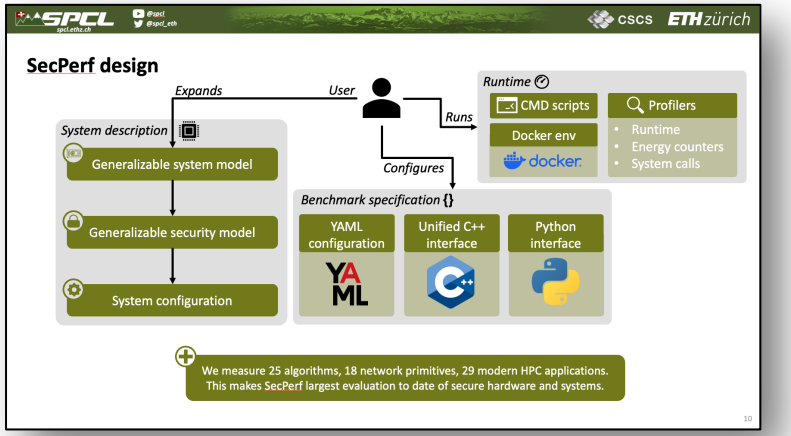
github.com/spcl — 5.6K+ Stars

… or spcl.ethz.ch

github.com/spcl/secperf

Conclusions





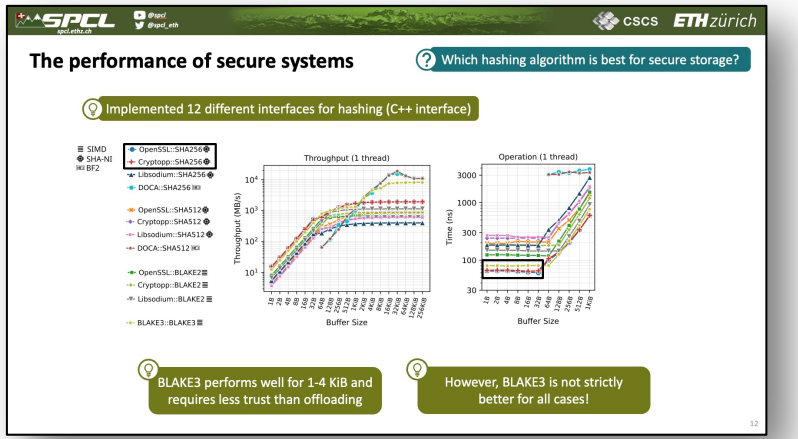
More of SPCL's research:

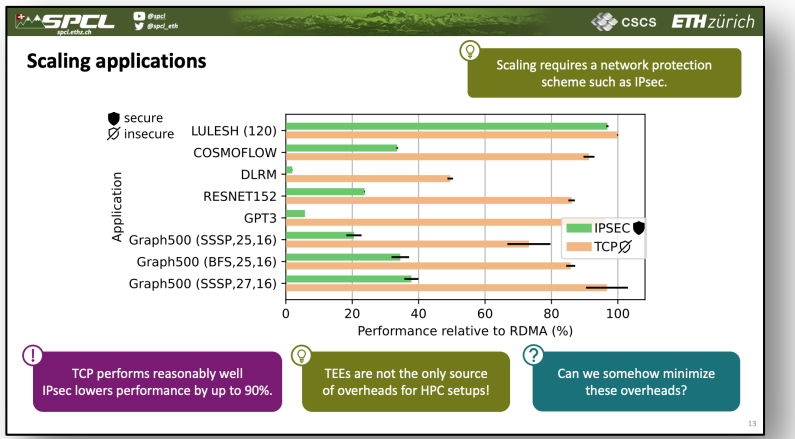
youtube.com/@spcl — 210+ Talks

twitter.com/spcl_eth — 1.6K+ Followers

github.com/spcl — 5.6K+ Stars

… or spcl.ethz.ch





github.com/spcl/secperf