Future Computing @ LRZ

June 8, 2021 | Josef Weidendorfer

Motivation

How can we improve and extend our compute services in the presence of

- Hardware specialization & divergence of programming models
- New workloads & user expectations

Scope

Focus on HPC systems, but keeping full service portfolio in mind

Challenges Ahead

Energy Consumption of ExaScale class systems

Heterogeneous: Integration of accelerators

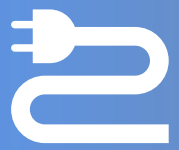

Programmability

Language (extension) of vendor X + vendor Y + vendor Z …

How to avoid porting efforts for our users?

(Answer: standards)

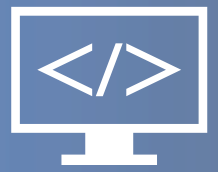

Usability

Researchers are best served when able to focus on own research

Future Systems

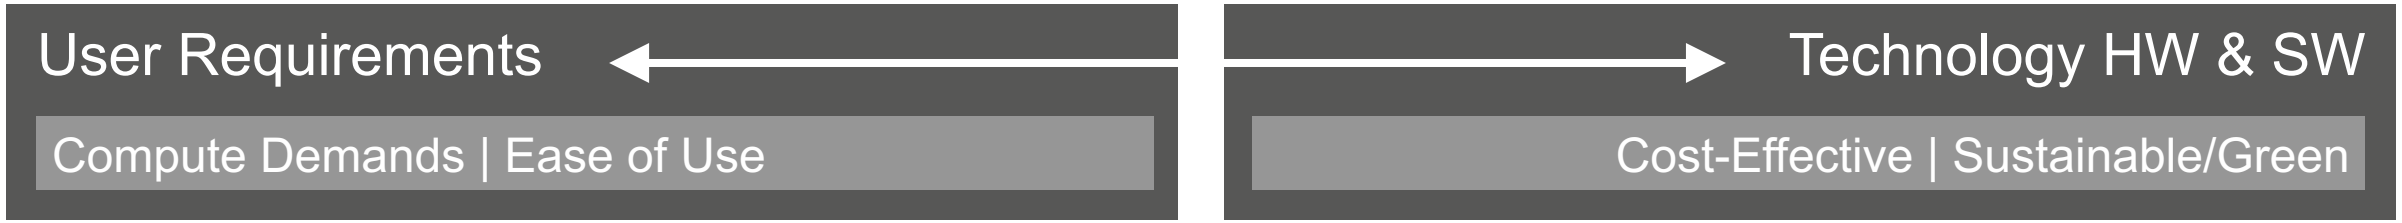

Aspects

User Requirements ← → **Technology HW & SW**

Compute Demands | Ease of Use

Cost-Effective | Sustainable/Green

Usage Models

Batch vs. interactive | Scalable vs. job farming

Hardware

Compute | Acc. | Memory | Network | Storage

Programming Models

Language | Library | DSL

System Software

Monitoring | Scheduling | Energy Awareness

Domain

HPC | Data Analytics | AI | QC Sim

Infrastructure

Cooling | Power Delivery

The Role of the FC Group

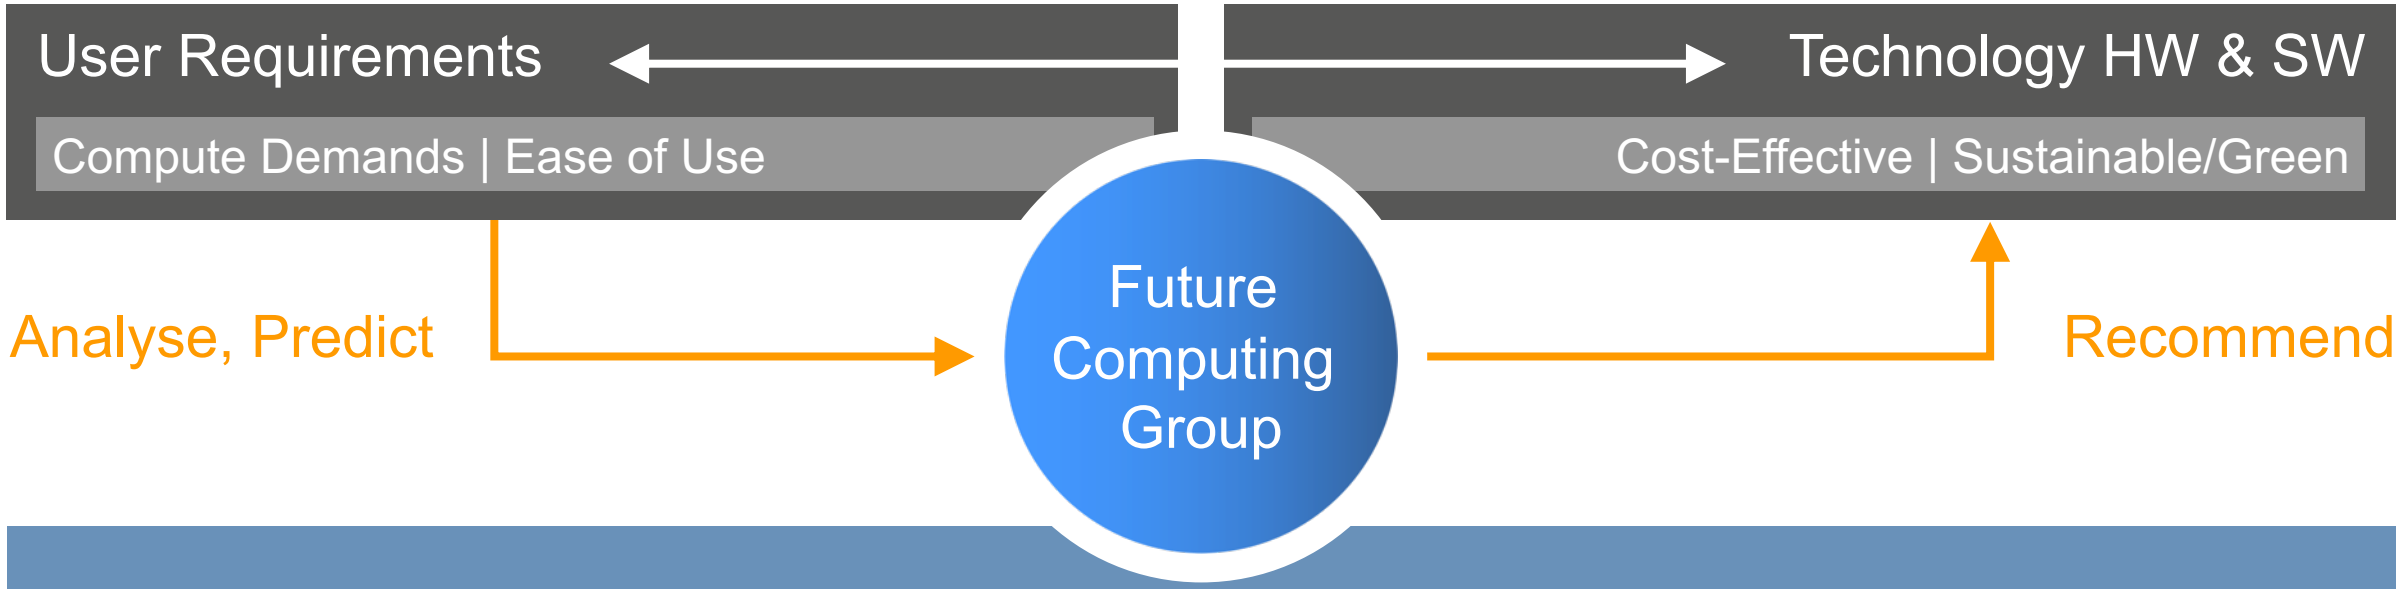

Understand best options – not just for the next system

Recommendations internally (for system purchase and operation) and externally (for supporting LRZ users)

Opportunity / Challange: Specialization

User Requirements ⟷ Technology HW & SW	
Compute Demands | Ease of Use	Cost-Effective | Sustainable/Green
Huge performance benefits possible ★★★	Enable ExaScale at reasonable Power ★★★
Restricted to some domains ?	Large design space: focus on what? ?
Often tradeoff: portability vs. efficiency ??	Increased complexity: heterogeneity ??

Opportunity / Challenge: New Workloads (eg. AI) + Expectations

User Requirements ⟷ Technology HW & SW	
Compute Demands | Ease of Use	Cost-Effective | Sustainable/Green
Opportunity: New User Classes ★★★	Often targeted by efficient frameworks ★★★
Users not used to SSH/Terminal ?	New tools to enable interactivity + steering ?
Demand for changing SW stack ??	Complexity of Virtualization Solutions ??

Future Systems: How?

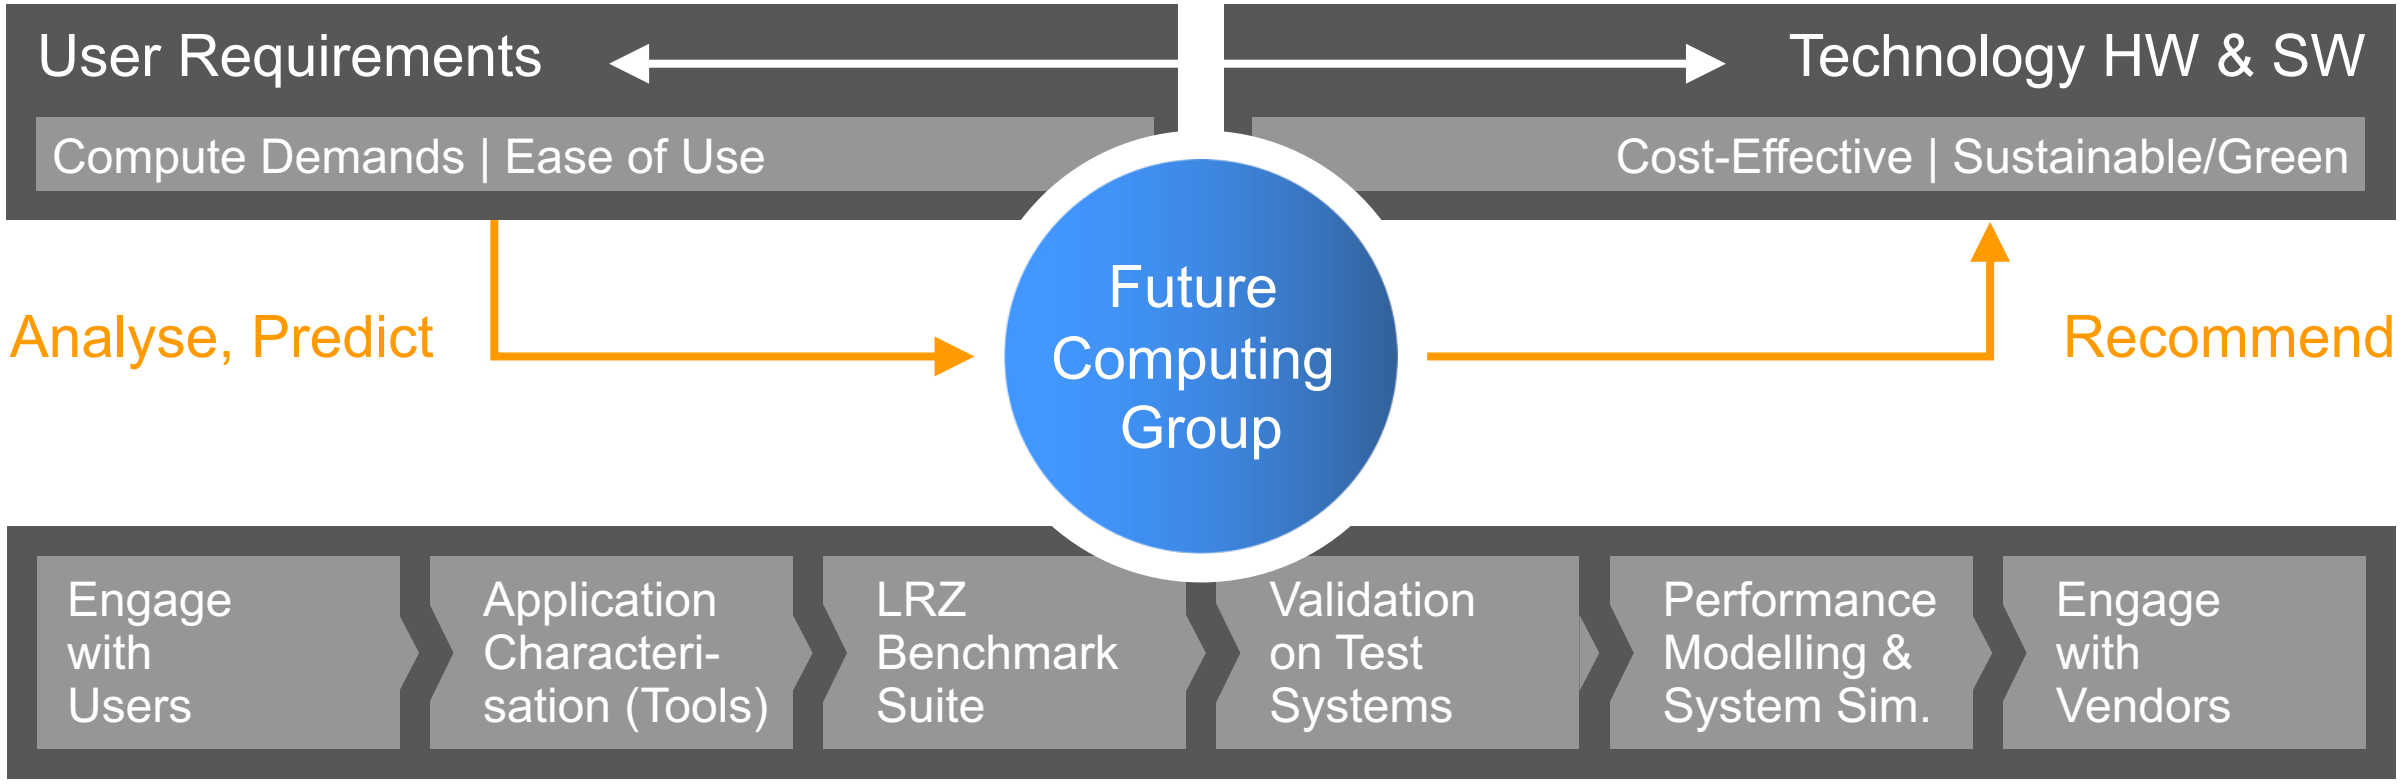

Future Systems: How?

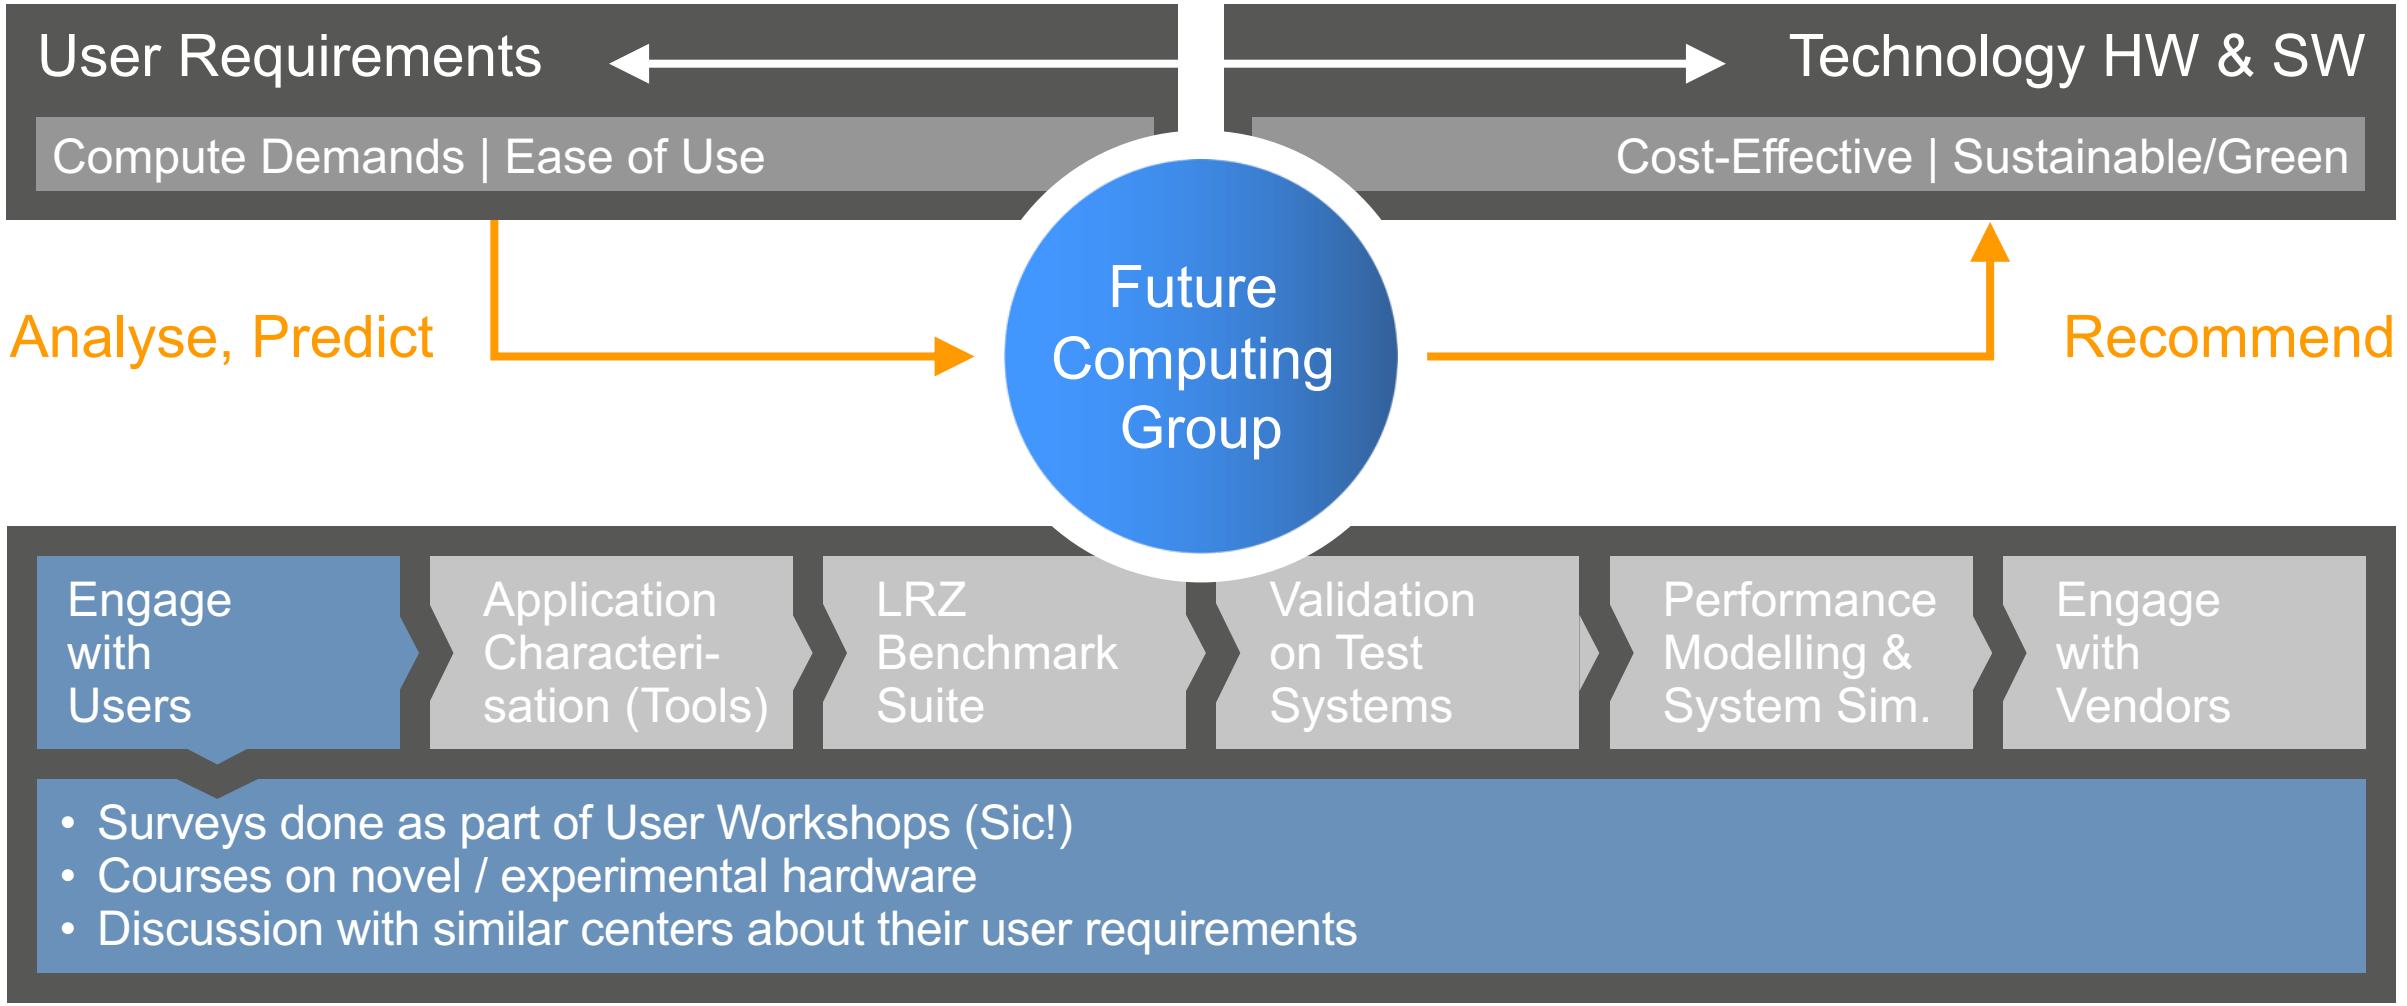

Future Systems: How?

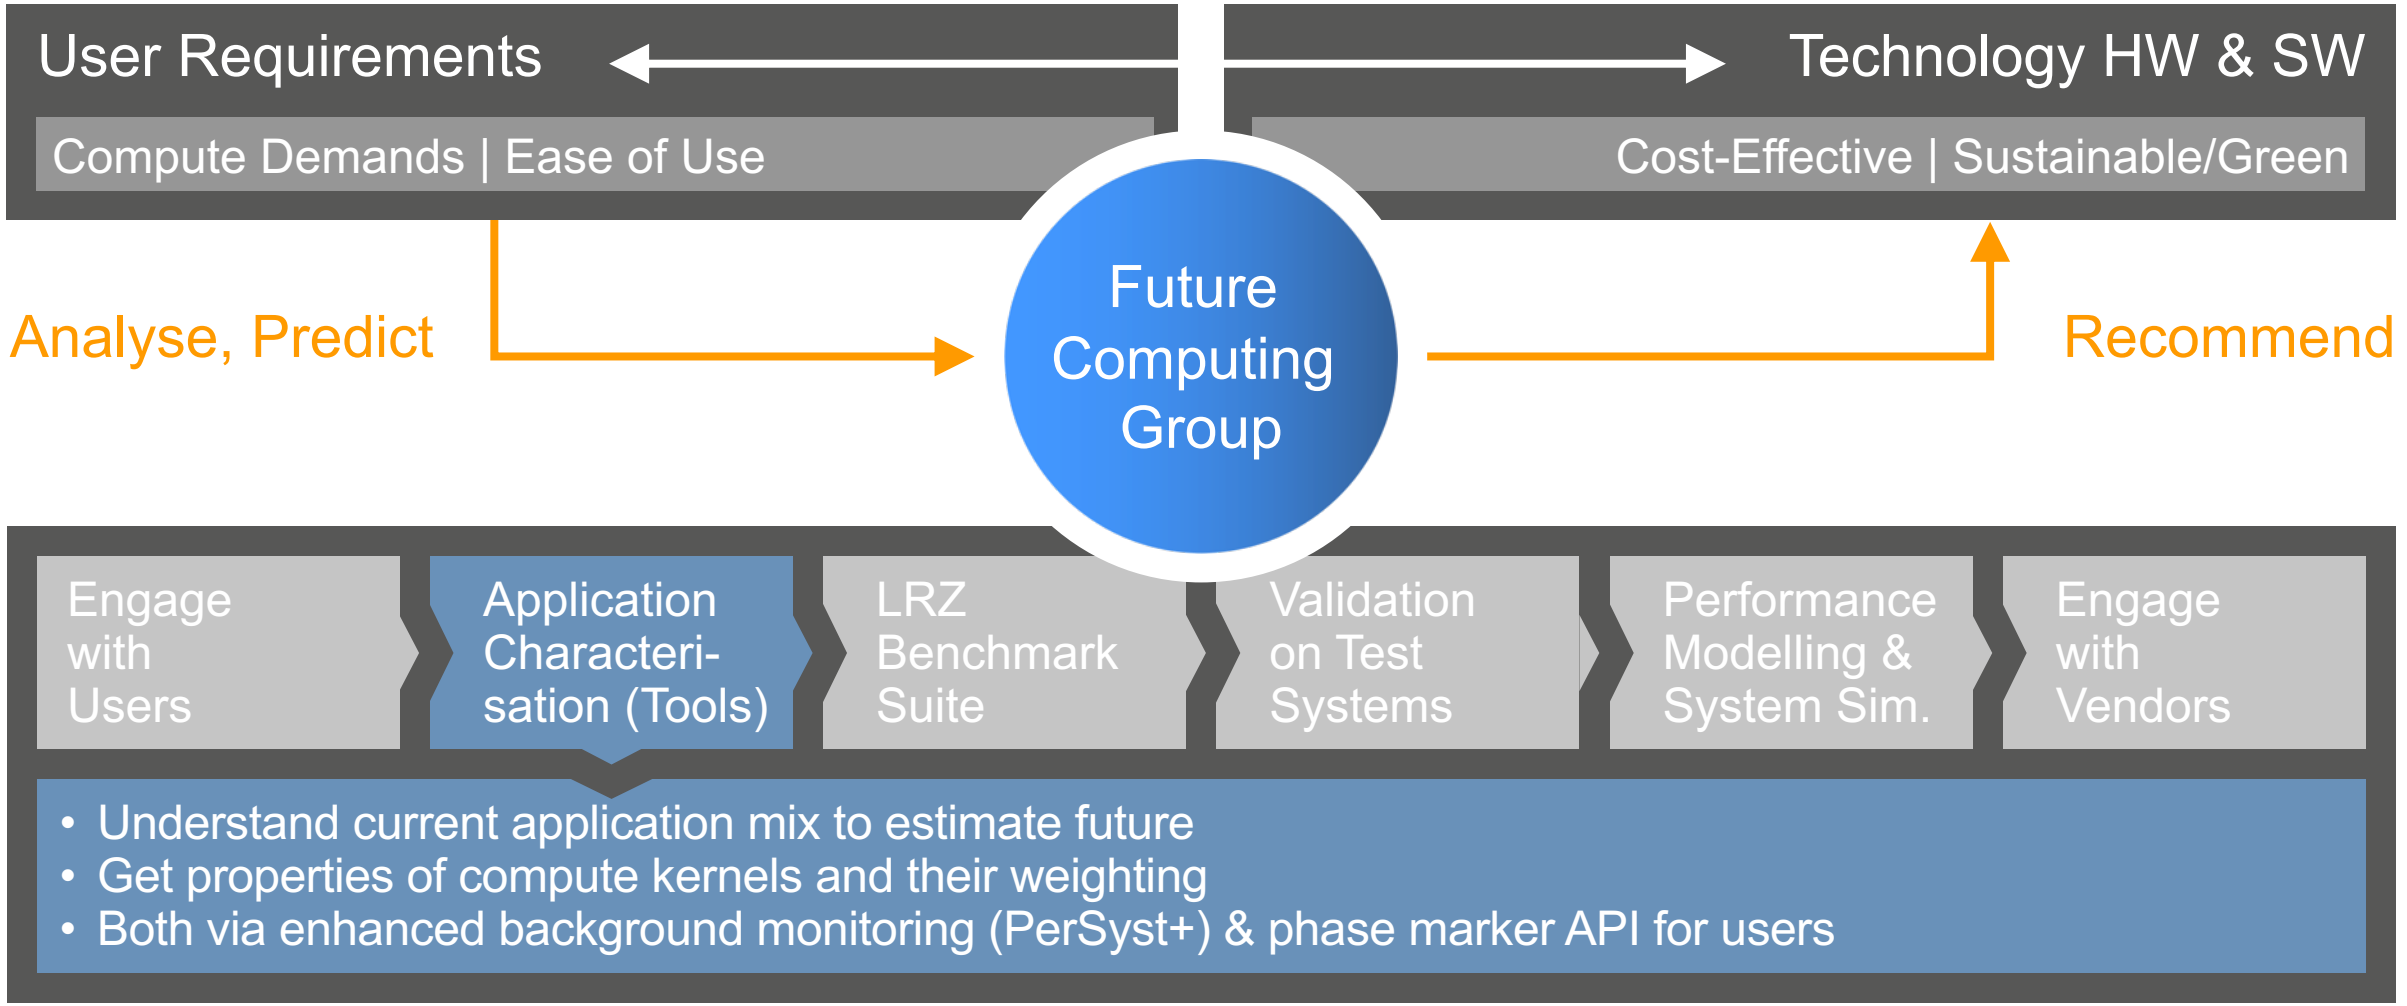

Future Systems: How?

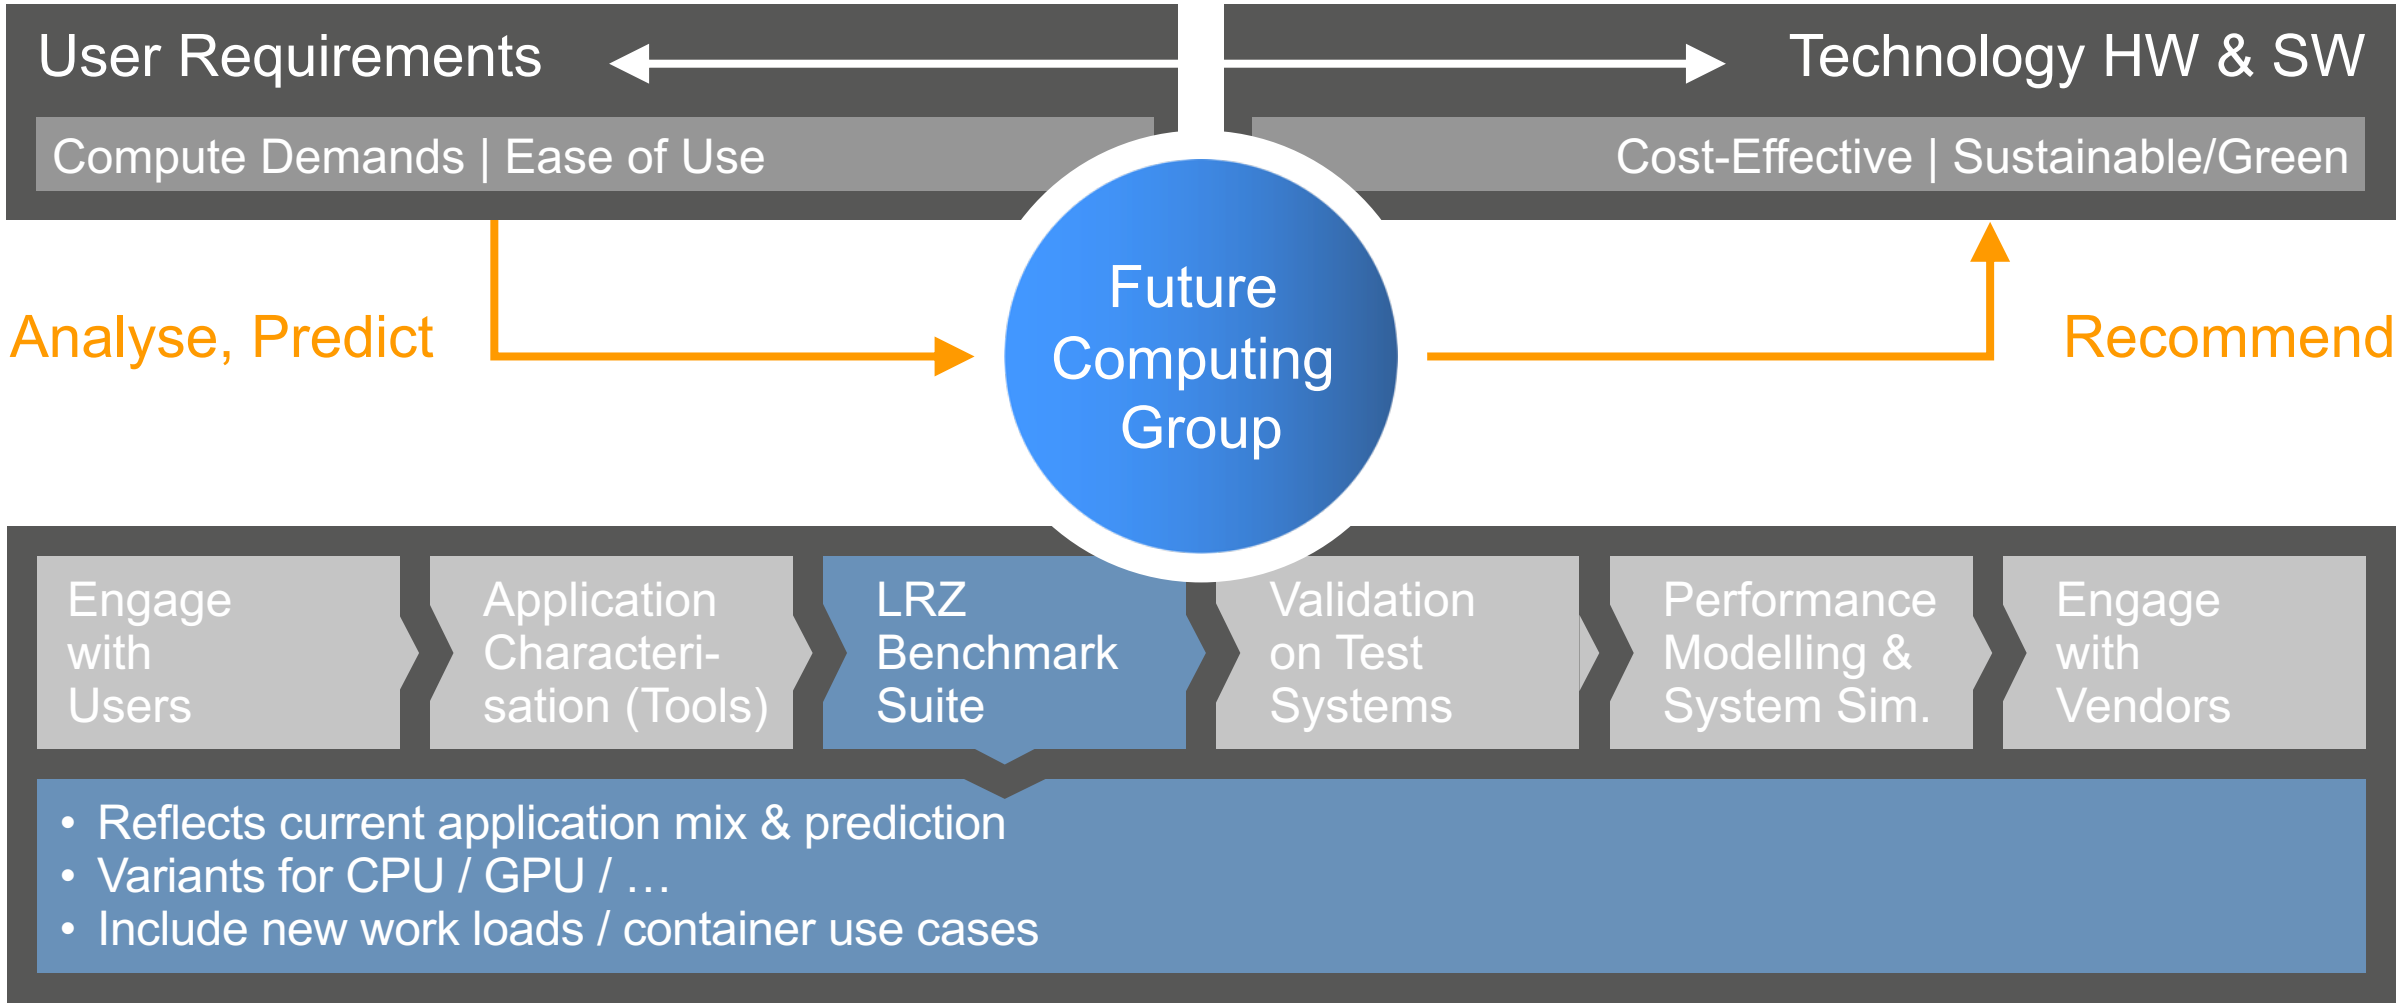

Future Systems: How?

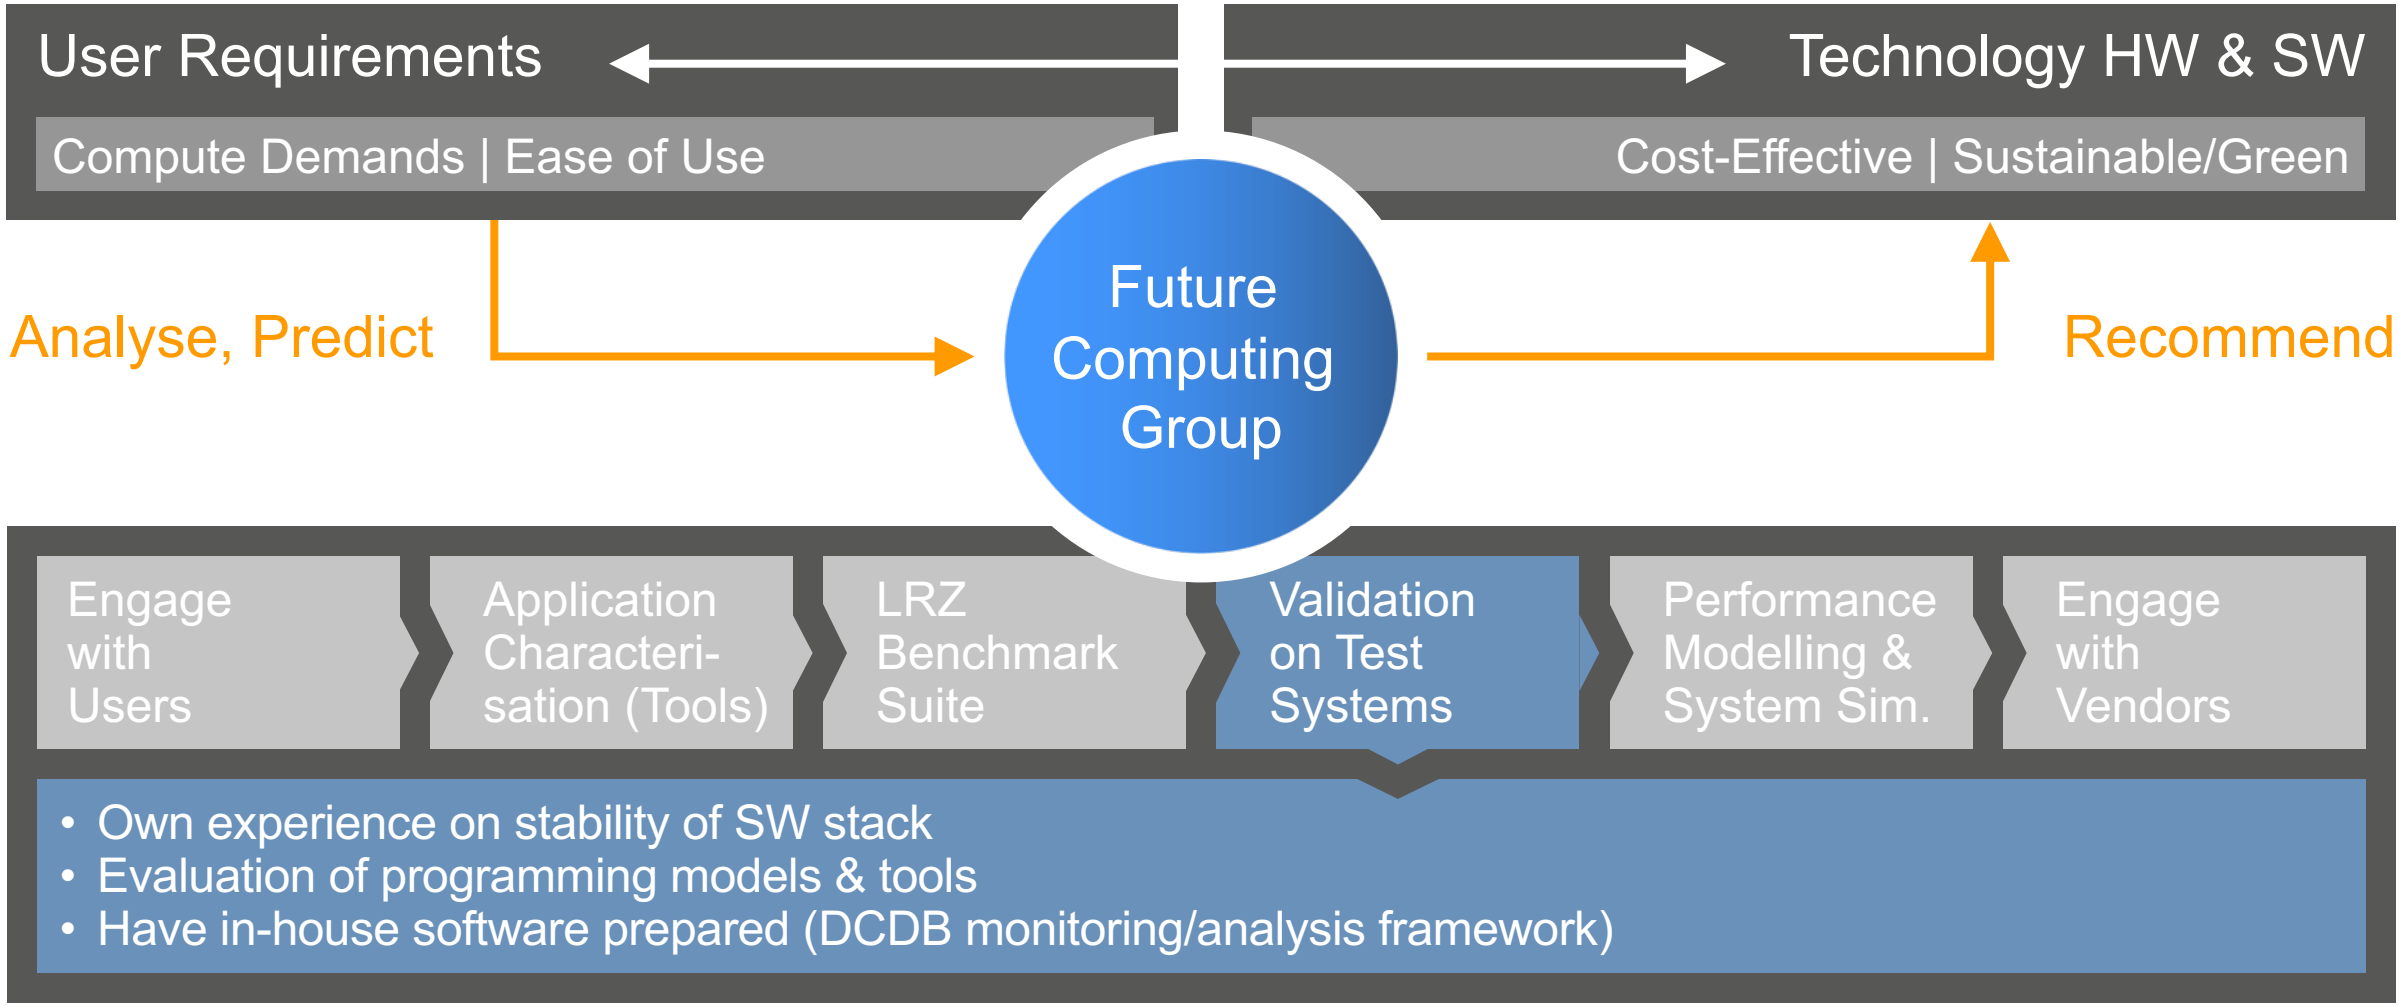

BEAST

"Bavarian Energy, Architecture, Software Testbed"

BEAST

Systems/Usage

Test Systems

V1
AMD Rome / MI-50
Th.X2 / Nvidia V-100

V2
HPE CS500 A64FX

V3
DAOS, CooperLake,
IceLake / Intel GPU

V4
AI Accelerator

Usage

Internal (no LRZ service)

- To benchmark codes on architecture options for future deployments
- Do LRZ SW adaptation

Selected partners

- Experience exchange

Access Modes

- Enable system experiments as root (custom installations)
- Run benchmark bundles (without account)

Future Systems: How?

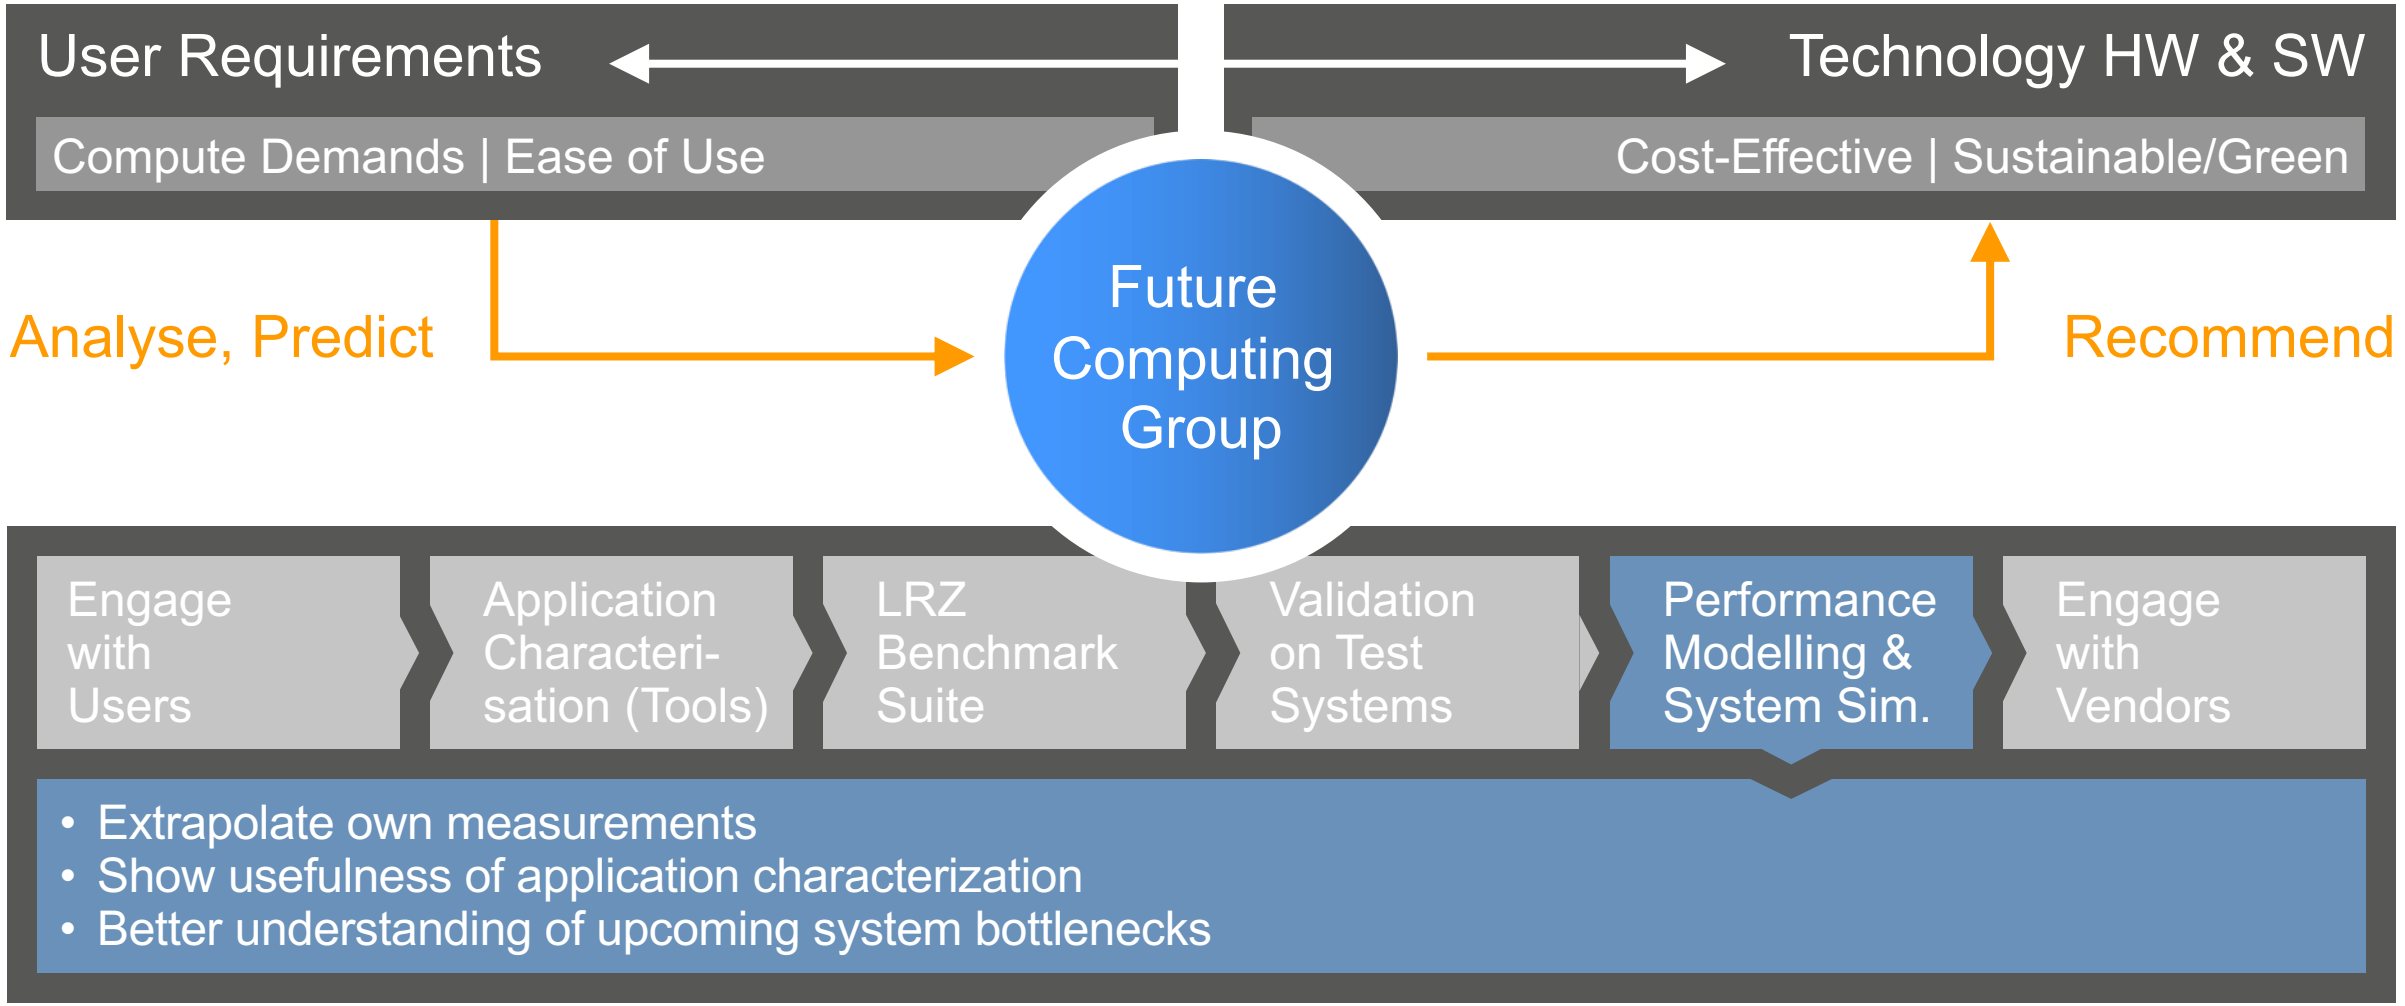

Future Systems: How?

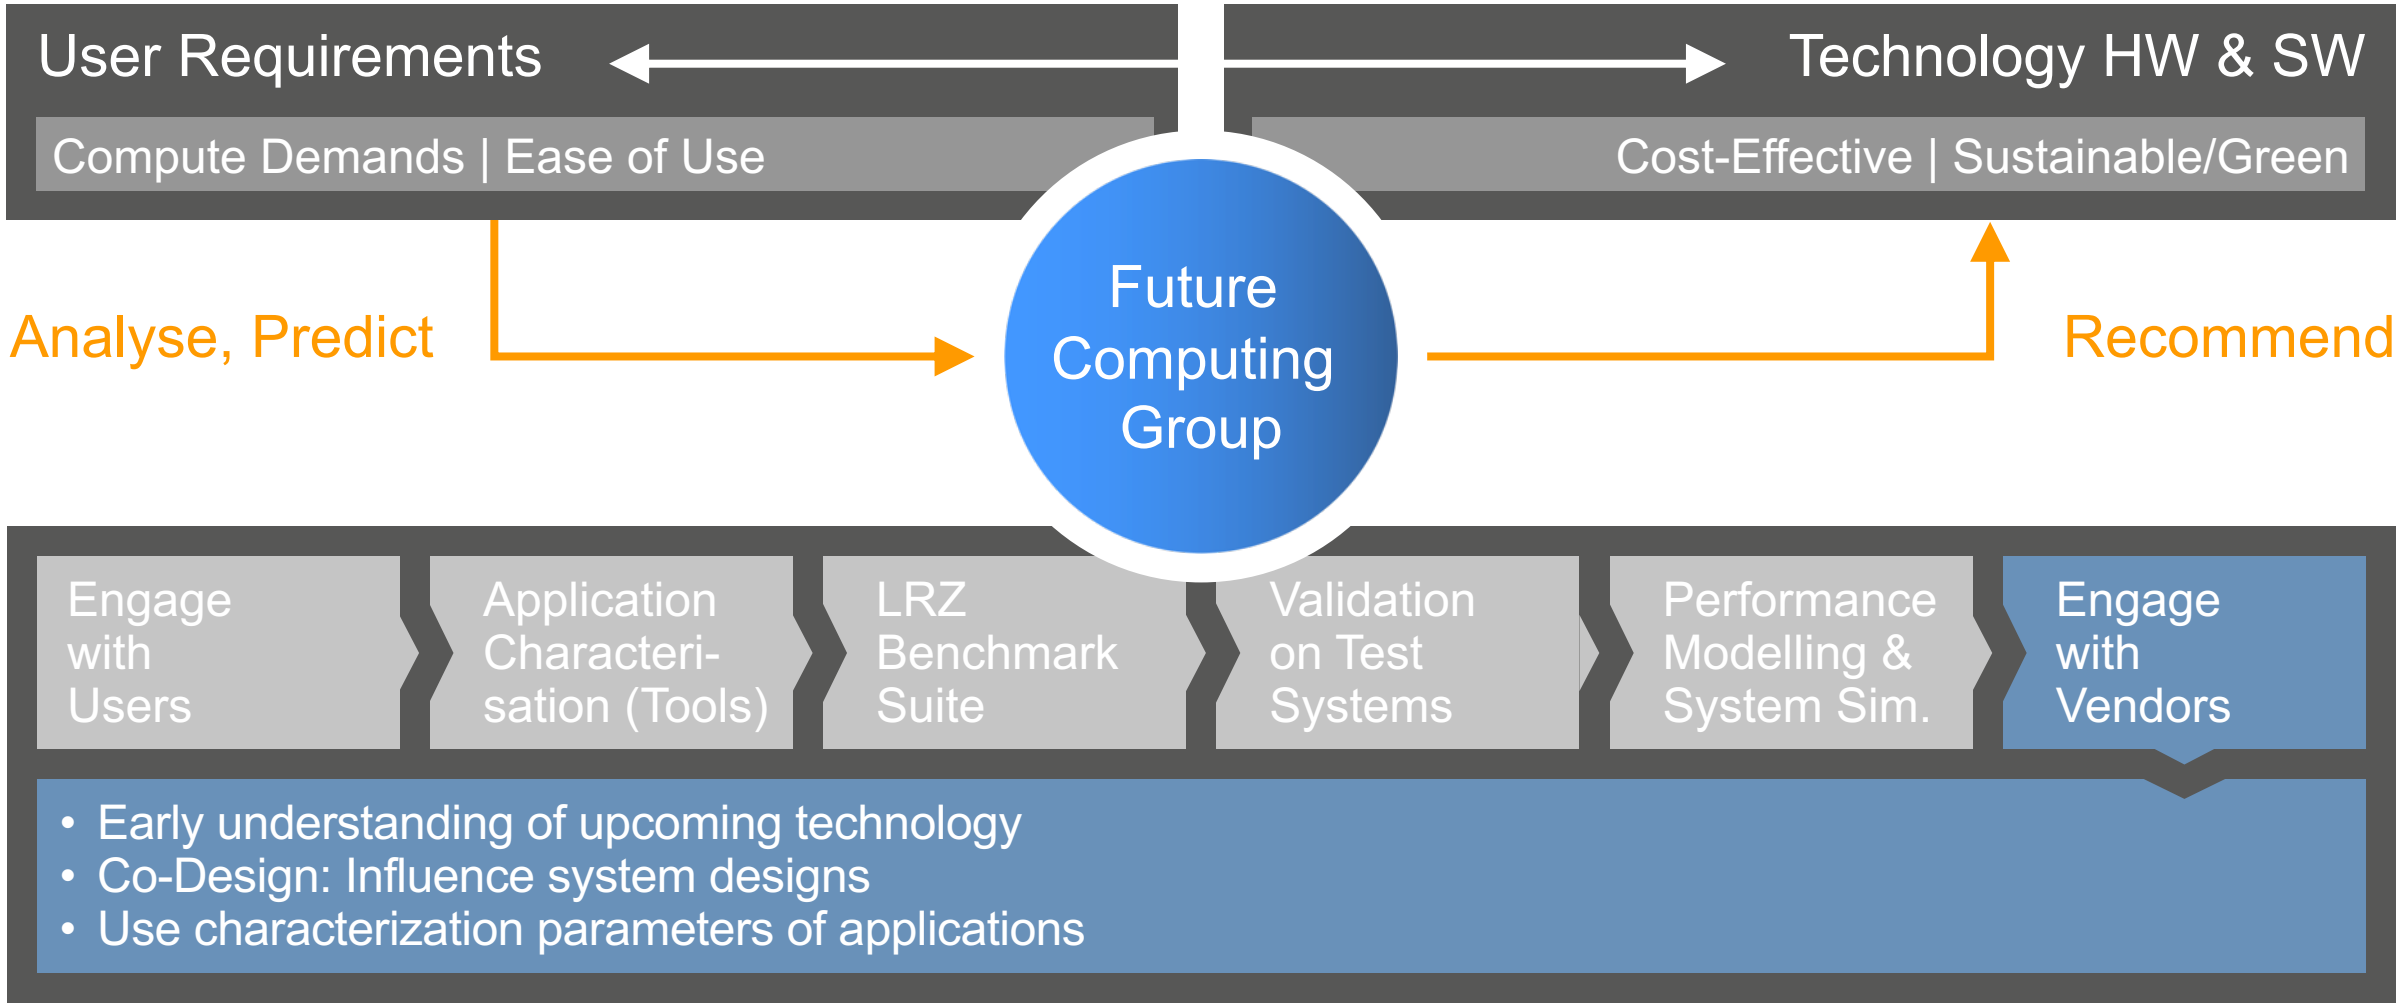

Help Us – You can shape the future LRZ Leadership Class System!

- Please take part in our User Survey

- Contact us if you want to share more details
- Your code may become part of the official LRZ Application Benchmark

futurecomputing@lrz.de

Leibniz Supercomputing Centre

of the Bavarian Academy of Sciences and Humanities