


Building the Foundation for Exascale Computing



2 INTEL XEON SCALABLE PROCESSORS

“Sapphire Rapids”

6 X^E ARCHITECTURE BASED GPU'S

“Ponte Vecchio”

ONEAPI

Unified programming model

LEADERSHIP PERFORMANCE

For HPC, data analytics, AI

UNIFIED MEMORY ARCHITECTURE

Across CPU & GPU

ALL-TO-ALL CONNECTIVITY WITHIN NODE

Low latency, high bandwidth

UNPARALLELED I/O SCALABILITY ACROSS NODES

8 fabric endpoints per node, DAOS

DELIVERED IN 2021

