Chip Design on Cloud - from Pipe dream to Preeminence?

Gartner predicts that by 2025, about 80% of the datacenters will move to the cloud. Inherent benefits of the cloud, such as elasticity, fault tolerance, and security cannot be matched by on premise data centers. Indeed, cloud is the datacenter of choice for several industries from retail to banking to manufacturing.

In addition to the widely known advantages of cloud, AI and AI-driven opportunities are other reasons to take a hard look at cloud. There has been a “cambrian” explosion in chips with the growing popularity of AI-enabled applications, and consequently, the number of design starts has grown. Many believe that for companies to stay competitive, chip designs cannot take the usual nine to twelve months - and cloud and data-driven flows will be crucial to scale quickly and efficiently. Further, most EDA tools embed ML-enabled engines which require GPUs for training. On paper - it appears that it will be challenging for on premise datacenters to meet the requirements of scalable, compute-intensive, AI-driven chip design flows.

Chip design and verification on cloud has been a topic for decades (many among us remember the efforts of EDA companies in the early 2000s). However, the cloud - as it is today - is stunningly different. But what is the state of chip design on cloud? Are companies designing successfully on the cloud? Are migration efforts underway at companies? What works? What doesn’t?

Is our industry inherently different in some way with regards to infrastructure needs? What are the lessons from our previous efforts? Hear from a panel of designers and infrastructure experts on their experience with cloud. Learn about what works, what doesn’t work or what cannot work!

Organizer: Sashi Obilisetty, Chief Architect, Silicon Solutions, Google Cloud

Moderator: Ann Mutschler, Executive Editor/EDA Semiconductor Engineering

Panelists:

Google: Richard Ho, Principal Engineer, Google
SiFive: Megan Wachs, VP of Engineering
Synopsys: Bob Lefferts, Synopsys Fellow
AMD: Eric Chesters, AMD Fellow

Ann Mutschler
Executive Editor/EDA at Semiconductor Engineering

Ann Mutschler has been an editor and journalist for more than 25 years focused on the semiconductor industry. She is currently Executive Editor/EDA at Semiconductor Engineering, the premier semiconductor industry magazine that provides deep insights into the increasingly complex task of designing, testing, verifying, integrating and manufacturing semiconductors, as well as insights into the market dynamics that make it all possible. She focuses on all aspects of the semiconductor design process with particular interest in the automotive segment as well as system design, and power/performance sensitivity in leading-edge applications. She holds a Bachelor of Arts Degree in Communications Studies and Journalism from Sonoma State University.

Megan Wachs
VP, Engineering, SiFive

SiFive's Core Designer product (https://www.sifive.com/core-designer) allows customers to customize and generate RISC-V cores and collateral on the cloud, without any local EDA tooling. Megan has dealt with the opportunities and challenges from the inside when creating a platform, and brings a perspective on ways that increasing tools available in a cloud environment could improve this experience for users. Megan has a Bachelor's degree from Brown University and a Masters and Phd in Electrical Engineering from Stanford University.

Bob Lefferts
Fellow, Synopsys

Dr. Bob Lefferts has a Ph.D. from Stanford University in Semiconductor Device Physics and 40 years of experience in the Semiconductor industry. Bob has been at Synopsys since 2004 and held positions managing a high speed SERDES IP team for his first 6 years and then managed Synopsys' CAD & Design enablement team for the past 10 years. Bob was recently promoted to a Synopsys Fellow and is currently engaged in working with Synopsys' full mixed-signal IP development design environment in the cloud for flexible IP development.

Eric Chesters
Fellow, AMD

Eric joined AMD 3 years ago and has spent the last couple of years working as a member of the Verification Methodology and Tools team. In this capacity he has been involved with driving a wide range of initiatives with a particular focus on improving the efficiency of AMDs front end development tools and flows. Recently Eric's focus has shifted to the challenges associated with transitioning AMDs legacy flows from the private data center into the cloud. Prior to joining AMD Eric spent 10+ years defining, architecting and verifying high end network processors for Cisco Systems.

Richard Ho
Principal Engineer, Google

Richard received his Ph.D. in Computer Science from Stanford University. He was a co-founder of 0-In Design Automation and an early technologist involved with assertions, formal property verification and CDC verification. Richard was also part of the team that built the Anton and Anton2 supercomputers for molecular dynamics simulation at D. E. Shaw Research. He is currently a Principal Hardware Engineer with the datacenter chip team at Google working on TPU and other cool projects.