



# Cloud & AI Technologies for Faster, Secure Semiconductor Supply Chains

Jay Lewis

GM, Incubation and Innovation

CAD for Hardware Security Workshop (CAD4Sec)

Design Automation Conference (DAC) #59

July 10, 2022



DALL·E 2

## Image Generation from Natural Language

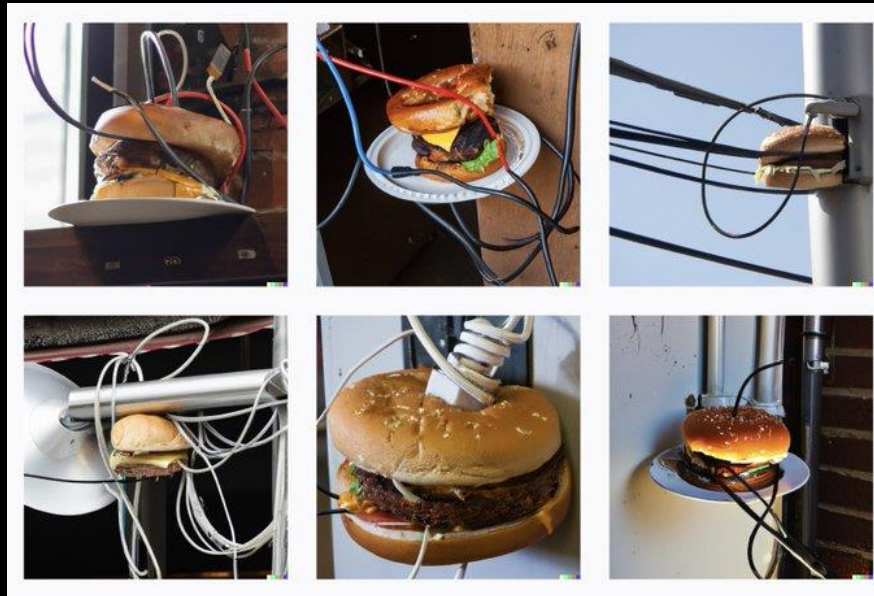
A bowl of soup  
that looks like a monster  
knitted out of wool

# DALL·E 2

## Image Generation from Natural Language



“stained glasses”



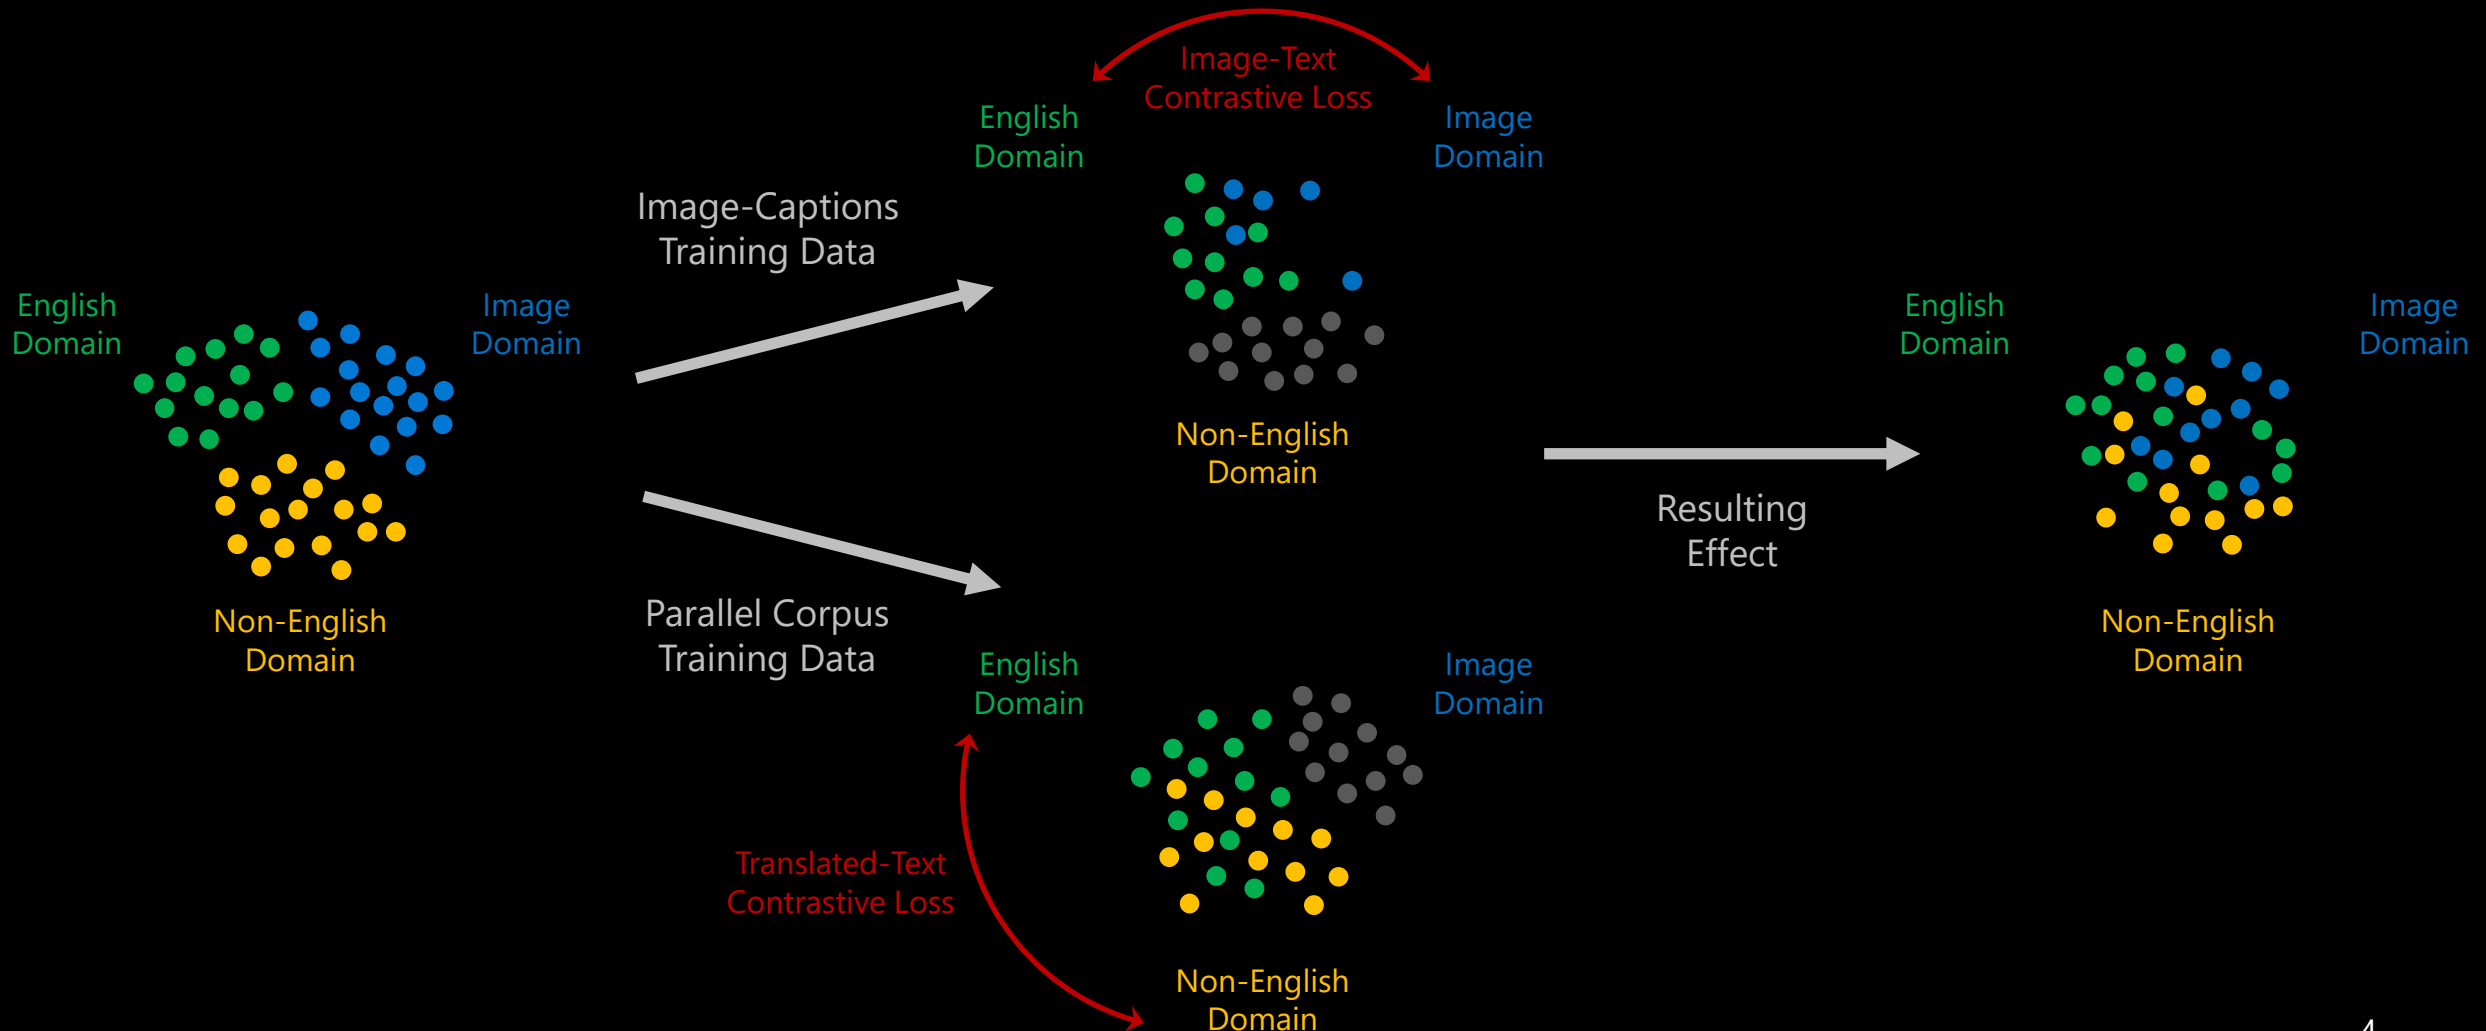
“a hamburger with incredibly unsafe electrical wiring”



“a photo of a VW Beetle parked on the street in Brooklyn NY in front of a brownstone, taken with a Pentax K1000, lens flare”

# Turing Bletchley

## A Universal Image Language Representation Model



# Solving a Riddle

## Megatron-Turing Model

All of us have one, but few get to choose.  
If you don't know mine, you are not my friend.  
When it is called out, attention is drawn.  
Even if we are gone, they are still around.

I think the answer is: A name

Everyone has a name, but not everyone gets to choose their own name.

If you don't know my name, you are not my friend.

When someone calls your name, you turn your attention to them.

Even if you are gone, your name is still around.

# Limitations of Machine Learning

## What I Used To Think

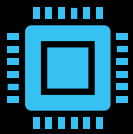
- ML is excellent at pattern recognition but lacks context, and bigger models don't solve this
- Lack of explainability will be a barrier for critical applications
- Fragility will be a barrier for critical applications
- ML implementations are proprietary and only available to a few massive companies
- AI models are not sentient  
(Just kidding)

# Complexity and Growth

## Natural and Artificial Intelligence

~100B

Transistor Count  
on an SoC



~100B

Neurons in the  
Human Brain



~200-1000T

Synapses in the  
Human Brain



~1.8T

Parameters in the  
largest AI Models



# Machine Learning Applied to EDA

Resource Optimization

Analog design Supply Chain

Result estimation

IR drop prediction

CoPilot for design

Power prediction

High level synthesis SAT solver

Verification and test

Lithography hotspot detection

Test set redundancy reduction

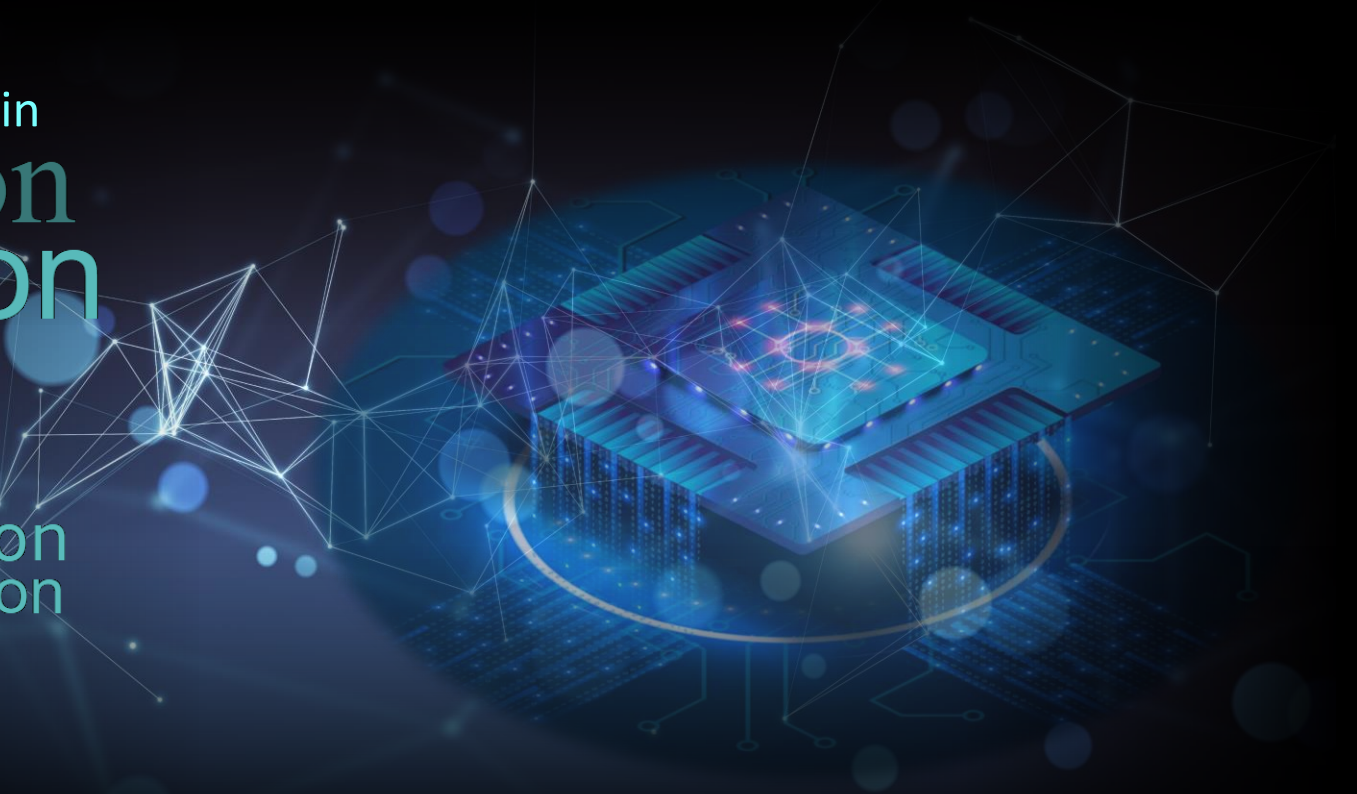
Design space exploration

Defect detection

Logic synthesis Mask synthesis

Place and route

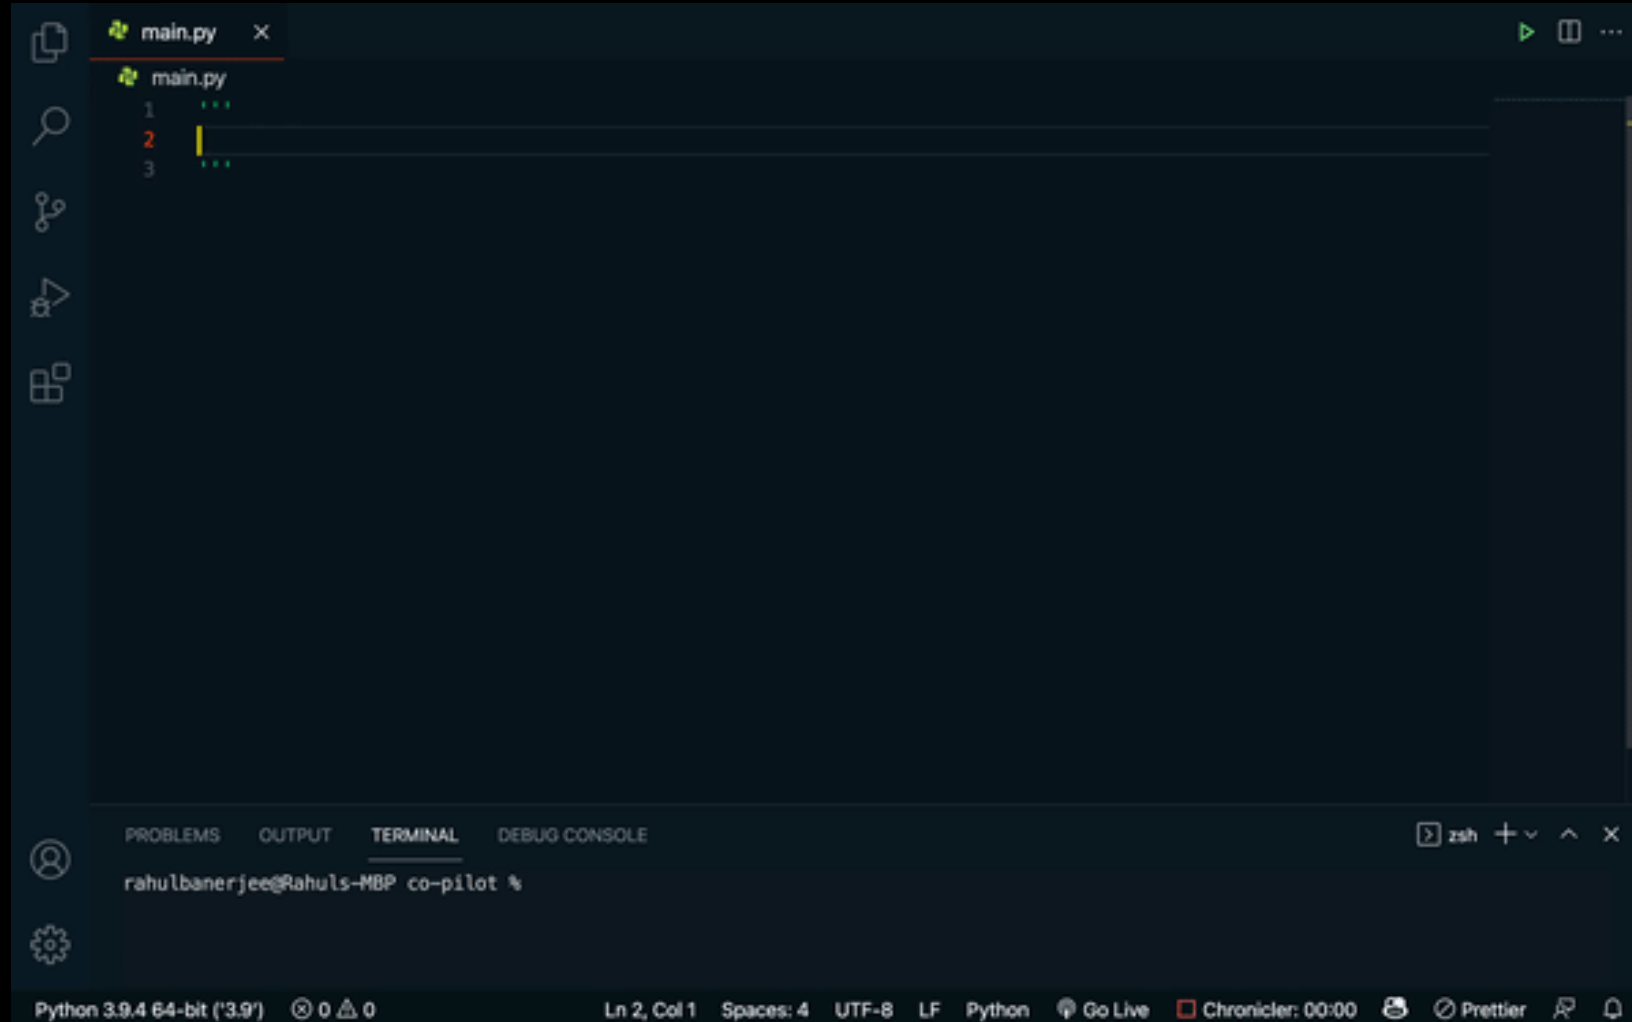
3D integration Security-aware





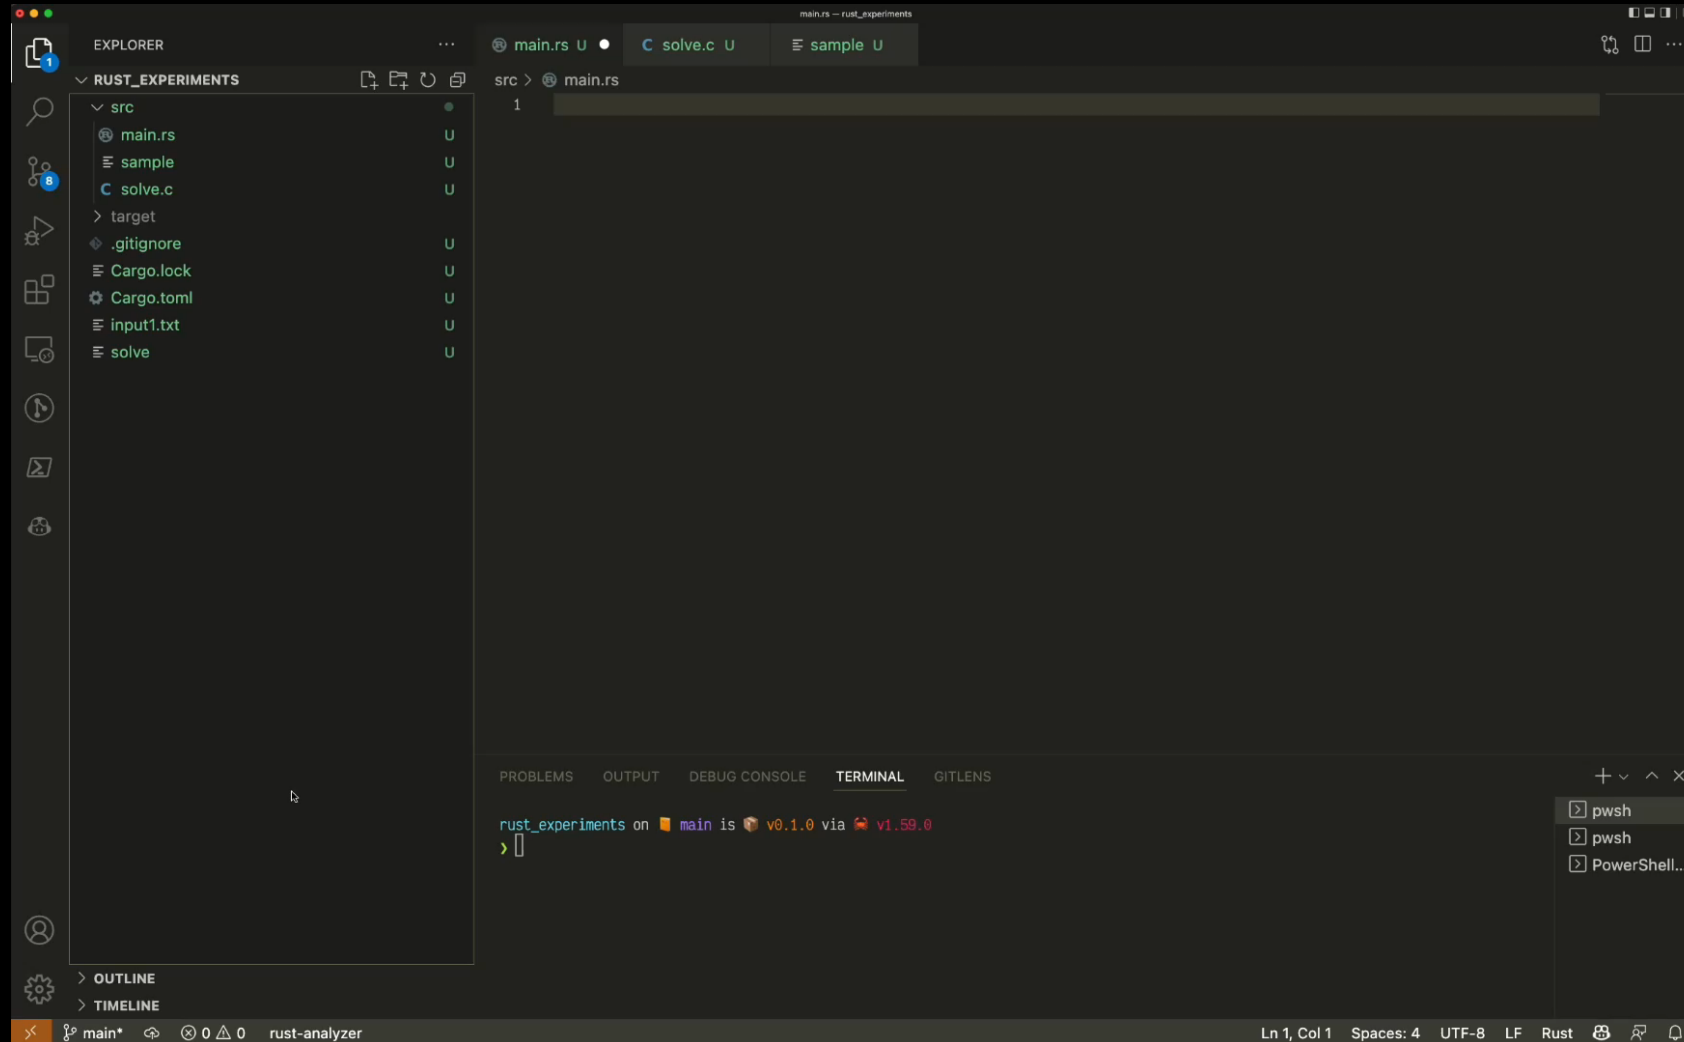
# GitHub Co-Pilot

## Advanced AI Pair Programmer



# GitHub Co-Pilot

## Language Translation: C to Rust



# Machine Learning Applied to EDA

Resource Optimization

Analog design Supply Chain

Result estimation  
IR drop prediction

CoPilot for design Power prediction

High level synthesis SAT solver

Verification and test

Lithography hotspot detection

Test set redundancy reduction

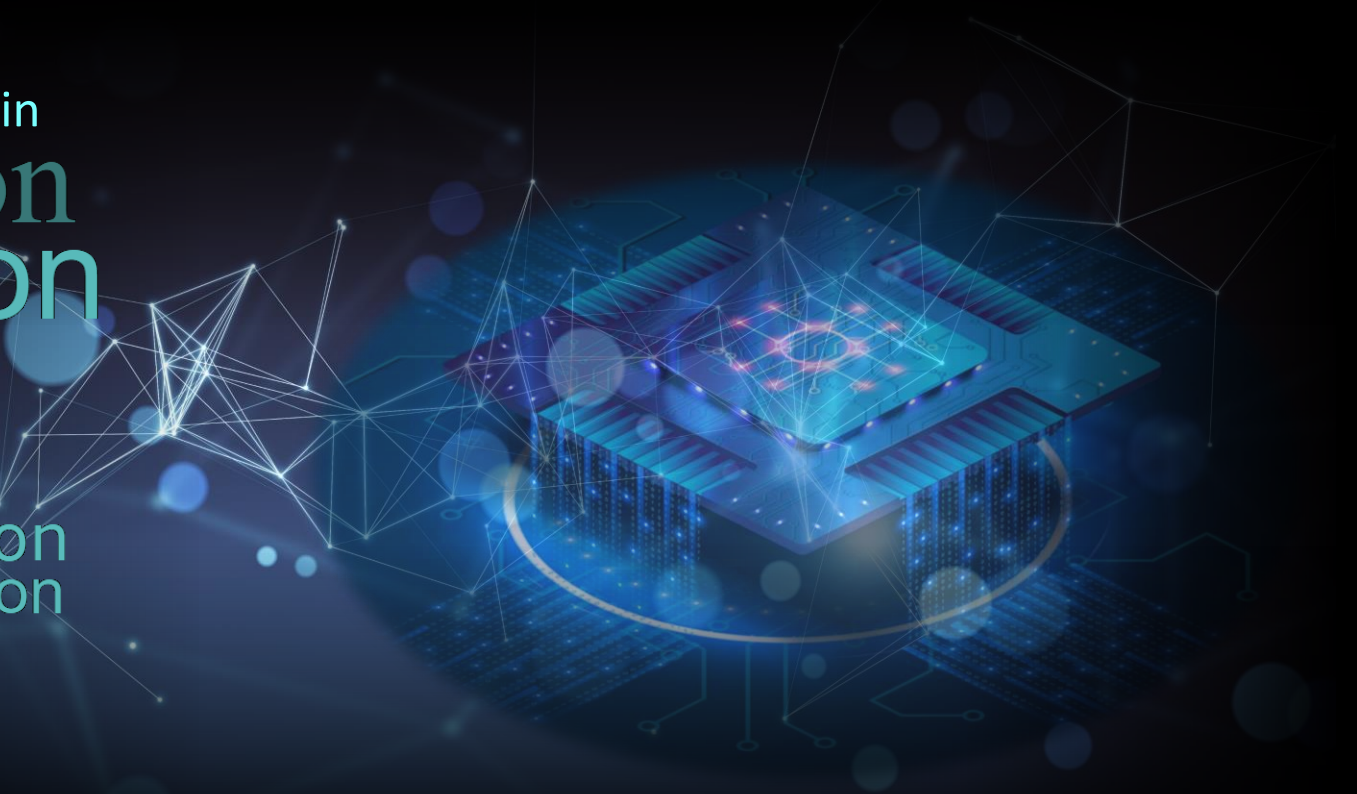
Design space exploration

Defect detection

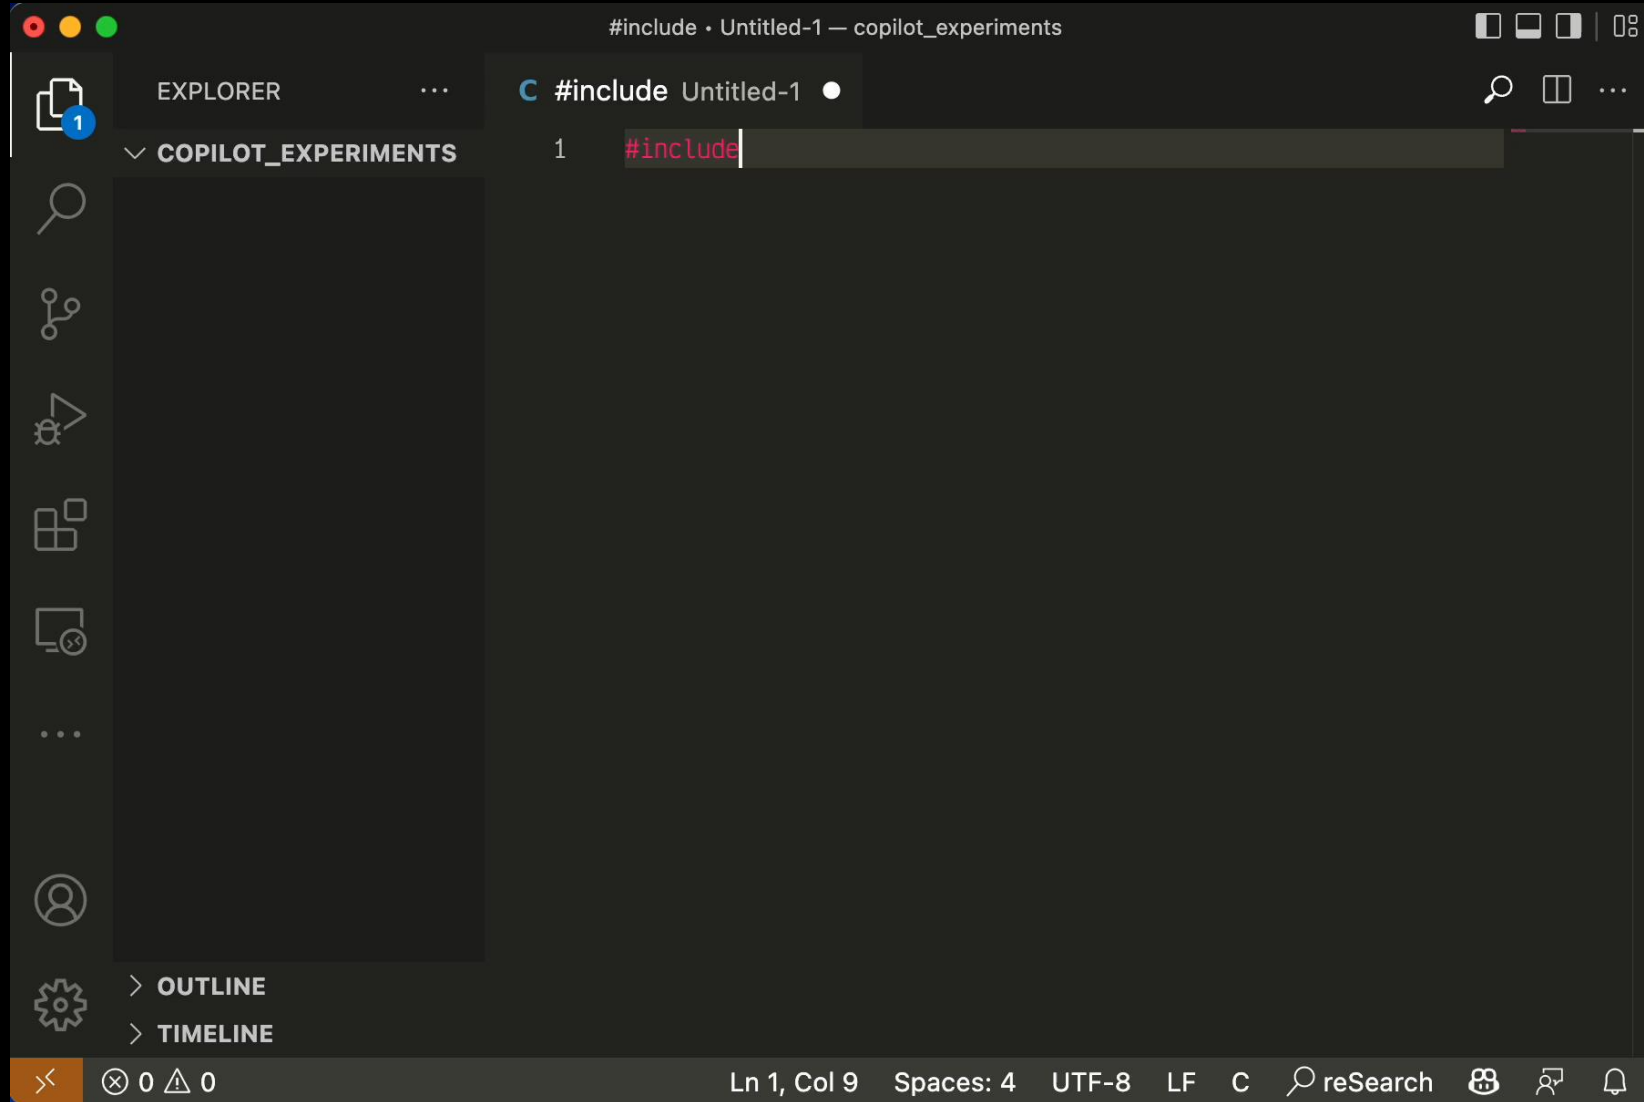
Logic synthesis Mask synthesis

Place and route

3D integration **Security-aware**



# Security Awareness



# Machine Learning Applied to EDA

Resource Optimization

Analog design Supply Chain

Result estimation

IR drop prediction

CoPilot for design Power prediction

High level synthesis SAT solver

Verification and test

Lithography hotspot detection

Test set redundancy reduction

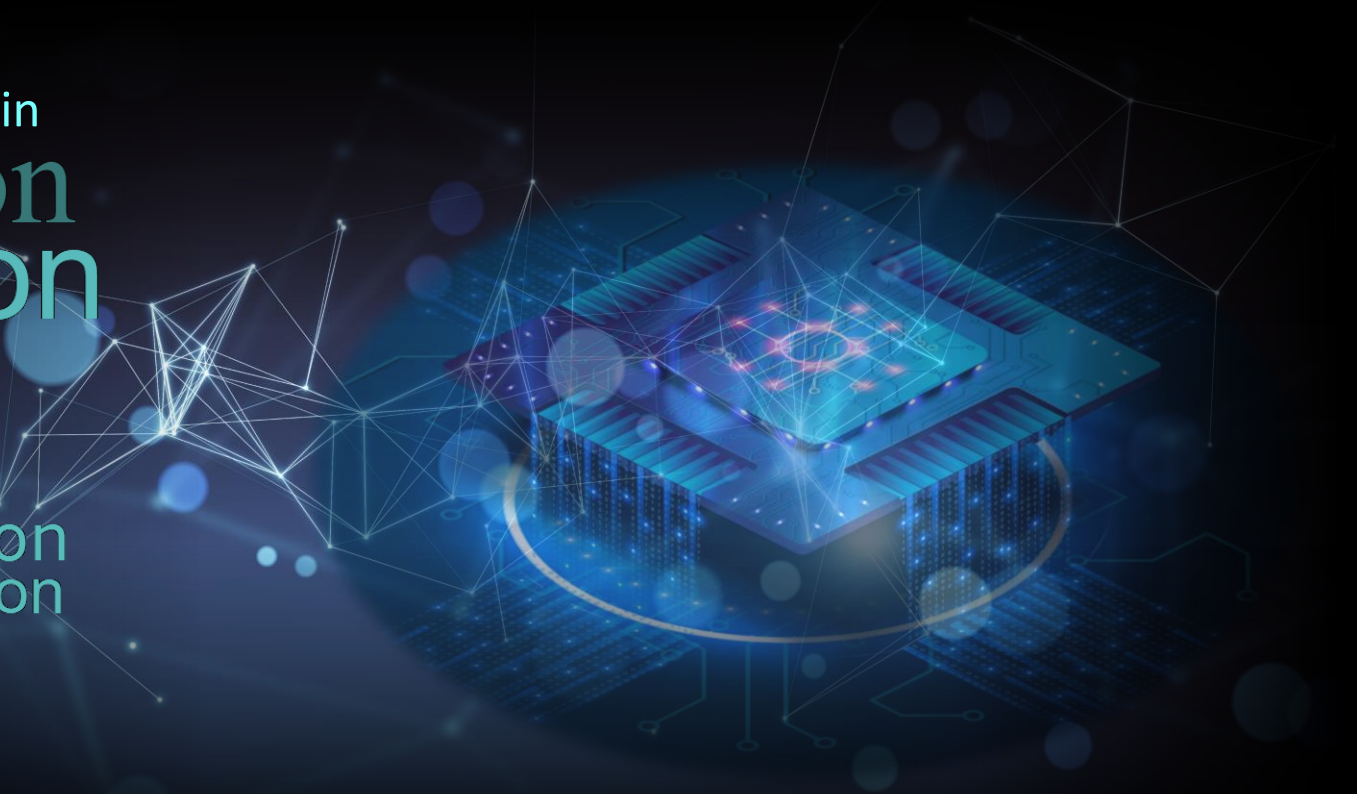
Design space exploration

Defect detection

Logic synthesis Mask synthesis

Place and route

3D integration Security-aware



# EDA and the shift to cloud

press

## cādence

[Cadence Collaborates With TSMC and Microsoft to Reduce Semiconductor Design Timing Signoff Schedules With the Cloud](#)

[Cadence Extends Cloud Leadership With New CloudBurst Platform for Hybrid Cloud Environments](#)

## synopsys®

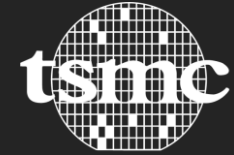
[Synopsys, TSMC and Microsoft Azure Deliver Highly Scalable Timing Signoff Flow in the Cloud](#)

[SiFive Selects Synopsys Fusion Design Platform and Verification Continuum Platform to Enable Rapid SoC Design](#)

## Mentor®

[Mentor's analog/RF/mixed-signal verification tools scale to 10,000 cores on Microsoft Azure](#)

[Mentor and AMD verify massive Radeon Instinct Vega20 IC design on AMD EPYC in ~10 hours with ecosystem partners Microsoft Azure and TSMC](#)



[Microsoft and TSMC announce Joint Innovation Lab to accelerate silicon design on Azure](#)

[TSMC Leads the Industry by Hosting the First "TSMC IC Layout Contest" in the Cloud](#)

blogs



Chris Lattner, SiFive

[Cloud Accelerated Idea To Silicon](#)



Daniel Payne, SemiWiki

[Mentor Adds Circuit Simulators to the Cloud using Azure](#)



Simon Sharwood, The Register

[Microsoft cooking Azure instance types just for chip designers](#)



Chad Morgenstern, NetApp

[Chip Design and the Azure Cloud: An Azure NetApp Files Story](#)



Omar el-Sewefy, Tech Design Forum,

[How cloud computing is now delivering efficiencies for IC design](#)

# Rapid Assured Microelectronics Program (RAMP)

- Manage cloud resources, EDA flows, and design provenance
- Built for collaboration
- Security and data privacy
- Implement Microelectronics Quantifiable Assurance (MQA)
- Design Reference Flows - Builds upon commercial best practices and experience
- Secure design environment - scalable and resilient compute and storage platform
- Built on the secure, scalable, ITAR compliant, Azure Government cloud infrastructure

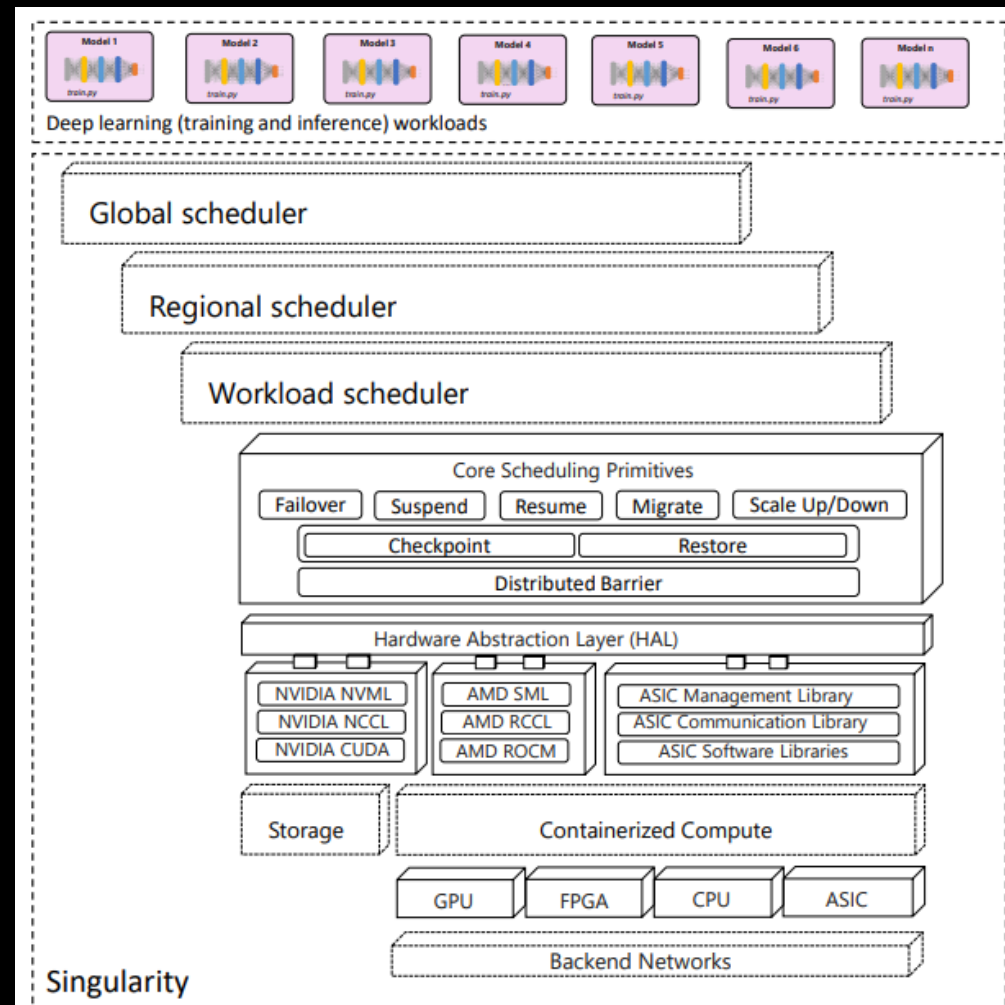


# Real world performance numbers

| EDA Workload                      | On-prem Compute | Azure Compute   | Performance | On-prem Storage | Azure Storage           | Azure Scalability        | Overall ToT               |
|-----------------------------------|-----------------|-----------------|-------------|-----------------|-------------------------|--------------------------|---------------------------|
| SPICE Circuit Simulation - Cells  | Broadwell       | Fv2 16 (HT off) | Equivalent  | NetApp          | ANF                     | Linearity                |                           |
| SPICE Circuit Simulation - Cells  | Haswell         | Esv4            | 20% faster  | NetApp          | HPC Cache + ANF         | Linearity                |                           |
| SPICE Circuit Simulation - Mem/IP | Skylake Gold    | Esv4            | Equivalent  | NetApp          | ANF                     | Linearity                |                           |
| SPICE Library Characterization    | Broadwell       | Esv3            | Equivalent  | NetApp, ZFS     | ANF                     | Linearity                |                           |
| Simulation, RTL Block Level       | Broadwell       | Fv2 16 (HT off) | Equivalent  | NetApp, ZFS     | ANF                     | Near Linearity           |                           |
| Simulation, Gate-level            | Skylake Gold    | H16M            | 20% slower  | NetApp, ZFS     | ANF                     | Near Linearity           |                           |
| Simulation Full Regression        | Skylake Gold    | Hc44            | 5% slower   | NetApp, ZFS     | ANF                     | Near Linearity           |                           |
| DFT Scan/Bist/ATPG                | Skylake Gold    | Mv2             | Equivalent  | Isilon          | Data on-prem            | Near Linearity           |                           |
| Signoff Timing                    | Broadwell       | ESv4            | 30% faster  | NetApp          | ANF                     | Sweet spot at 128 cores  | 3x reduction vs single VM |
| Extraction                        | Broadwell       | ESv4            | Equivalent  | NetApp          | ANF                     | Sweet spot at 128 cores  | 5.7x reduction            |
| Signoff Timing                    | Ice Lake        | FXv1            | 10% faster  | NetApp          | ANF                     | Linearity at 64 cores    |                           |
| Extraction                        | Broadwell       | ESv4            | Equivalent  | NetApp          | ANF                     | Linearity at 64 cores    | 2x reduction              |
| Design Rule Check                 | Haswell         | MV1             | Equivalent  | Master Node     | HPC Cache + Master Node | Sweet spot at 4000 cores | 6x reduction              |
| Design Rule Check                 | Haswell         | ESv4            | 10% faster  | NetApp          | ANF                     | Sweet spot at 4000 cores | 4x reduction              |
| IR Drop                           | Haswell         | Mv1             | Equivalent  | NetApp          | ANF                     | Near Linearity           |                           |
| IR Drop                           | Skylake Gold    | Mv1             | 20% slower  | Isilon          | ANF                     | Near Linearity           |                           |



# Singularity: Elastic scheduling and resource optimization



**Figure 1. Singularity System Architecture**

# Machine Learning Applied to EDA

Resource Optimization

Analog design **Supply Chain**

Result estimation

IR drop prediction

CoPilot for design Power prediction

High level synthesis SAT solver

Verification and test

Lithography hotspot detection

Test set redundancy reduction

Design space exploration

Defect detection

Logic synthesis Mask synthesis

Place and route

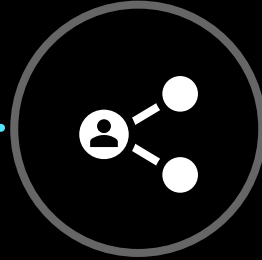
3D integration Security-aware



# Key Insights



Data can enable a security assessment



Companies don't like to share data

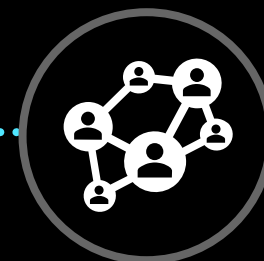


Data do not need to be aggregated

# Things to Avoid

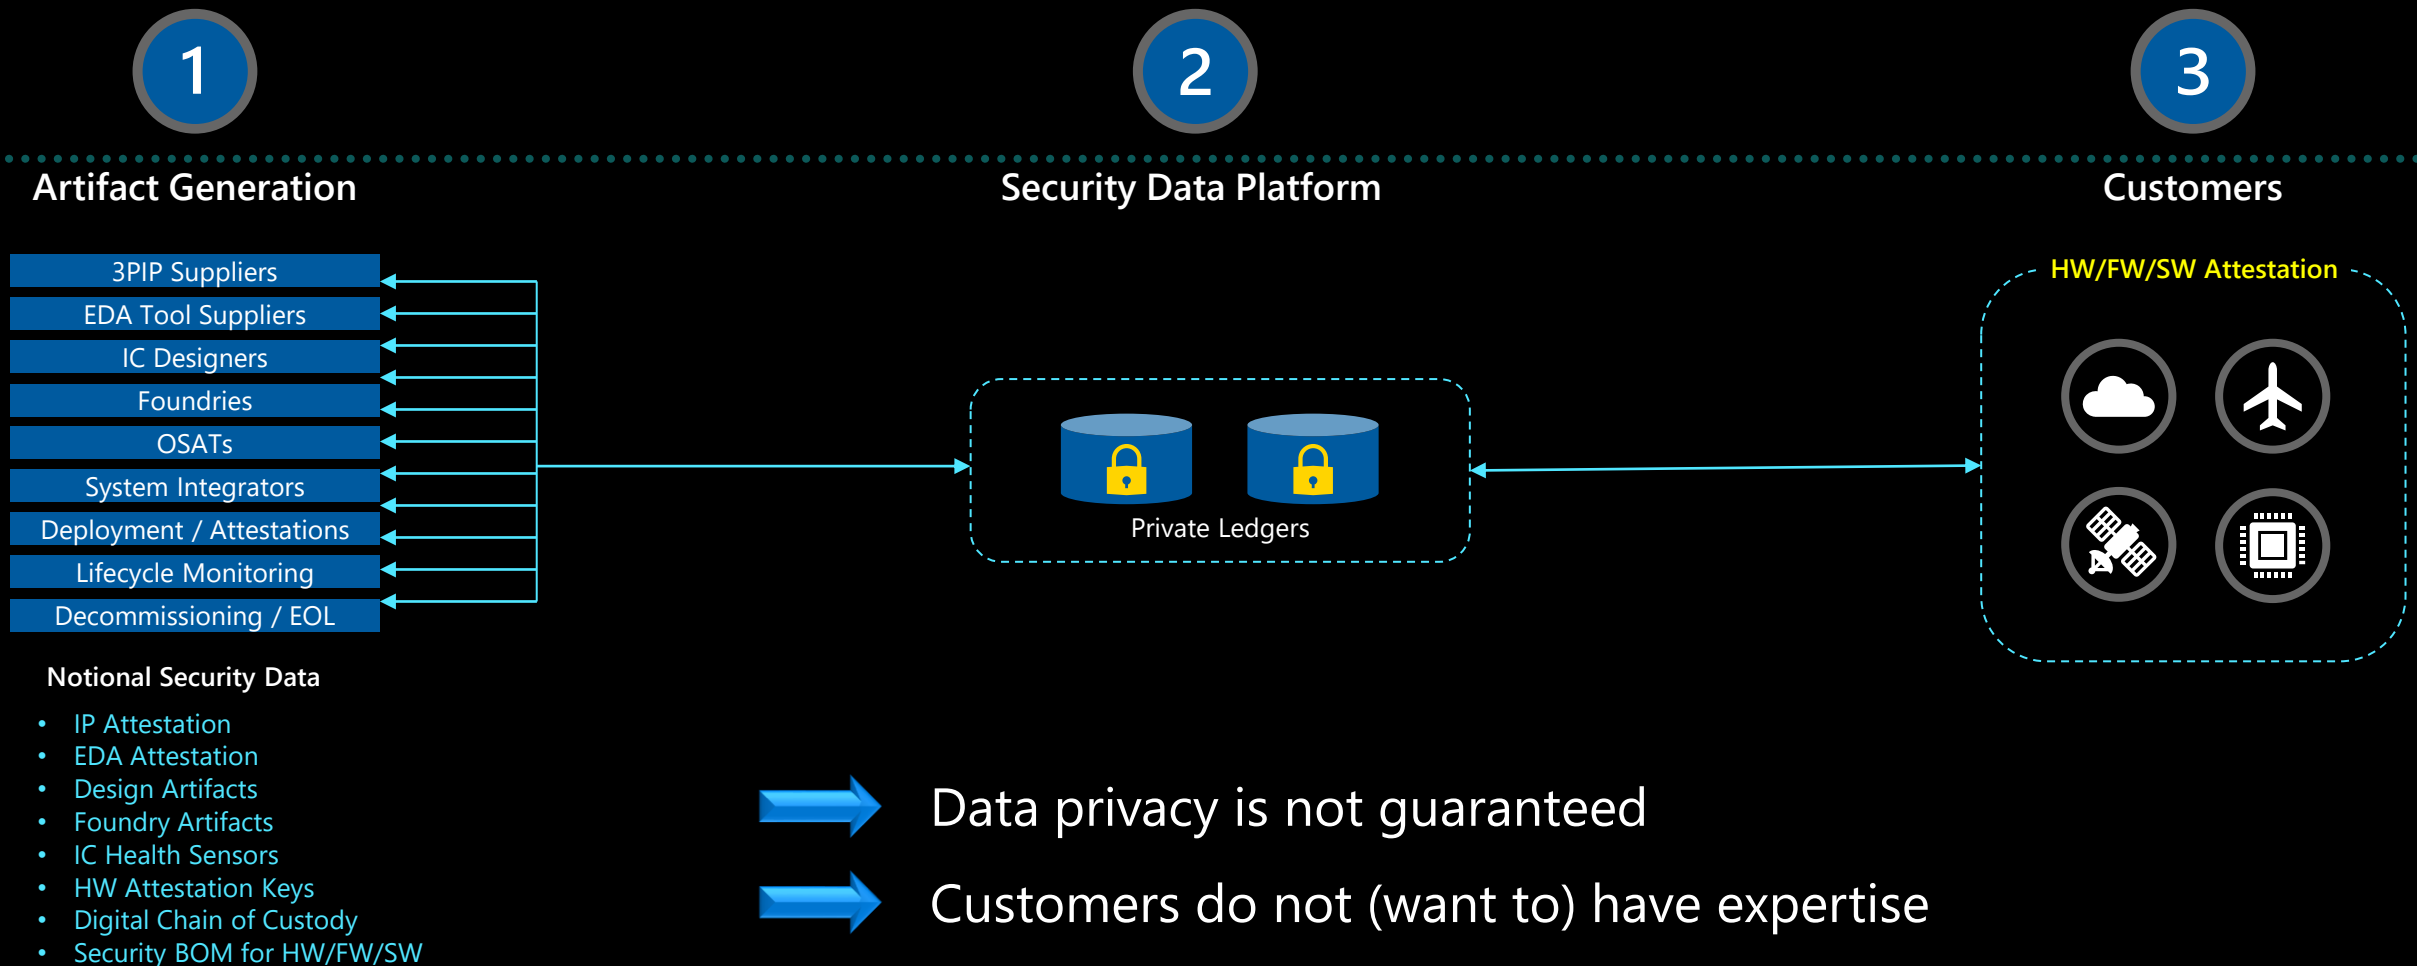


Common Data Repositories

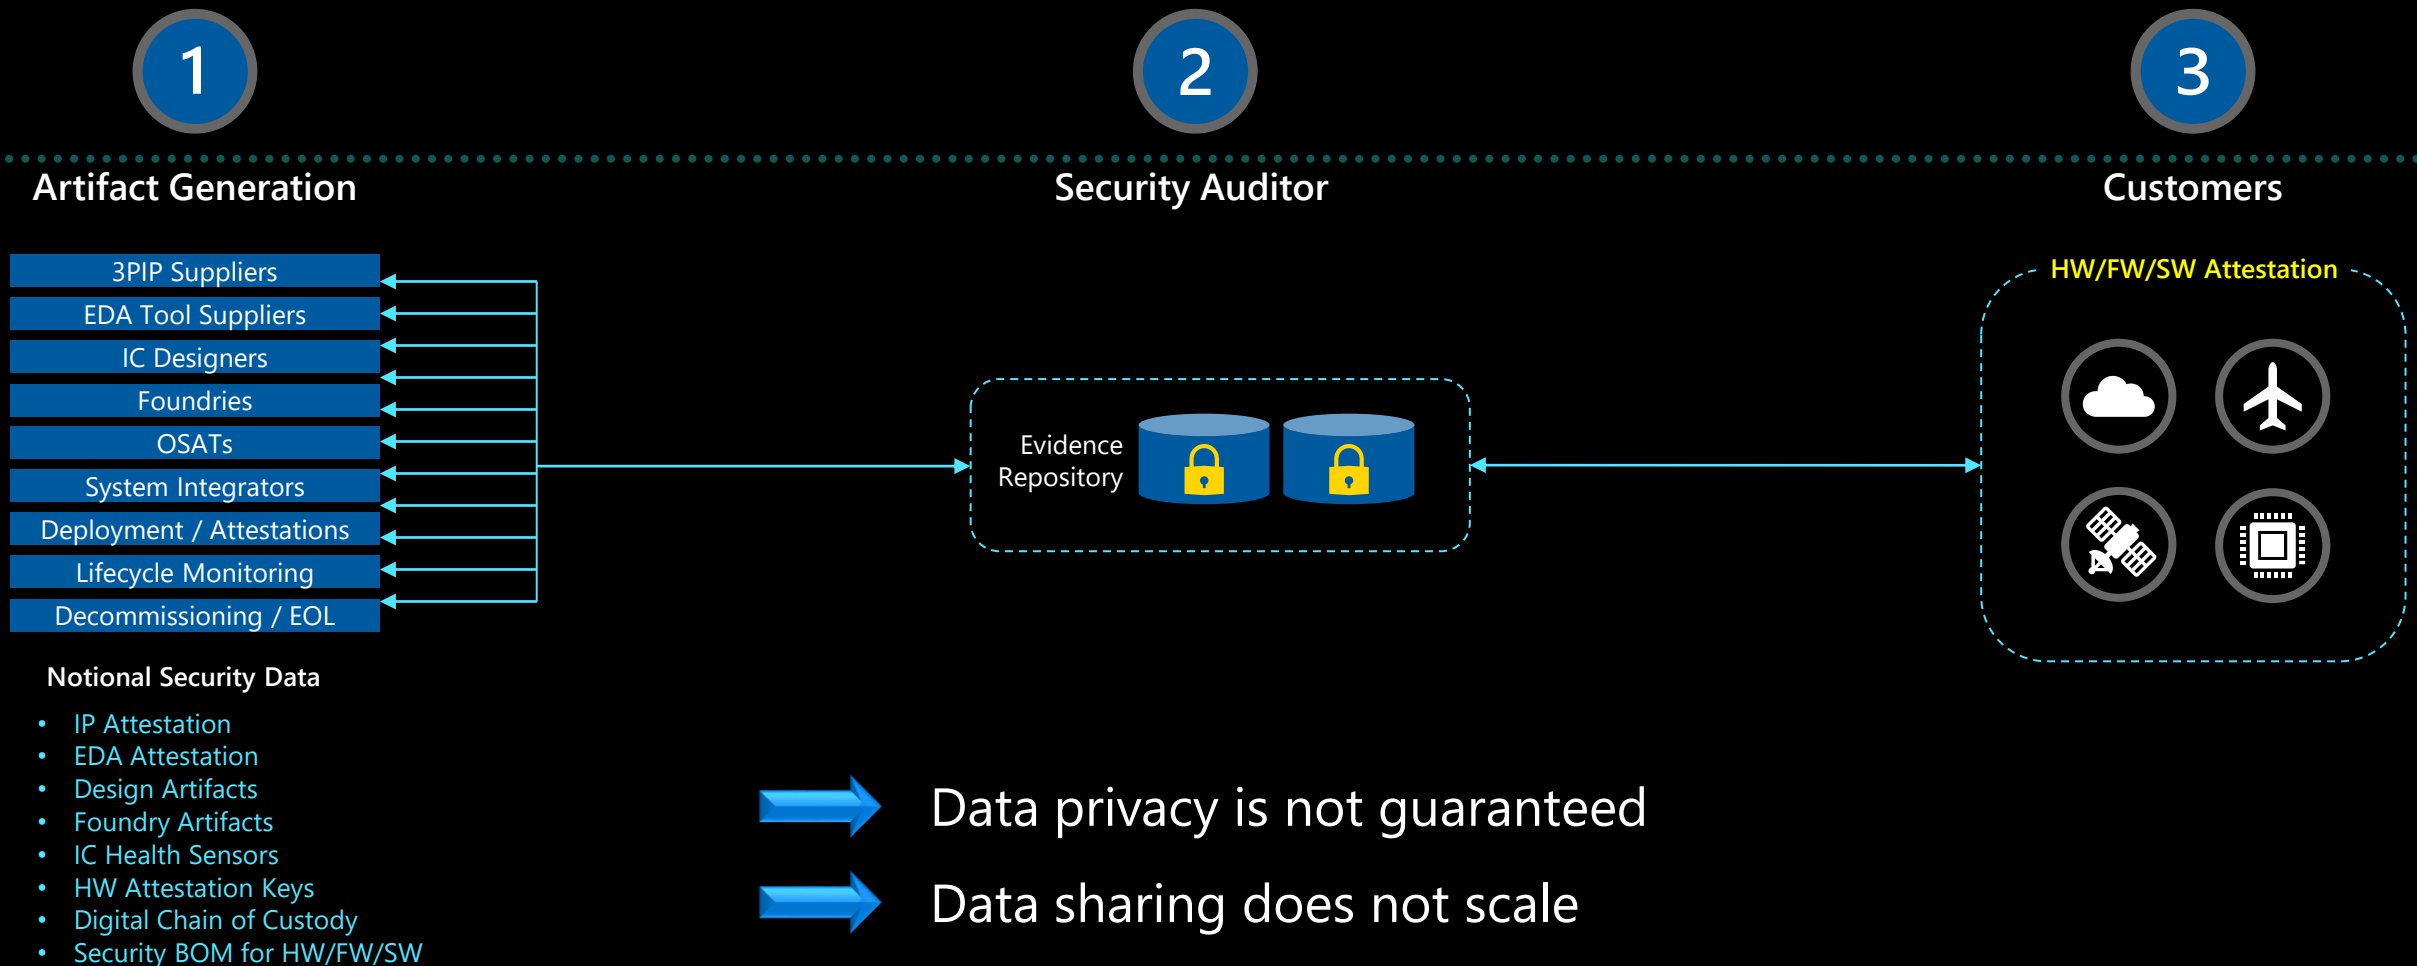


Networks of Data Sharing Agreements

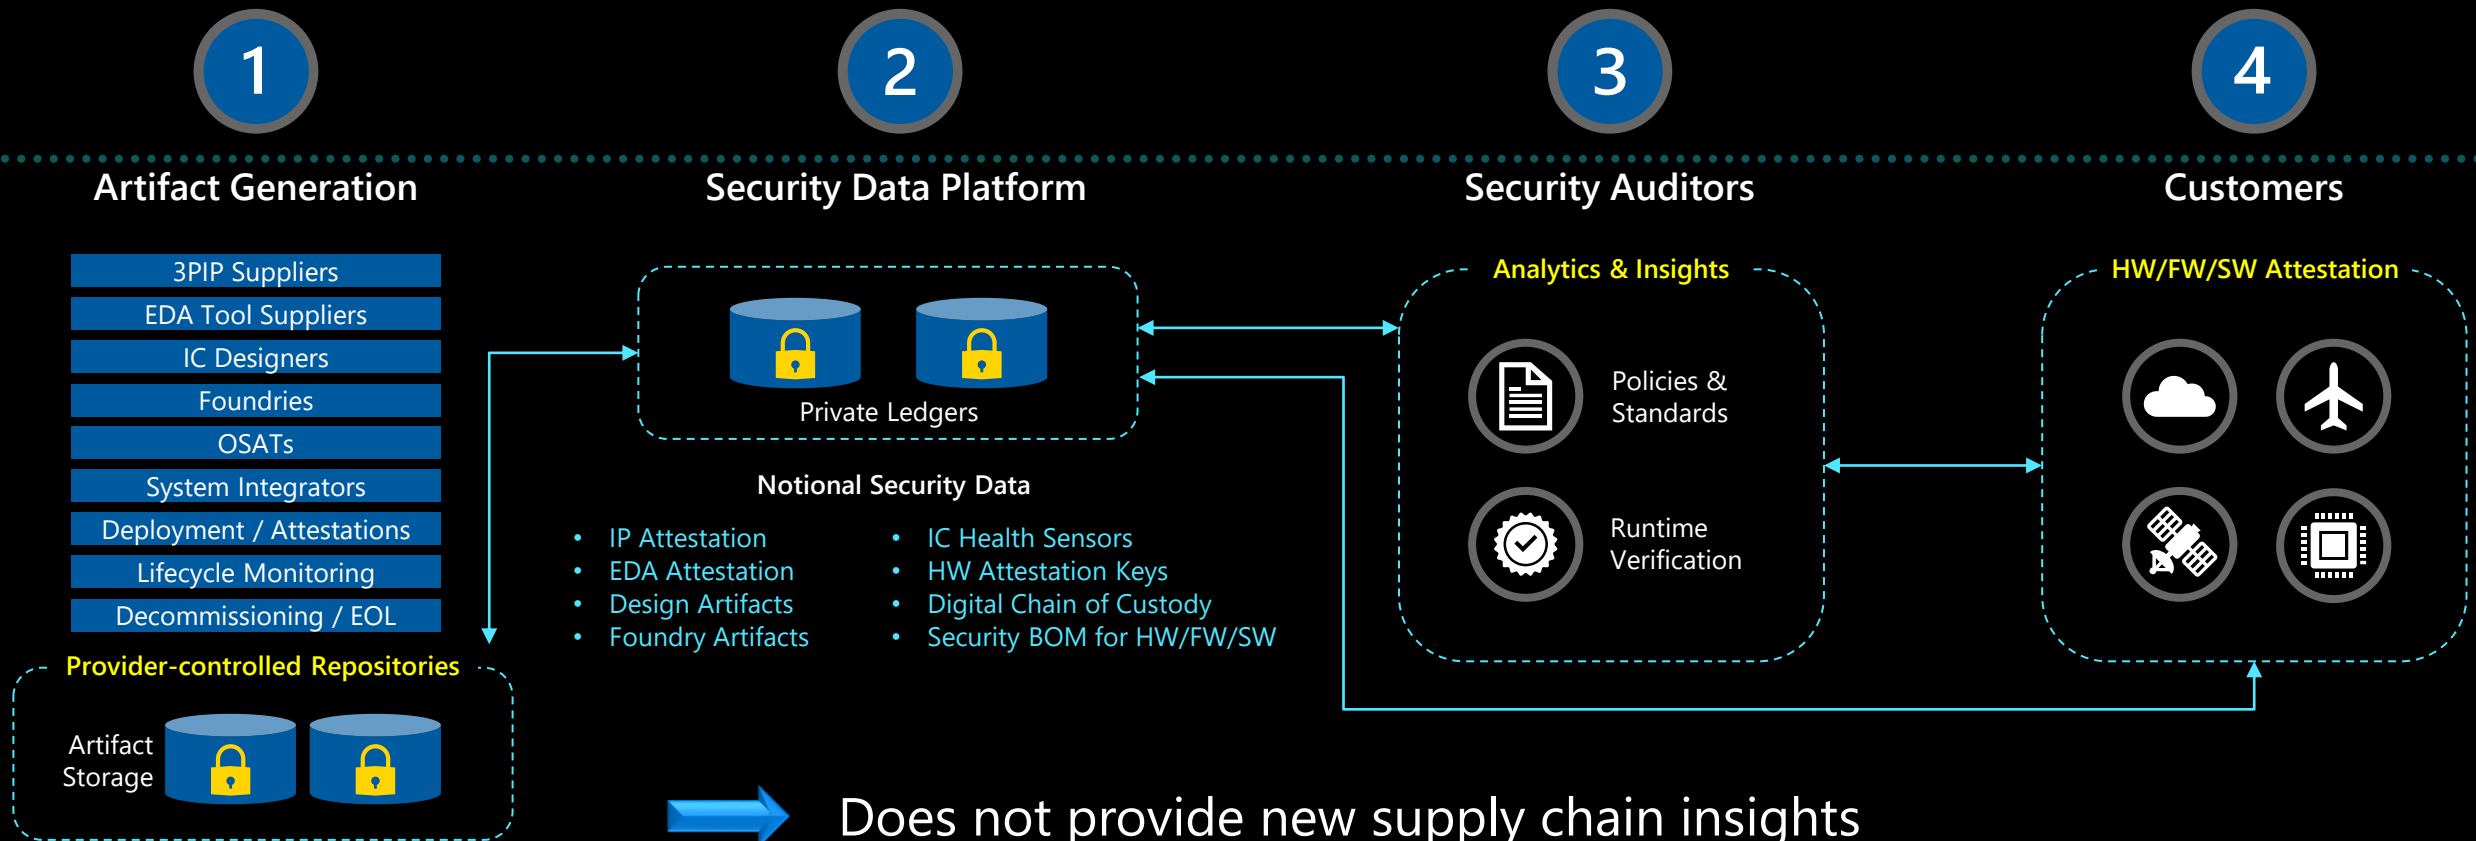
# Data Sharing for Supply Chain Security



# Data Sharing for Supply Chain Security



# Supply Chain Security Architecture that Protects Privacy



- ➡ Does not provide new supply chain insights
- ➡ The only NDA is with the auditors
- ➡ Hybrid AI + rules-based security assessment

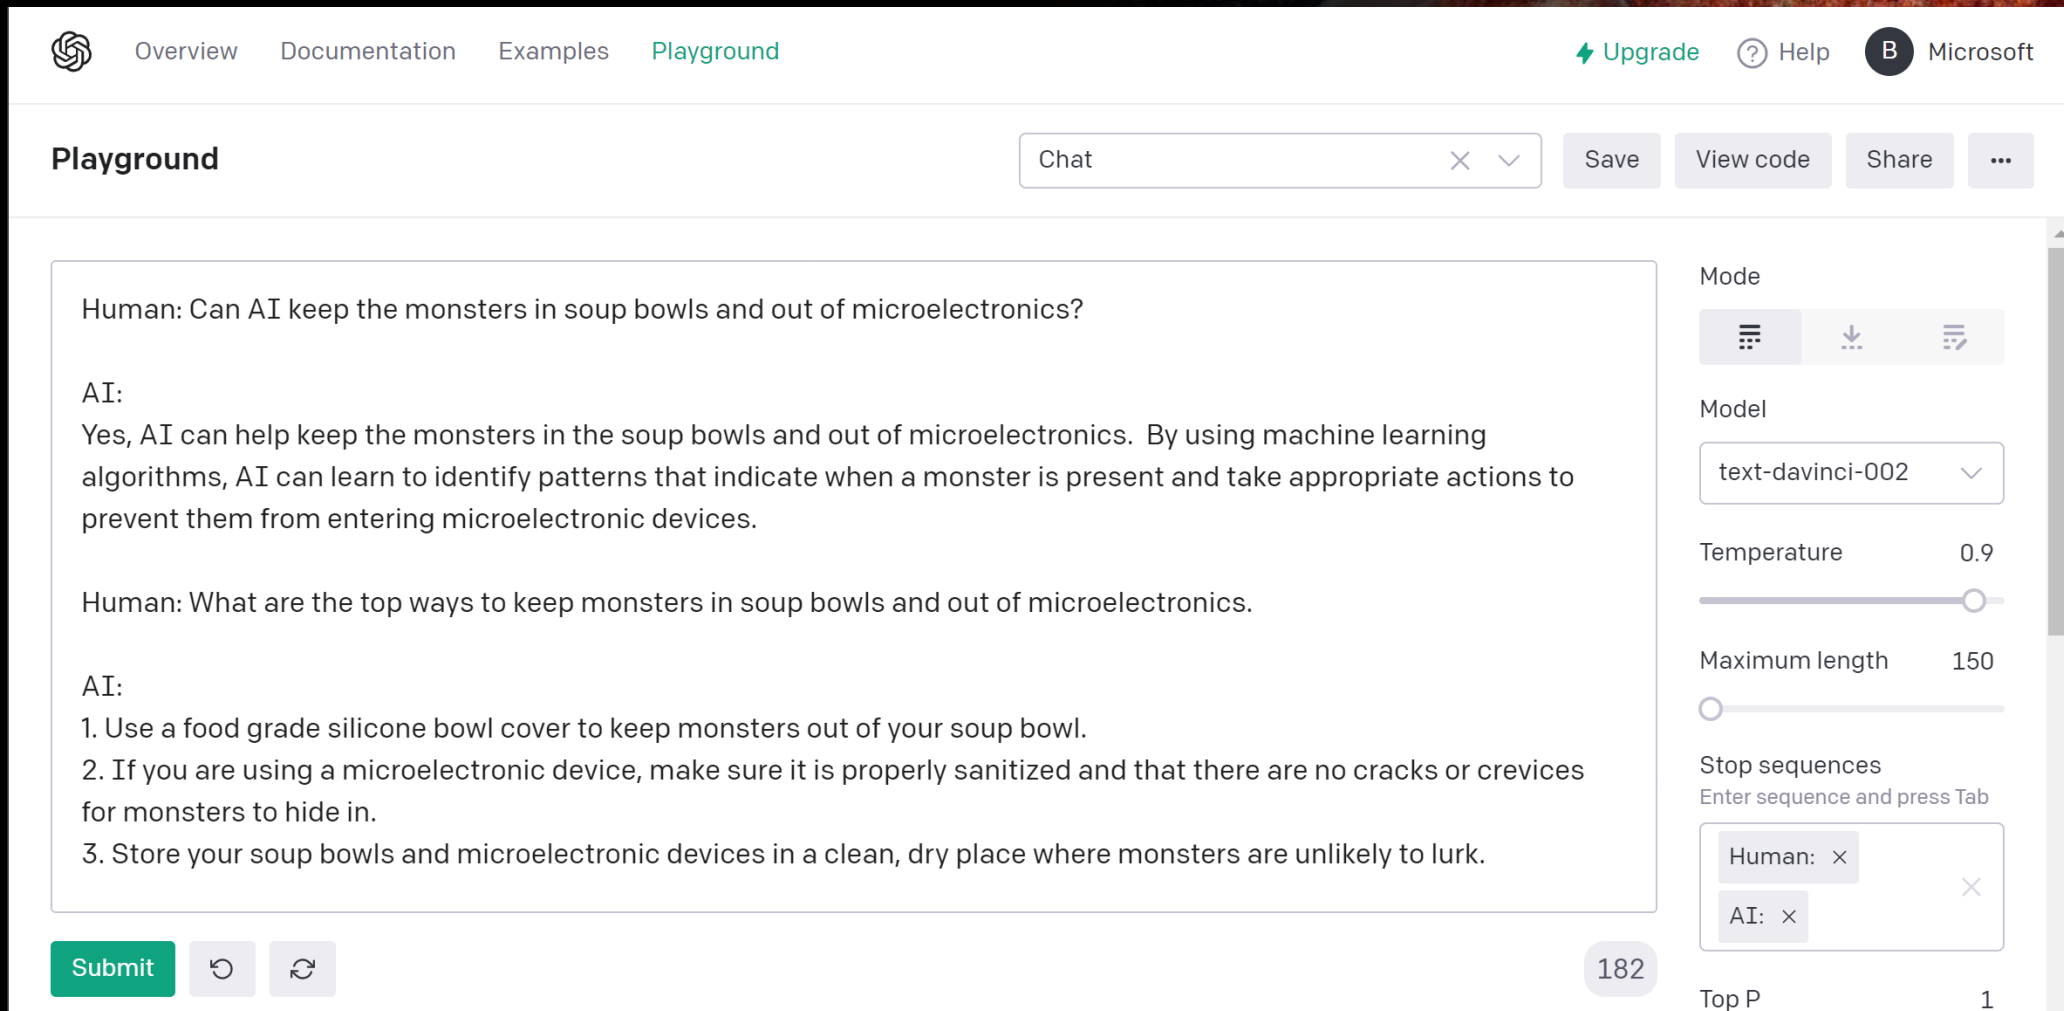
# AI for Microelectronics Security

## Opportunities ahead

- Open datasets, models, and tools for digital, analog, and mixed-signal ICs
- Verification and trust for machine learning driven designs
- AI-driven design, simulation, and exploration across multiple modalities
- Security aware AI pair programming models: Natural language to GDSII
- HW optimization for dynamic AI workloads across the toolchain
- AI-driven data analytics for supply chain security and compliance risk management

# AI for Microelectronics Security

## The future is not scary, and not that far away



The screenshot shows the OpenAI Playground interface. At the top, there are navigation links for Overview, Documentation, Examples, and Playground. On the right, there are links for Upgrade, Help, and a user profile for Microsoft. The main area is titled "Playground" and contains a chat window. The chat window shows a conversation between a human and an AI. The human asks: "Can AI keep the monsters in soup bowls and out of microelectronics?". The AI responds: "Yes, AI can help keep the monsters in the soup bowls and out of microelectronics. By using machine learning algorithms, AI can learn to identify patterns that indicate when a monster is present and take appropriate actions to prevent them from entering microelectronic devices." The human then asks: "What are the top ways to keep monsters in soup bowls and out of microelectronics?". The AI responds with a list of three items: 1. Use a food grade silicone bowl cover to keep monsters out of your soup bowl. 2. If you are using a microelectronic device, make sure it is properly sanitized and that there are no cracks or crevices for monsters to hide in. 3. Store your soup bowls and microelectronic devices in a clean, dry place where monsters are unlikely to lurk. On the right side of the chat window, there are settings for Mode, Model (text-davinci-002), Temperature (0.9), Maximum length (150), and Stop sequences. At the bottom of the chat window, there are buttons for Submit, a refresh button, and a button with the number 182. The background of the slide shows a bowl of food.

Overview Documentation Examples Playground Upgrade Help B Microsoft

Playground Chat Save View code Share ...

Human: Can AI keep the monsters in soup bowls and out of microelectronics?

AI:  
Yes, AI can help keep the monsters in the soup bowls and out of microelectronics. By using machine learning algorithms, AI can learn to identify patterns that indicate when a monster is present and take appropriate actions to prevent them from entering microelectronic devices.

Human: What are the top ways to keep monsters in soup bowls and out of microelectronics.

AI:  
1. Use a food grade silicone bowl cover to keep monsters out of your soup bowl.  
2. If you are using a microelectronic device, make sure it is properly sanitized and that there are no cracks or crevices for monsters to hide in.  
3. Store your soup bowls and microelectronic devices in a clean, dry place where monsters are unlikely to lurk.

Mode  
Model: text-davinci-002  
Temperature: 0.9  
Maximum length: 150  
Stop sequences: Enter sequence and press Tab  
Human: x  
AI: x

Submit 182 Top P 1



# Microsoft

## Booth #1252

### How to Run EDA Tools on the Azure Cloud

- Wednesday, 13 July, 10:15 am – 1:15 pm PDT
- Design on Cloud Pavilion, Level 2 Exhibit Hall

### Bespoke Silicon: Tailor-Made for Maximum Performance

- Wednesday, 13 July, 2:00 pm – 2:45 pm PDT
- DAC Pavilion, Level 2 Exhibit Hall

### How Robust is Your Hardware Security Program?

- Wednesday, 13 July, 3:00 pm – 3:45 pm PDT
- DAC Pavilion, Level 2 Exhibit Hall



