



Big Data and Machine Learning for Chip Design

Electronic Design Process Symposium 2017

John Lee

Vice President & GM, Semiconductor Business Unit



- ANSYS, Inc. product names are trademarks or registered trademarks of ANSYS, Inc. or its subsidiaries in the United States or other countries.

Have **margins** outlived their
usefulness?



Silos == Margins


<https://en.wikipedia.org/wiki/Silo>

Chip

Package

System

Domains



<http://regmedia.co.uk>



X

Physics

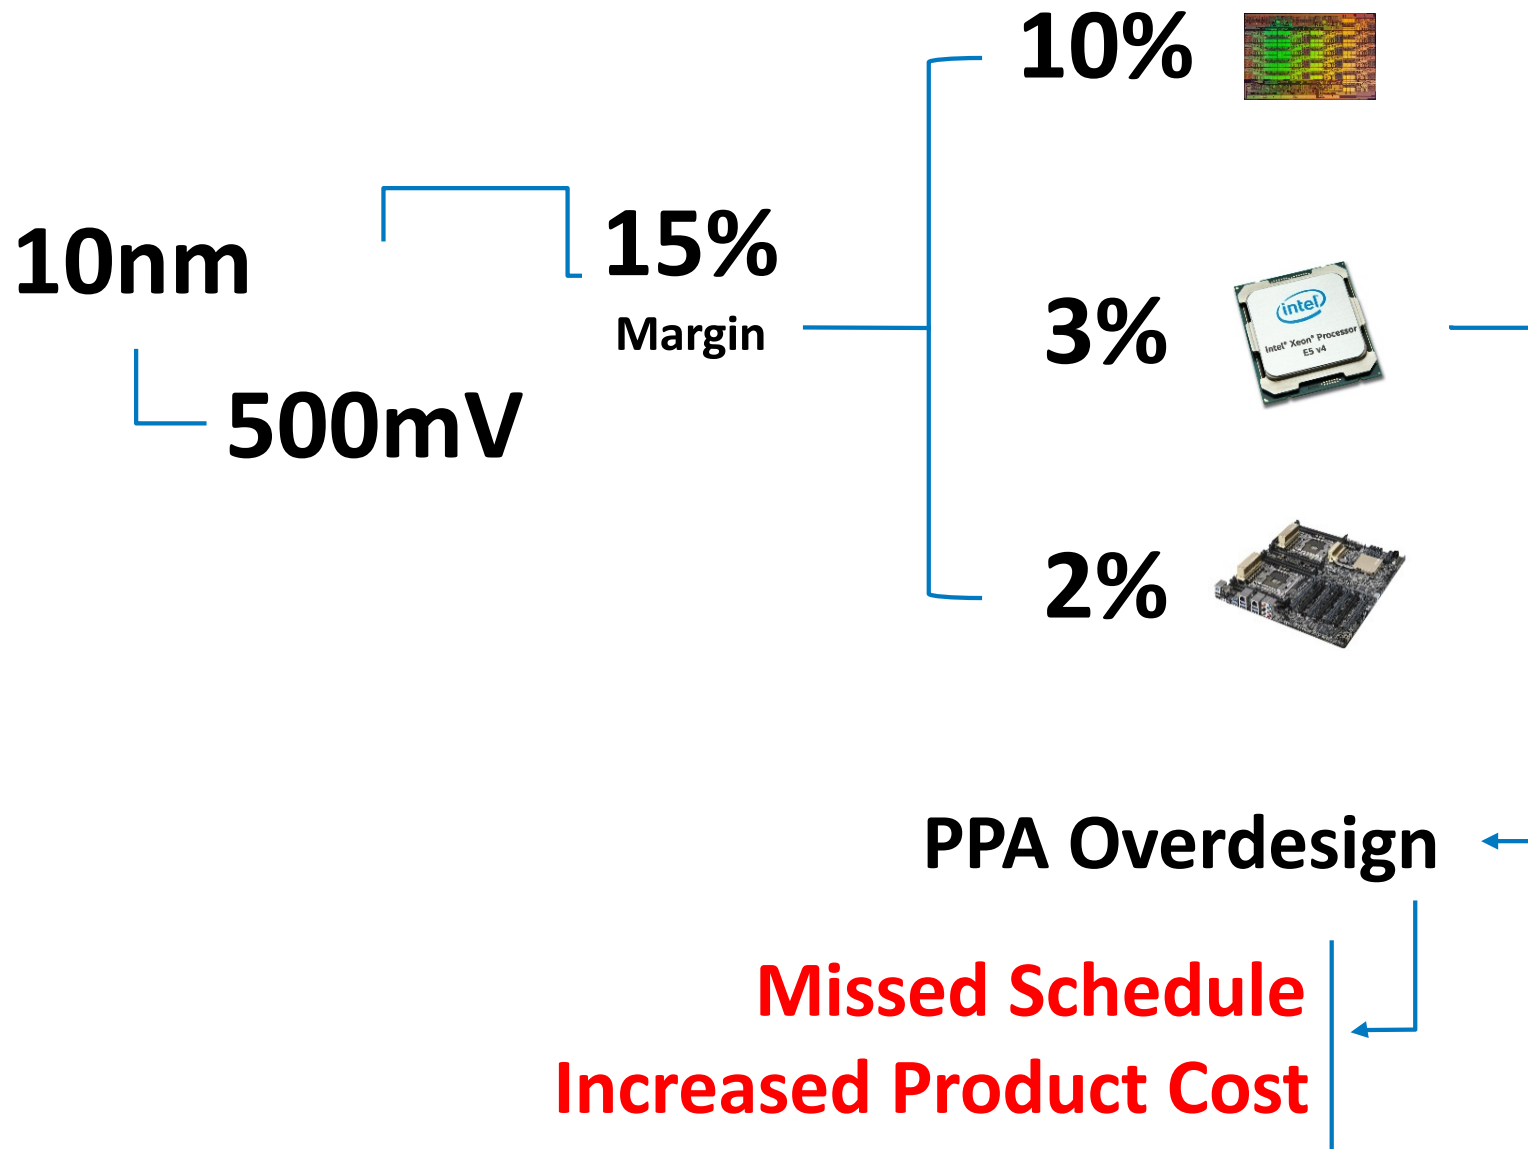
Timing

Power

Thermal

EMI


Mechanical



Margins have outlived their usefulness.

<http://semiengineering.com/have-margins-outlived-their-usefulness/>

Big data has changed the world.

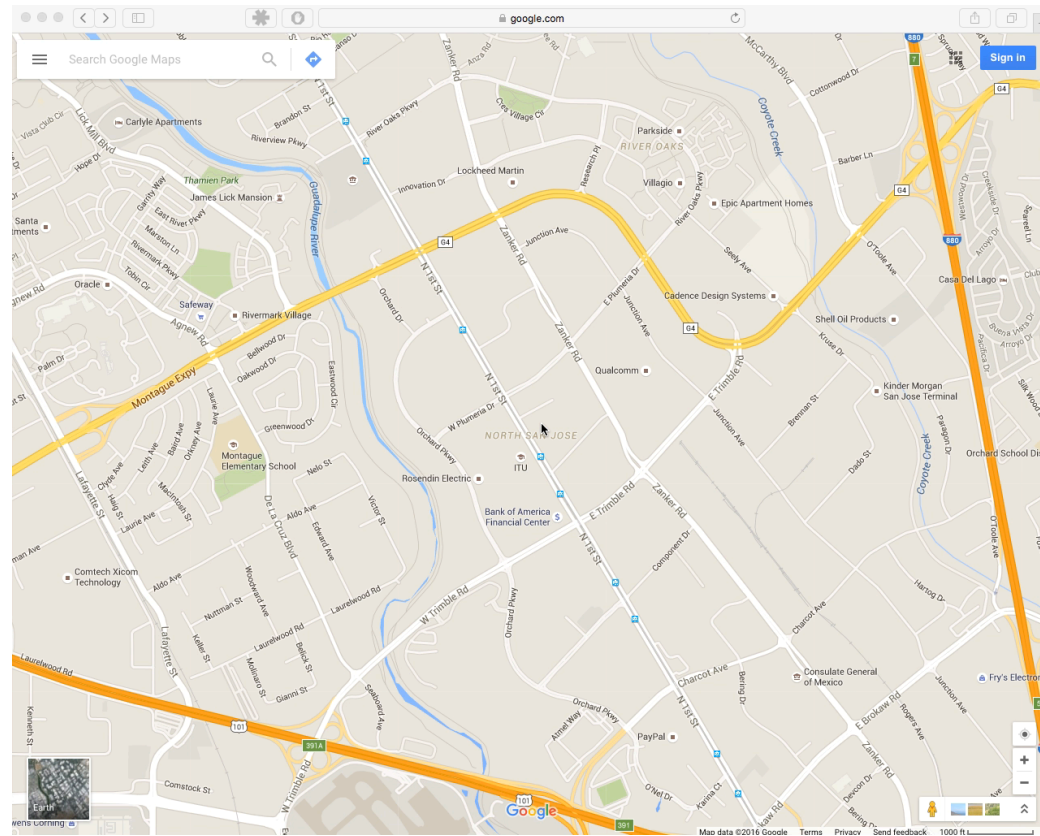


Pre-Google™

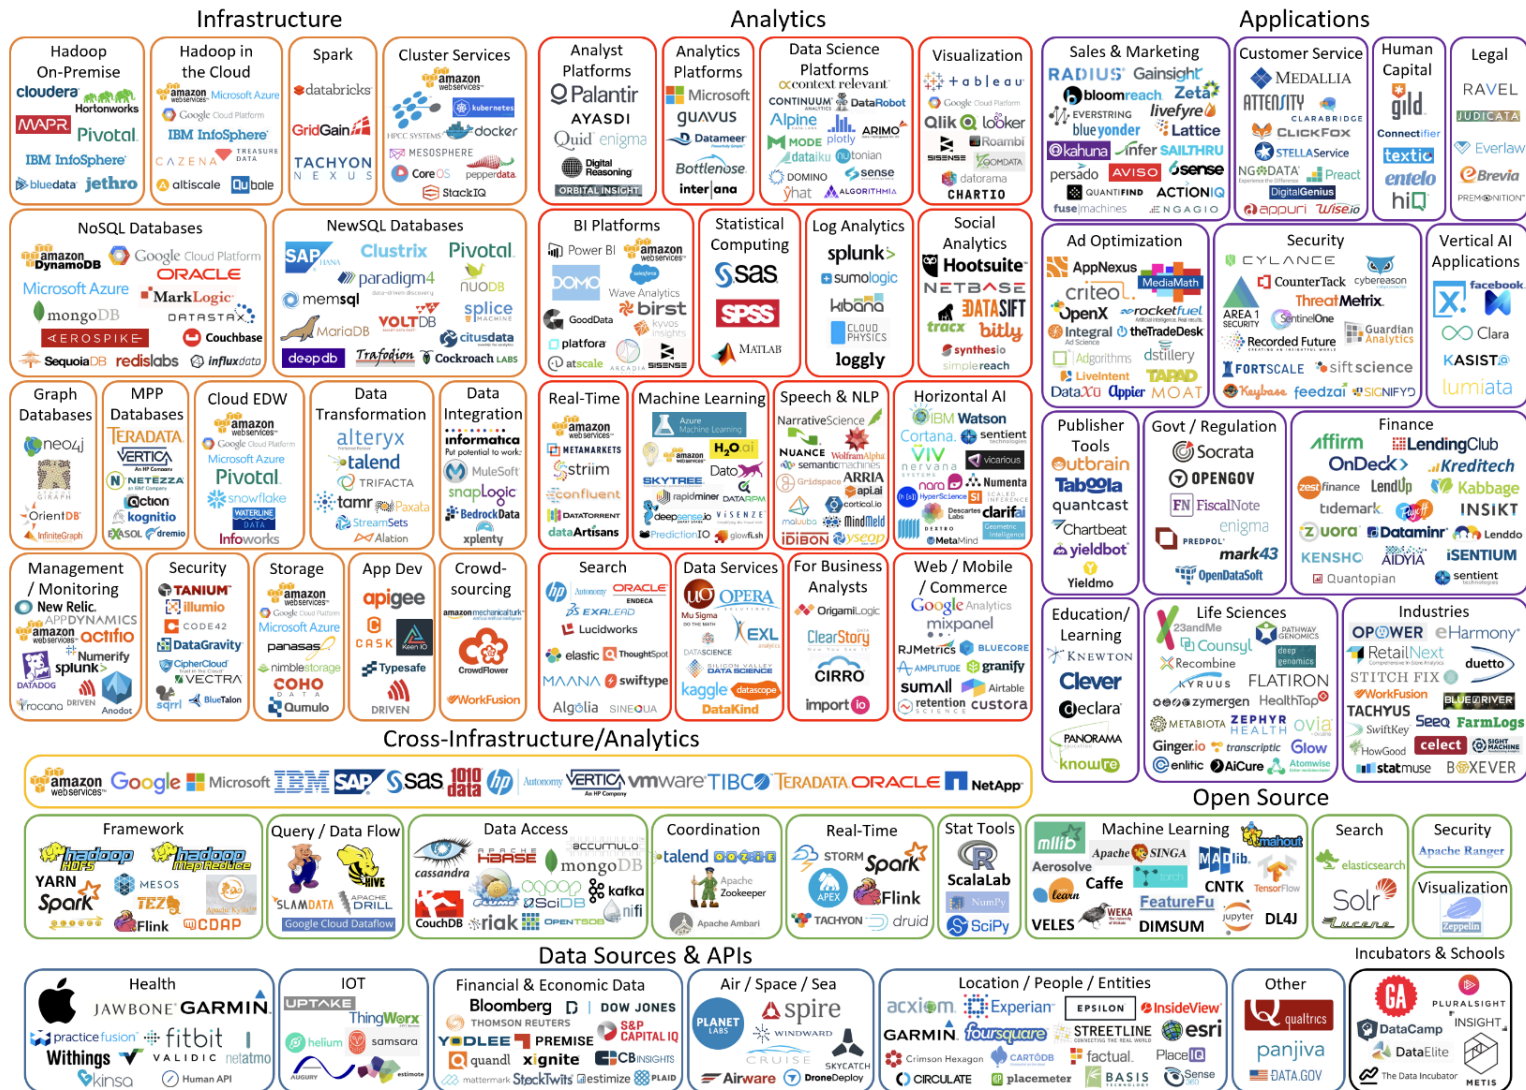
The data was there, but you had to work hard to make it useful.

Google™ Maps

Big data
MapReduce
Hadoop



Big Data Landscape 2016 (Version 3.0)



Last Updated 3/23/2016

© Matt Turck (@mattturck), Jim Hao (@jimrho), & FirstMark Capital (@firstmarkcap)

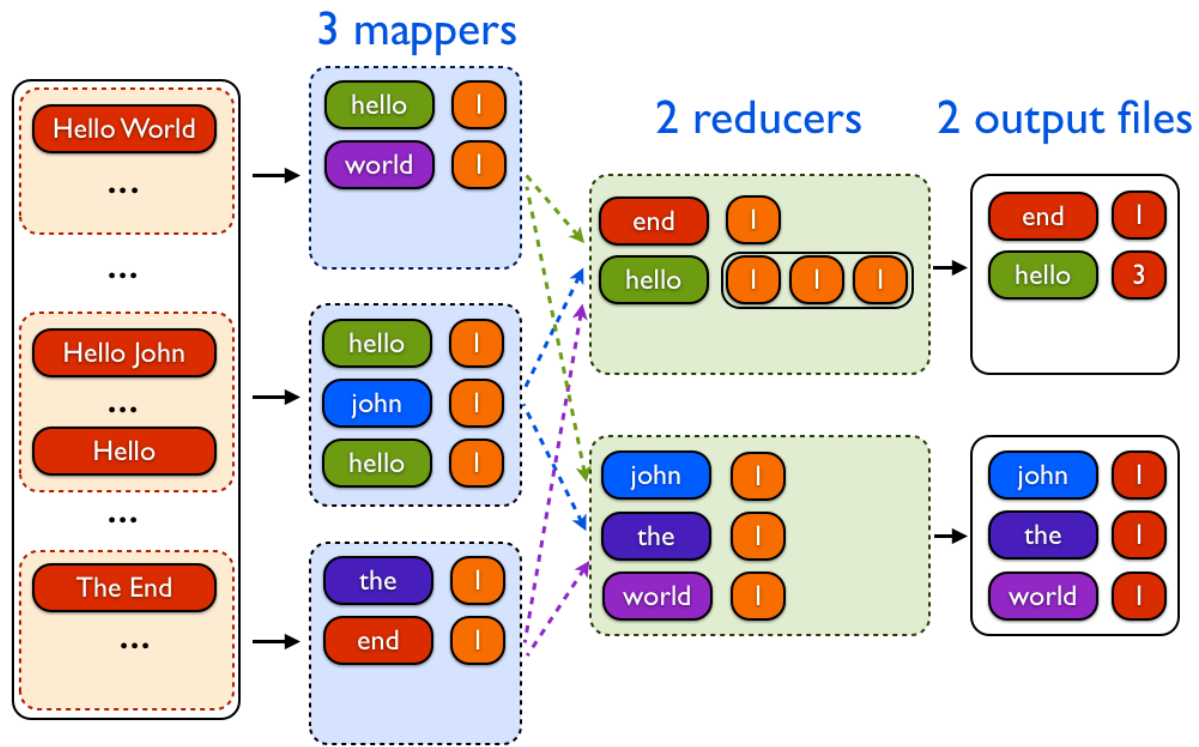
FIRSTMARK

Big data startups received \$6.64B VC investment [2015]

<http://mattturck.com/2016/02/01/big-data-landscape/>




Hadoop MapReduce




Big data works great for many things...

<http://logicheartblogs.blogspot.com/>



Big data for chip design?



21B transistors
(Source: NVDA, 2017)

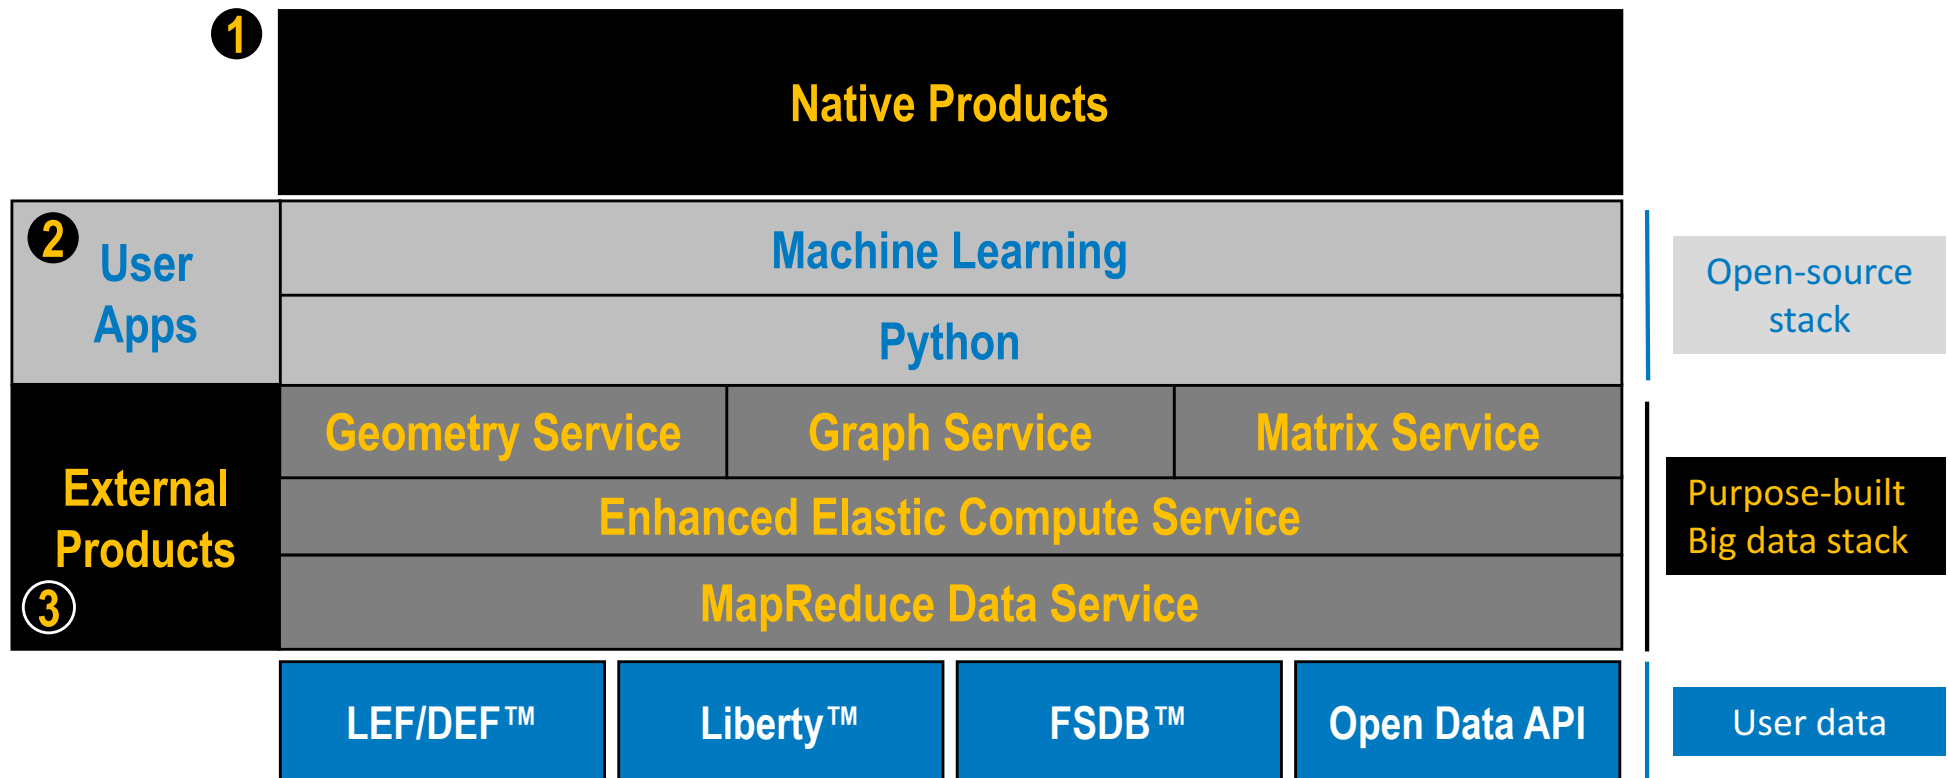
1 millions vectors
1 billion logical instances
10 billion transistors
100 billion geometries
1 trillion model elements

The complexity of a chip **exceeds** that of the internet.

Challenges


- **Silo'ed data is produced by separate tools / vendors**
- **Data sizes exceeds that of existing commodity hardware**
- **Over-design due to margins has cost for designs today**
- **Advanced computational sciences – such as MapReduce, ML/DL – are applied ad hoc in chip design**

ANSYS SeaScape™: A purpose-built big data platform for chip design



- 1- Native products
- 2- Open, actionable analytics
- 3- Integrated workflows with non-native products

Elasticity




Maximize the use of commodity hardware; elastically scaling as needed, with resiliency and efficiency


Power Integrity Analysis

- **Goal**
 - Ensure VDD and VSS does not impact parametric yield, across all operating conditions ("black out", "brown out", etc...).
- **Challenge**
 - Power grid attaches to every device in the chip
 - Dynamic conditions, including package / board resonance, can not be exposed by static techniques
 - At low voltages (e.g. mobile processors), the impact on timing is extremely important
- **Known to be one of the most computationally intense analysis steps for chip design**

Power Integrity: Worst Case? Coverage?




A large set of RTL and/or gate vectors



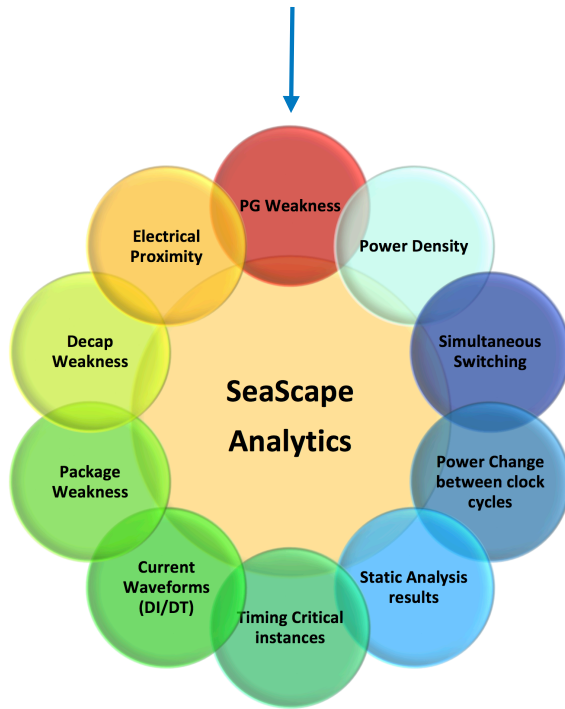
What vectors are needed for power integrity? – worst case, coverage, interaction with timing, etc...?

Scenario Scoring and Coverage




A large set of RTL and/or gate vectors

SeaScape™ based products can process an order of magnitude more data than before



Scoring toolkit

Using SeaScape MapReduce, processing is rapid (minutes)
Standard power, timing, vector data easily processed.




Scoring and coverage

Each vector is scored.

Good (or partial) vectors selected to provide coverage.

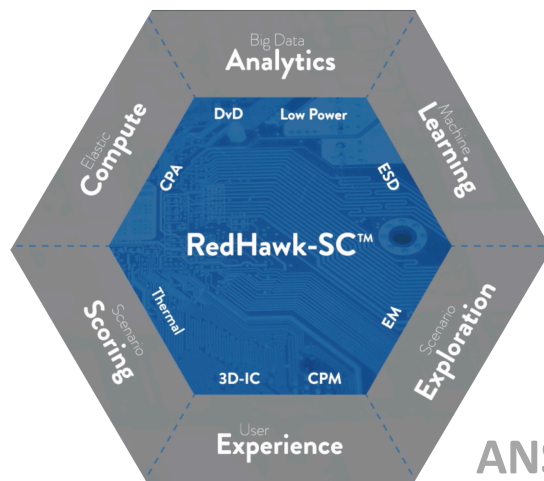
Auto-generated (directed) vectors can be added.

ANSYS Scenario Scoring and Coverage



Multi-Scenario Simulations


Once scored, a set of RTL, gate and vectorless scenarios are run as single set.



ANSYS RedHawk-SC™

SeaScape™ Elastic Compute

Each scenario can be processed in parallel, or in series, depending on # of CPU cores available at that time. The # of cores can be dynamically adjusted, based on need and availability.

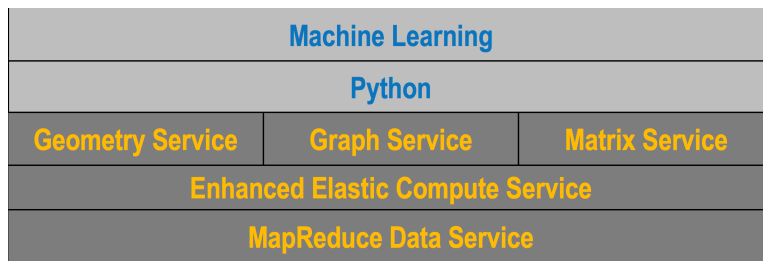
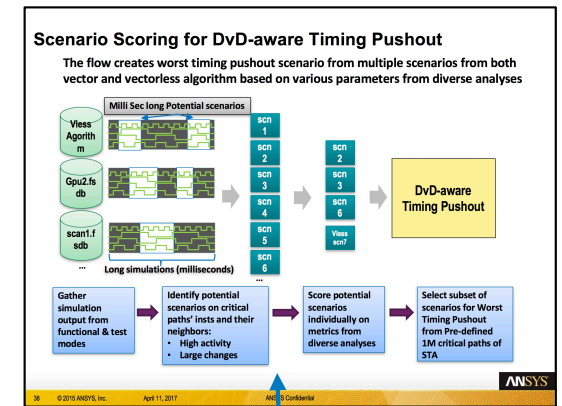
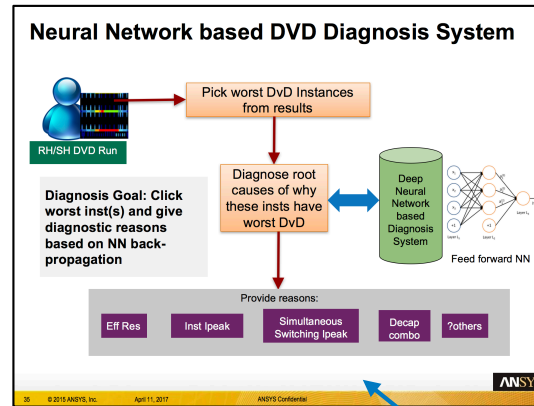


User data intentionally obfuscated for presentation purposes.

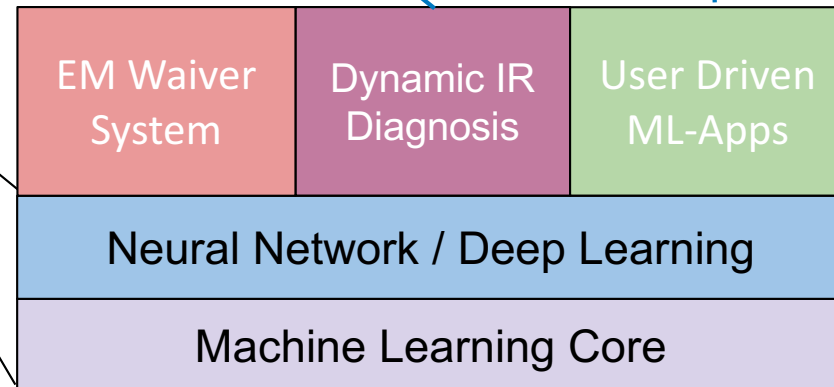
SeaScape™ User Experience

Each result can be viewed, queried and reported individually, or as a holistic single result. No large memory machines needed, due to

SeaScape™ Machine Learning



SeaScape Platform



Machine Learning Layer / Services

Google TensorFlow
Numpy K-means Clustering, ...

ANSYS has an integrated ML stack for scientific computing, and it natively understands EDA data



Conclusion

- **A purpose-built big data platform has been developed as a platform for chip design**
- **The platform leverages open source software for advanced computational sciences.**
- **Opportunities exist to reduce design margins, and improve PPA using such techniques**