ITC 2023 PO.19: Scan Diagnosis On The Cloud

John Abraham, Simon Springall, Sumanth Shindgi, Jayant D'Souza, Randy Klingenberg

SIEMENS

Abstract Summary

Performing volume scan diagnosis on advanced node semiconductor designs warrants memory intensive compute resources and ability to scale such resources on demand. A cloud platform like Amazon Web Services (AWS) offers flexibility to scale such compute resources on demand with pay-as-you go pricing compared to on-prem resource limitations. We demonstrate 50% faster turnaround times for volume scan diagnosis workloads on AWS cloud compared to on-prem.

Scan Diagnosis Overview

Scan diagnosis uses failure data from manufacturing test, scan test patterns and design information to identify the location and classification of the defect causing the failure as shown in Figure 1.

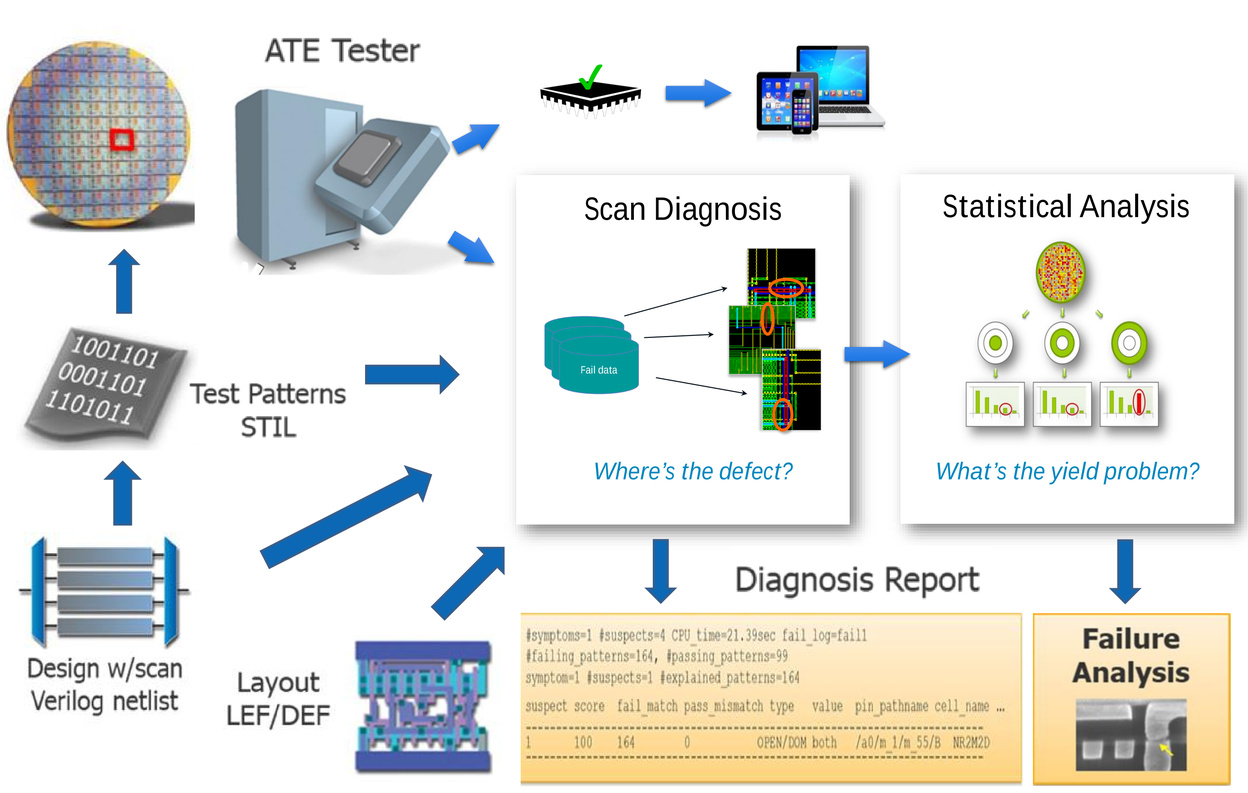

Figure 1: Scan Diagnosis Overview

A leading-edge customer Design A use case with 25 million simulation gates, 9600 scan test patterns and 4010 scan failure logs for volume diagnosis is shown in Figure 2. A single failure log requires 1 analyzer machine to diagnose.

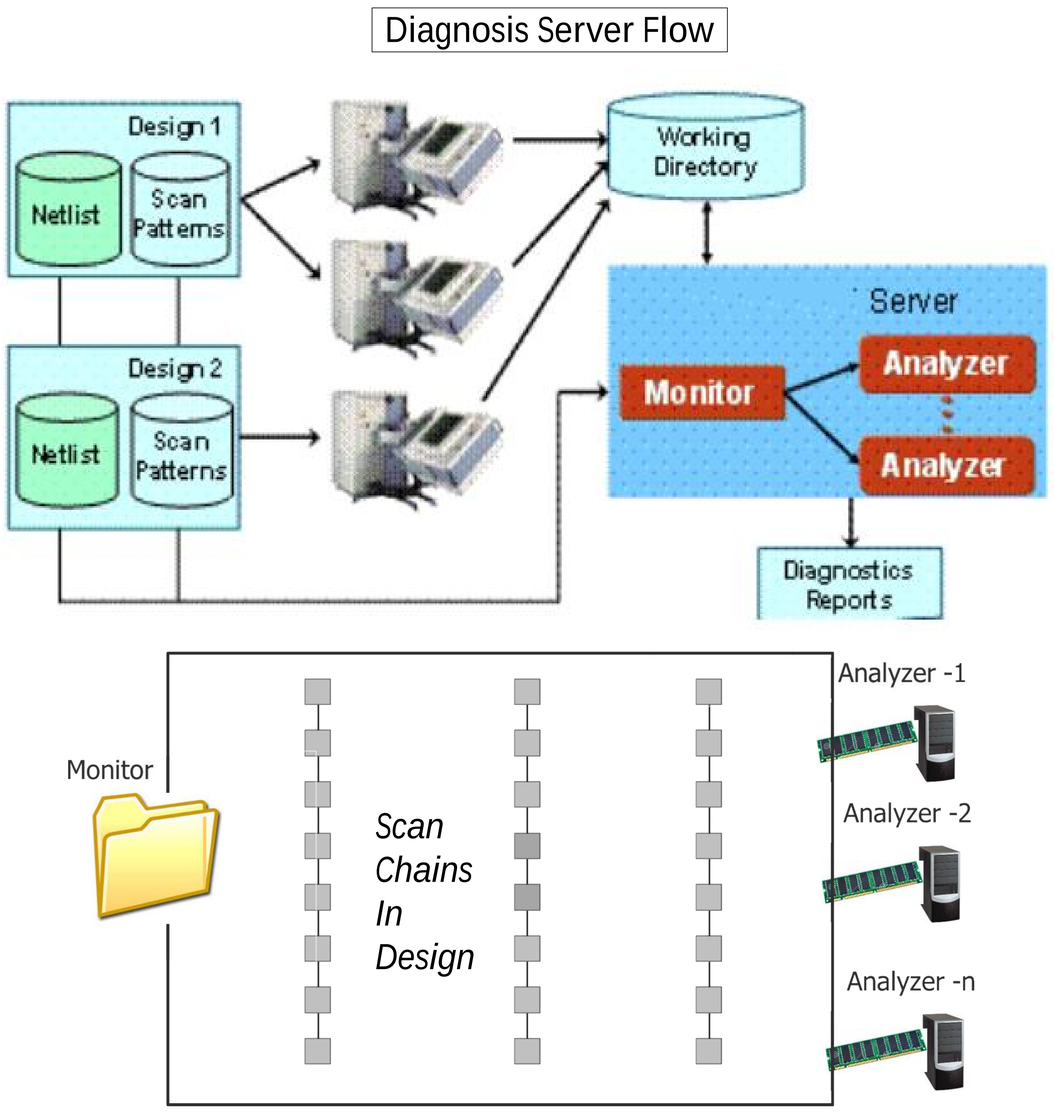

Per analyzer memory requirement : 45GB

Figure 2: Volume Scan Diagnosis Memory Intensive Workload

Section I : Lift and Shift

- ❑ Design B containing 160,000 simulation gates and 978 scan test patterns used for proof of concept
- ❑ 32 elastic compute cloud (EC2) R6a memory-optimized instances using AMD 7000 series EPYC processors as analyzer machines within head node
- ❑ Successful diagnosis of 200 failure logs to complete the lift-and-shift exercise

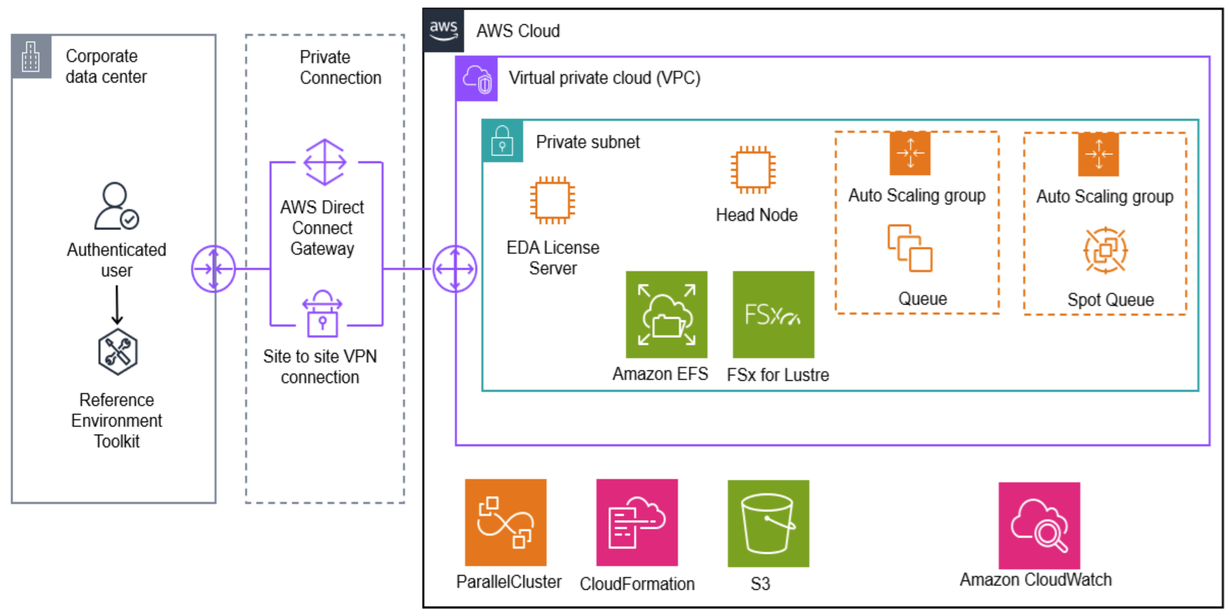

Figure 3: AWS Cloud Architecture for EDA Workloads

AWS cloud architecture for running leading volume scan diagnosis EDA software is shown in Figure 3.

- ❑ Provision cluster using AWS ParallelCluster
- ❑ Configure cluster using AWS CloudFormation through templates optimized for scan diagnosis workloads
  - ➢ Memory optimized EC2 instances for analyzer machines
  - ➢ Compute optimized EC2 instances for head node
  - ➢ FSx for Lustre file systems for EDA software
  - ➢ Elastic file systems (EFS) for design collateral
  - ➢ EDA license server residing within same virtual private cloud (VPC)

Section II : Spot Instance Cost Savings

- ❑ Spot vs on-demand AWS EC2 memory-optimized machines for volume diagnosis studied
- ❑ Design B containing 160,000 simulation gates and 978 scan test patterns exercised over a 24-hour continuous scan diagnosis run
- ❑ Spot EC2 instances offer cost savings that scale with larger machines

Machine Config.	Memory	CPU Cores	On-Demand Cost	Spot Cost	Cost Savings
Small	32 GB	4	$18	$12.5	31%
Medium	256 GB	32	$57	$25	56%
Large	768 GB	96	$163	$37.5	77%
Extra Large	1536 GB	192	$278	$47	83%

Table 1: Cost Savings Using Spot Instances (24 Hour Run)

Section III : Run Time Savings

- ❑ Efficient memory-based scheduling technique using Simple Linux Utility for Resource Management (SLURM) on AWS exercised
- ❑ Advanced node semiconductor Design A containing 25 million simulation gates, 9600 scan test patterns and 4010 scan fail logs used
- ❑ Resulted in 50% faster turnaround time using AWS compared to on-prem

Experiment Run	Number of Machines	Time To Diagnose	Time Savings
On-Prem	50	12D:21H:28M:57S	6D:7H:14M:45S
AWS Cloud	100	6D:14H:14M:12S	

Table 2: Time Savings AWS Cloud vs On-Prem

Conclusion

- ❑ Successful lift and shift from on-prem to AWS cloud
- ❑ Achieved cost savings of up to 83% using spot over on-demand instances on AWS
- ❑ Demonstration of 50% faster turnaround times for volume scan diagnosis workloads on AWS cloud compared to on-prem