

IBM Cloud HPC EDA Solution for Cadence – Case Study

Nov 12, 2023

Ravi Konuru, Principal Solution Architect IBM Cloud, IBM

Gregory Gelles, Storage Solutions Architect Cadence Design Systems

Electronic Design Automation – Challenges

At this level, the requirements overlap with multiple industries, not just EDA

COMPUTE

Sustained use: Ability to efficiently use existing computing resources

Elastic use: Ability to creating computing resources on demand and execute jobs when existing compute pool is insufficient

DATA HYDRATION & SYNCHRONIZATION

Hydration: Data needs to be available with low latency where the job is executing for high performance Synchronization: Data needs to be kept consistent across

Synchronization: Data needs to be kept consistent across geos for global collaboration

SECURITY & FOUNDRY CERTIFICATION Core Security: Given the high IP involved in chip design, data leakages can be catastrophic, so all security and especially of data is critically important

Foundry approval: Chip manufacturing foundries will only allow their process development kits to be used only in HPC environments that they approve



IBM Cloud HPC Offerings

IBM Cloud offerings with automated provisioning of HPC clusters using

- Workload schedulers: IBM Spectrum Symphony (financial workloads) & IBM Spectrum LSF (non-financial workloads), SLURM (opensource scheduler)
- High performance storage: IBM Spectrum Scale

Symphony + Spectrum Scale



Financial Services Sector (FSS)

- Risk Analytics
- Regulatory compliance
- Monte Carlo
- FSS Asset Pricing

Based on IBM Research benchmarks & current capabilities targeting four workload types 4 Hybrid cloud connectivity & secure configurations

³ Automated deployment, elastic scaling, tools for data management & single vendor for support

LSF + Spectrum Scale



Electronic Design Automation (EDA)

- Optical Proximity Correction
 (OCP)
- Design Rule Checking (DRC)
- Simulations (like timing analysis)



Life Sciences

- Genomic sequencing
- Drug discovery
- Molecular modeling
- Protein docking



Weather

٠

٠

- Climate modeling
- Forecasting and Severe storm prediction
- Ocean Modeling



IBM Cloud HPC | Hybrid Architecture

Goal: Faster time to market with Lower operational cost



A Solution Pattern on IBM Cloud Capabilities

- VPC infrastructure with SGs, ACLs, Transit Gateway, Private DNS
- Physical firewall with HW support for IPSec
- IBM Spectrum LSF
- IBM Storage Scale File System
- Key Protect, Hyper Protect and GKLM for various key management functions
- Schematics and Private Catalogs for controlled provisioning of cloud resources
- Activity Tracker & Security and Compliance Center (SCC) for compliance checking and audit reports
- Code Engine & Notification for enhanced security checking & reporting
- Cloud Object Storage & IBM Storage Scale AFM for transparent tiered storage
- · Event Streams to send events to on-premise
- Context Based restrictions to control access to IBM Cloud & within IBM Cloud
- Custom Roles for fine-grained access control
- Integration with CP4S and/or Qradar for more sophisticated threat Intelligence and user behavior analysis





Using Storage Scale on IBM Cloud for EDA: Cadence Perspective

Gregory Gelles, Storage Solutions Architect Cadence Design Systems

About Cadence Design Systems

30 Years of leadership in EDA & computational software





Provider of complete design flow including design implementation, system & physical verification, system design analysis & packaging



Leader in Hardware & IP for chip design & verification and system design technologies



Cadence Cloud Portfolio



Flexible Business Models | SaaS | Hybrid | Native



Storage related requirements

NON-FUNCTIONAL

Support Scale of both Compute & Storage

Secure Transmission into & out of cloud (IPSec)

Boot drive and Storage encryption with customer managed keys with ability to rotate keys

Detailed storage access logs, both success & failure, with the ability to forward to on-premise SIEM

Security groups, ACLS can set up on a need-to-access basis both for users & infrastructure

Ability to set guard rails around provisioning, & management of resources

Support aggregate performance up to 200,000 IOPS for some of the data sets because of their very high frequency of metadata calls

Avg. Latency less than 2 ms

Support NFS V3, V4 File Protocols

Support quotas for space usage

Support access control

FUNCTIONAL

Support integration with enterprise LDAP

Support hybrid networking with on-premise

Generate fine-grain and coarse-grain performance reports for a time period

Granular role and access control definition

Support monitoring and alerting on different entities including system, nodes, data sets, storage, network

Support Snapshots

Backup solutions can be integrated

Generate audit reports



Storage Scale: Features Liked

Functionality

- Deployable to Bare metals or Virtual Servers
- Expose Storage Scale Native Filesystem via NFS, SMB, COS, HDFS
- Expansion of storage is self managed and can be done on-demand without disruption.
- Integration with LDAP and support more than 16 UNIX groups (NFS v3)
- Encryption
 - Create multiple file systems and different Encryption policies
 - Use External key manager for for Key management via KMIP (CMK)
- Enable auditing on select file systems Infosec
- Auto-incrementing inode limit to accommodate extremely large number of small files.
- Establish quotas on different file sets
- Scale health and audit log forwarding (Centralized logging)
- Control of Meta data and data replication based on reliability and performance needs.
- AFM for data hydration and synchronization (remote cache of FS data / hybrid cloud with no lift/shift)
- Transparent Tiering to Cloud Object Storage to keep only hot data in filesystem and keep overall costs low

Performance & Other Observations

- Performance
 - Various EDA workloads were run across various stages of IC design and verification processes
 - Power & Power integrity Analysis
 - Specialized RFIC parasitic analysis
 - RTL Synthesis & Implementation
 - Physical Sign-off.
 - Logic simulation with machine learning
 - Timing analysis
 - Extraction
 - Performance of these workloads on IBM Cloud with Encryption enabled has been comparable to on-premise turn-around times.
 - More testing with higher number of cores is needed for additional evaluation.
- Complexity of the Scale:
 - Steep Learning curve
 - Mitigated by very responsive help from IBM Cloud team
- Upgrade process
 - Allows both new and old versions to co-exist until commit
 - However, no way to rollback after commit



Storage Scale: Future Plan & Feature Requests

- Future Plan in using Scale Features
 - Multi-tenancy
 - Transparent Tiering of Cloud Object Storage
 - More aggressive Remote caching using AFM
- Feature requests
 - HA levels without 2x data footprint. (We heard about Scale Erasure Code Edition (ECE) is on the roadmap.)
 - More visibility into AFM cached data through command line and UI
 - Ability to adjust trigger levels for inode expansion, right now it very close 100% utilization.
 - Unix df command on a client node should display file system details in the context of their quotas





Thank you! Any Questions?

