

Best Practices for Electronic Design Automation (EDA) Workloads with Amazon FSx for NetApp ONTAP

Optimizing EDA and semiconductor workloads on Amazon Web Services (AWS)

First published Feb 16, 2023



Notices

Customers are responsible for making their own independent assessment of the information in this document. This document: (a) is for informational purposes only, (b) represents current AWS product offerings and practices, which are subject to change without notice, and (c) does not create any commitments or assurances from AWS and its affiliates, suppliers, or licensors. AWS products or services are provided “as is” without warranties, representations, or conditions of any kind, whether express or implied. The responsibilities and liabilities of AWS to its customers are controlled by AWS agreements, and this document is not part of, nor does it modify, any agreement between AWS and its customers.

© 2023 Amazon Web Services, Inc., or its affiliates. All rights reserved.

Contents

Introduction	4
EDA Storage Requirements.....	5
Anatomy of an EDA Cluster	5
EDA Volumes and Directories	6
EDA Data Types	8
EDA Performance.....	12
FSx for ONTAP	14
FSx for ONTAP Introduction	14
FSx for ONTAP Sizing Considerations for EDA Workloads	16
FSx for ONTAP Disk Selection and Throughput Requirements	16
EDA Storage Architecture Patterns on AWS.....	17
Amazon FSx for NetApp ONTAP Setup	24
FSx for ONTAP file system Creation Setup	24
Post FSx for ONTAP Creation Setup	36
Contributors	45
Further reading	45
Document revisions	45

Introduction

NetApp ONTAP-based storage has been a premier choice for customers running semiconductor workloads on-premises. With Amazon FSx for NetApp ONTAP (**henceforth referred as FSx for ONTAP**), AWS and NetApp have collaborated to bring feature rich ONTAP to customers with the operational ease of an AWS managed service. FSx for ONTAP provides fully managed shared storage in AWS Cloud with the popular data access and management capabilities of ONTAP. This document is intended to capture best practices for using FSx for ONTAP for Electronic Design Automation (EDA)/silicon design workloads. The document includes a short background on typical EDA storage use cases and requirements.

EDA workloads are HPC (High Performance Computing) workloads which can see hundreds of thousands of compute cores running parallel jobs on a single project design tree. EDA workloads include frontend verification simulations running tools such as Synopsys VCS, Cadence Xcelium or Siemens EDA ModelSim simulators to backend chip timing analysis jobs like Synopsys Primetime or Fusion Compiler. There could be as many as twenty-five or more EDA design tools used between design specification, design capture, constrained random simulation, synthesis, Place and Route (P&R), Power, Performance, and Area (PPA) optimization, floor planning and chip finishing.

Faster Time-To-Market (TTM) at lower design cost is becoming critical to be relevant in the industry. Consumers expect smaller, faster, and cheaper devices. Factors mentioned above are triggering quicker adoption of advanced silicon manufacturing technologies. Advanced chip designs are constantly demanding increased compute and storage for EDA workloads. Amazon Web Services (AWS) with its virtually infinite capacity and wide selection of compute and storage options, is a natural fit for EDA. EDA vendors, like Synopsys, Cadence, and Siemens EDA are optimizing their EDA workflows to take advantage of cloud and AI techniques. Vendors are utilizing massive parallel computing resources to speed up job completion time, reducing time to results, time to bug identification and fix – realizing better quality tape-outs in shorter design cycles.

EDA storage solutions demand dealing with high metadata and highly parallelized accesses. As customers architect EDA clusters on AWS, this document is intended to recommend set of configuration options for FSx for ONTAP to achieve the best performance at scale for EDA storage.

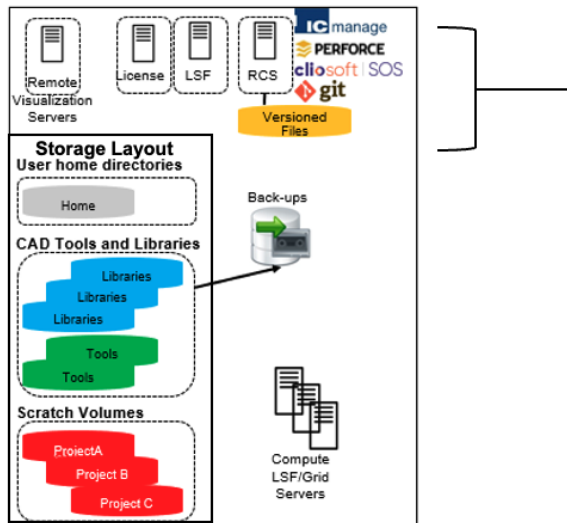
EDA Storage Requirements

Anatomy of an EDA Cluster

Some of our largest customers utilize more than hundreds of thousands of parallel compute cores. High parallelism, high file count, high metadata operations in EDA workloads make optimizing for scale-out performance critical. EDA is not a single workload. Sometimes more than fifty different tools could be used in a single semiconductor chip design project. These tools vary from providing functionality for design specification, verification to timing/power/area optimization and chip finishing.

NetApp has been a data management and storage leader in the semiconductor industry from its earliest days. On premises, NetApp customers use NetApp A700s and A800 class All-Flash storage appliances running NetApp's latest ONTAP storage operating system. Over the years customer workloads have demonstrated that EDA workloads benefit from All-Flash storage due to the ability drive highly concurrent operations at lower latencies. Customer workloads have also shown that scale-out performance of FlexGroup volumes can dramatically improve the number of parallel jobs design teams can run, which in turn leads to faster total turn-around time for EDA jobs. Typical on-premises EDA cluster is depicted in Figure below:

Typical On-Prem Design Infrastructure



- IBM LSF/Grid Servers
 - Bare metal Linux Servers – very uniform configurations
Optimized for compute performance
 - Various Core counts and memory sizes
- Load Sharing (LSF) – manages job queuing
 - Ensures optimal utilization of resources and fair share
 - IBM Spectrum LSF, Univa Grid Engine (Sun Grid Engine), SLURM
- FlexLM License Servers
 - Most common EDA license manager.
 - EDA jobs check-out licenses at runtime. Once all licenses are in use, additional jobs will either wait or fail.
 - After People, EDA tool licenses are the next most expensive component of chip design. EDA license can run from \$5k-\$100k per license depending on the tool.
- Revision Control System (RCS)
 - Perforce and ICMANage are very popular
 - Clio SOS, Synchronicity Design Sync, IBM Clearcase
 - Less common: Git/Artifactory class source and artifact mgmt.
- Data logging, databases, data mining and AI
 - Tools like Splunk, ElasticSearch, etc. for log file mining and reporting
 - Databases for job and metrics collection
 - Increased use of AI for design and design flow optimization
- Remote Desktop Visualization Servers
 - Hummingbird, VNC, Terminal services

Figure-1: Typical On-premises EDA Cluster

Cloud uniquely offers the ability for EDA design flows to scale compute with almost no limit using the latest and greatest compute cores available. As such, underlying storage and data management system must be designed to support cloud-scale parallel workloads.

The recommendations in this document are focused on configuring FSx for ONTAP for maximum parallel compute performance while maintaining latency suitable for EDA.

Customers can cost-optimize their deployments by right-sizing the IOPS and throughput, and enabling intelligent data tiering.

EDA Volumes and Directories

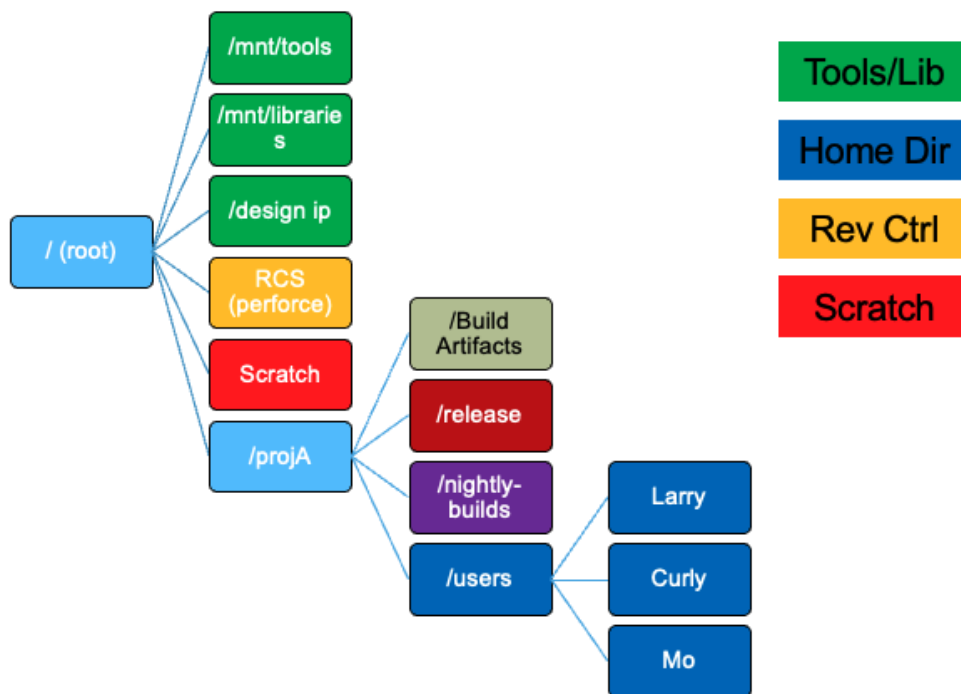
EDA workflows mount and use data from many different volumes in a single workflow. Not all data volumes have the same Read/Write characteristics. Knowing how each volume is uniquely used in EDA flows can provide opportunities for runtime performance and utilization optimization.

For instance, /mnt/tools and /mnt/libs or /mnt/libraries are common storage volume names representing the location of the EDA Tools and Design Libraries used in a design flow. Within these volume(s), there will often be tens to hundreds of tools and tool versions installed. Tools and libraries are often installed, centrally managed by CAD teams, and then shared with multiple teams/ projects.

The typical directories in an EDA environment, along with description of contents and function of each is as show below in Figure-2:

Storage and Directory Organization

Many volumes are used in an EDA flow, each with different IO and data management characteristics



- **Read Only Tools and Libraries Volumes**
 - EDA Tools (Synopsys, Cadence, Siemens, Ansys, etc)
 - Technology Libraries (TSMC, Global Foundry, etc)
 - Design IP (ARM, Synopsys, Cadence, etc)
 - Read-Only - except when new tools or libraries are installed.
 - Caching or Volume replication with DNS mount load balancing can improve read performance.
- **Revision Control System (RCS)**
 - Perforce, ICMange, Clisoft SOS, Clearcase, Synchronicity
 - Like Git, but good for very large and small files, text and binary
 - Combination of local attached SSD (or SAN) for dbase, logging, journaling NFS for versioned file store.
 - Very large volumes of both design files and build artifacts. Can grow up to PB of data.
 - Long code check-out times can be a day-to-day challenge as is multi-site deployments/replication. As build artifacts get larger, so do data transfer times.
- **Build Artifacts**
 - Build artifacts are re-usable outputs of the EDA and Software build process.
 - Netlists, design databases, compile images, generated file lists, log files, reports might all be considered Build Artifacts
 - Management of Build Artifacts is handled different from team to team or company to company. Some version and re-distribute artifacts using tools like Perforce/ICMange. Others store them in a file system.
- **Release Build Volumes**
 - This is where release candidates of the final chip output is generated. This can be one or many volumes. Typically managed by design leads or central build teams. Often automated by tools like Jenkins
 - All or many of the design tools are run against a single copy of the source files
 - Data management is critical – DR, Backup, HA, etc.
- **Nightly or CI Builds**
 - These are regular automated builds typically managed by a design lead or central build teams. These jobs are often automated with CI tools like Jenkins, Bamboo or CircleCI.
 - These builds will use either Scratch volumes or Normal protected volumes.
- **User Workspaces**
 - User workspaces are almost always provisioned on "Scratch Volumes". This is where developer do most of their work creating and editing designs, writing and running tests and iterating to improve design quality and functionality.
 - IT will often provision a large volume and then have multiple users share the volume. Individual quotas are often created to ensure one user does not fill up the whole volume.
- **Note:** Users are typically expected to check available storage capacity before launching new jobs. User level and management reporting is key.

Figure-2: Directory Organization in EDA

IT/Storage teams typically provision storage based on needs of Business Units, Teams, Projects and/or by specific Design flows. Each company's design flows are unique yet similar in a lot of ways. With FSx for ONTAP, customers get the agility and scalability of the cloud and only pay for provisioned storage. Customers do not need to make upfront investments like on-premises which can lead to stranded capacity if business objectives change.

EDA Data Types

Tool and Library Volumes – Optimized for Read Performance

EDA tools like those from Synopsys, Cadence, and Siemens EDA are typically installed in tool directories mounted like `/mnt/tools/synopsys`, `/mnt/tools/cadence`, etc. Each of these installation directories can contain many different tools and many different tool versions.

Similarly, 3rd party design technology libraries are typically installed in directories mounted like `/mnt/libs/tsmc`, `/mnt/libs/arm`, `/mnt/libs/dolphin`, etc. These libraries contain either design IP

(like a USB or CPU core) or they could contain EDA vendor libraries from device manufactures like TSMC, Intel, Global Foundries, UMC or Samsung which describe the physical power, area, and performance data associate transistor/cell representation of the process node. Tech libraries are typically counted in thousands of files containing an ever-increasing amount of data.

Internal libraries or design IP are another class of data used in designs. These might be subsystems or blocks developed by one team within the company and delivered in another team for inclusion in a full chip design. For instance, an ARM core designed by the processor team might be delivered and used by the team who will integrate the core into a mobile device.

The key characteristic of tools and libraries is that the data is typically READ-ONLY or read-mostly. Writes occur when new versions of tools or libraries are installed, but seldom are tools and libraries over-written. New tool or library versions are typically installed next to the prior versions.

As such, we can take advantage of this by optimizing the infrastructure for read performance. This might include NFS mount options to maximize compute server data caching to improve IO latency.

Tools and Libraries will be accessed by the EDA design flows and may see very high concurrent accesses from the 10k-100k parallel jobs simultaneously running in the compute farm. As such, IO performance and scale is critical. Tools like Synopsys PrimeTime, timing analysis tool, will load thousands of library files simultaneously from thousands of servers resulting in high read-loads and may result in performance bottlenecks. The reference architecture, proposed in this document, will address how to use FlexCache to improve read performance by load balancing NFS mounts across multiple FlexCache tool and library volumes. Figure-3 below depicts the above-mentioned directories.

EDA Tools (/mnt/tools/...)

```
jmichael@JMICHAEL101-PC:~/apple_csg$ find tools/
tools/
tools/synopsys
tools/synopsys/primetime_2020.09.1
tools/synopsys/primetime_2021.06.2
tools/synopsys/primetime_2020.03.2
tools/synopsys/many_other_tools_and_versions
tools/mentor
tools/mentor/many_other_tools_and_versions
tools/cadence
tools/cadence/xcelium_2021.09.1
tools/cadence/xcelium_2022.06.1
tools/cadence/xcelium_2021.09.3
tools/cadence/voltus_2021.03.1
tools/cadence/xcelium_2020.10.4
tools/cadence/many_other_tools_and_versions
tools/cadence/voltus_2022.06.1
tools/cadence/voltus_2021.09.1
```

EDA Libraries (/mnt/libraries/...)

```
jmichael@JMICHAEL101-PC:~/apple_csg$ find libraries
libraries
libraries/tsmc
libraries/tsmc/c05
libraries/tsmc/c05/stdcell_lib_c05_and_many_more_versions
libraries/tsmc/c03
libraries/tsmc/c03/stdcell_lib_c03_and_many_more_versions
libraries/tsmc/c10
libraries/tsmc/c10/stdcell_lib_c10_and_many_more_versions
libraries/tsmc/c07
libraries/tsmc/c07/stdcell_lib_c07_20220101
libraries/tsmc/c07/stdcell_lib_c07_and_many_more_versions
libraries/synopsys
libraries/synopsys/usb
libraries/synopsys/usb/usbcore_many_more_versions
libraries/synopsys/mac
libraries/synopsys/mac/mac_many_versions
libraries/arm
libraries/arm/arm10
libraries/arm/arm10/arm10_many_versions
libraries/arm/arm10/arm_many_versions
libraries/arm/arm11
libraries/arm/arm11/arm_many_versions
```

Figure-3: Tools and Libraries Directory structure in EDA

Scratch Volumes – Optimized for Maximum Read/Write Performance

Scratch volumes are used for user workspace, nightly regressions, CI/CD build flows and many other development purposes. The term “Scratch” comes from the idea that the data is temporary or has such a high change rate that traditional back-up and archiving is unnecessary.

Scratch is where active development is happening. Engineers launch EDA jobs and wait for them to complete so they can check to see if a test passed or failed after a design change or a new functional test. For example, designers might be waiting for the results of 1000s of jobs to check the PPA (power performance and area) characteristic of the chip after a functional change or optimization. Performance is critical and both storage capacity and compute scale are critical attributes.

Scratch volumes should be optimized for maximum read/write performance. Modern chips design flows can use as many as 100 different tools leading a wide range of IO profiles (reads to writes ratios, big vs small files, sequential vs random access) and may vary dramatically. These environments are often shared by multiple flows and multiple teams. EDA scratch area snapshot is show below in Figure-4:

```

jmicahel@JMICHAE101-PC:~/apple_csg$ tree -L 4
.
├── mx_chip
│   ├── nightly_builds
│   │   ├── sims_20220802
│   │   │   └── mx_chip
│   │   ├── sims_20220803
│   │   │   └── mx_chip
│   │   └── sims_20220804
│   │       └── mx_chip
│   ├── release_builds
│   │   ├── build_20220901
│   │   │   └── mx_chip
│   │   ├── build_202209015
│   │   │   └── mx_chip
│   │   ├── build_202209015a
│   │   │   └── mx_chip
│   │   ├── build_202209015b
│   │   │   └── mx_chip
│   │   └── build_20220908
│   │       └── mx_chip
│   └── users
│       ├── curly
│       │   ├── devtest_20220812 ← $PROJECT_DIR (top of a design tree)
│       │   ├── devtest_20220815
│       │   └── devtest_20220817
│       ├── larry
│       │   ├── workspace1
│       │   ├── workspace2
│       │   └── workspace_n
│       └── mo
│           ├── workspace1
│           ├── workspace2
│           └── workspace_n
└── ...

```

```

jmicahel@JMICHAE101-PC:~/apple_csg/mx_chip/users/curly$ tree devtest_20220812
devtest_20220812
├── mx_chip ← $PROJECT_DIR (top of a design tree)
├── arm_core
├── audio_ctrl
├── bus_interconn
├── chiptop
│   ├── netlist
│   │   ├── chiptop.flat.netlist.v
│   │   ├── chiptop.netlist.v
│   │   ├── chiptop_bist.netlist.v
│   │   ├── chiptop_io.netlist.v
│   │   └── chiptop_test.netlist.v
│   ├── rtl
│   │   ├── chiptop.v
│   │   ├── chiptop_bist.v
│   │   ├── chiptop_io.v
│   │   └── chiptop_test.v
│   └── simulation
│       ├── results
│       │   ├── many_dirs_and_files
│       │   ├── test1
│       │   │   ├── test1.debug-fsdb_large_waveform_file
│       │   │   ├── test1.errors
│       │   │   └── test1.log
│       │   ├── test2
│       │   ├── test3
│       │   └── test_n
│       ├── testbench
│       │   └── chiptop_tb.v
│       └── tests
├── synthesis
├── timing
│   ├── chiptop.sdc
│   ├── chiptop_bist.sdc
│   ├── chiptop_io.sdc
│   └── chiptop_test.sdc
├── results
│   └── many_dirs_and_files
├── ethernet
├── i_core_b
├── i_core_c
├── ip_core_a
├── memctrl
├── pcie
├── tools
├── usb
└── video_ctrl

```

Figure-4: Typical EDA scratch space directory hierarchy

Build Artifacts

Build Artifacts can be any number of generated data from the EDA or software development workflow. A chip or IP netlist generated from Register Transfer Level (RTL) code using Synthesis tools is a build artifact. A generated gate level netlist might be reused by the verification team for Gate Level simulations, reused by the Emulation team, reused by the timing analysis and floor planning team to name a few. Modern semiconductors are made up of both Soft, Firm and Hard IP which is developed by one team and delivered to another team for integration into a design. The artifacts of an IP or subsystem development team might become reused IP deliverables provided to another team for integration.

Companies manage and share build artifacts differently. Some teams use design data management tools for revision control and intellectual property (IP) management tools to maintain build artifacts versions and distribute IP to teams and development sites. Other

companies distribute build artifacts by replicating data volumes or directories. Others use symbolic links to shared files and directories into development areas. Build artifact volumes are typically designed for balanced R/W. Many NetApp customers use data replication technologies like SnapMirror or FlexCache to distribute artifacts across multiple sites.

EDA Performance

EDA workloads demand high IOPS and throughput from its shared network storage. The more performant a file system is, the more EDA job can be run in parallel and the faster design teams can find bugs, iterate on design optimization, etc. The sizing goal is to optimize the configuration for the maximum possible IOPS and throughput.

EDA workloads are HPC workloads and require maximum single job runtime performance – shortest possible wall clock time to job completion. Modern EDA workflows will have 20k to 100k compute cores running in parallel on a single design (set of volumes). Optimizing FSx for ONTAP to maximize the number parallel cores is critical. Storage IO performance should be carefully monitored as the number of parallel jobs is increased, which can be observed by latency rapidly increasing above 3 ms. – this is observed as the knee of the performance curve.

Single Job Performance measured by Wall Clock time

EDA job performance is typically measured in wall-clock time (i.e., duration from the start and end of a job as measured by an end-user). How long did it take to run a single chip design simulation? How long did it take to complete my timing analysis or floor plan?

An improvement in job runtime is critical since it improves the productivity and utilization of design engineers and EDA licenses. Engineers and EDA licenses make up over 90% of chip development costs. Wall clock improvements of 5-10% are considered huge wins and are worth spending money on. EDA infrastructure cost is only ~10% of the chip development costs and storage is roughly 4% of the IT spend – so spending more on storage performance pays for itself if it improves the productivity and utilization of the engineers and licenses (the other 90%).

Performance measured by Total Turn-around Time (TTT) of All Jobs

EDA performance is also measured in “Total Turn-around Time (TTT). This is the time it takes for a collection of jobs to complete in its entirety. For instance, how long did it take for all 1000 simulations to run start to finish?”. The answer to that question is highly dependent on whether the jobs tests were run serially or parallelly or some combination. Load sharing tools like LSF, Grid Engine, Slurm, etc. allow engineers to submit large sets of jobs to run in parallel. The more jobs that can run in parallel, typically the faster all the jobs complete.

Example: Running 20k jobs

	Total Jobs	Max parallel jobs	Ave Job Runtime (min)	TTT (min)	TTT (hrs.)
Example 1	20,000	1,000	10	200	3.33
Example 2	20,000	5,000	10	40	0.67 (5x faster)

As you can see from the table above, running 5000 jobs in parallel results in a 5x improvement in job turn-around time, which in turns means 5x faster to find a bug, 5x faster design closure iterations, 5x more productive engineers.

Optimizing for maximum parallel performance and scale is a key when considering EDA design architectures. Cloud computing offers the opportunity of unlimited parallel compute scale since the cloud is vast and elastic vs the fixed resources in traditional datacenters. And in turn Cloud offers the opportunity for significant improvement in Total Turn-around Time (TTT) and productivity.

The ability to run more jobs in parallel is as important a consideration as optimizing single job wall clock time performance.

FSx for ONTAP

FSx for ONTAP Introduction

[Amazon FSx for NetApp ONTAP](#) is a fully managed service that provides highly reliable, scalable, high-performing, and feature-rich file storage built on NetApp's popular ONTAP file system. FSx for ONTAP combines the familiar features, performance, capabilities, and API operations of NetApp file systems with the agility, scalability, and simplicity of a fully managed AWS service.

FSx for ONTAP provides feature-rich, fast, and flexible shared file storage that's broadly accessible from Linux, Windows, and macOS compute instances running in AWS or on premises. FSx for ONTAP offers high-performance solid-state drive (SSD) storage with sub-millisecond latencies. FSx for NetApp ONTAP file system has two storage tiers: primary storage and capacity pool storage. Primary storage is provisioned, scalable, high-performance SSD storage that's purpose-built for the active portion of your data set. Capacity pool storage is a fully elastic storage tier that can scale to petabytes in size and is cost-optimized for infrequently accessed data. On each volume, data is automatically tiered to the capacity pool storage when automatic data tiering is enabled on a volume. The volume's tiering policy determines when data is tiered off to capacity pool storage. With FSx for ONTAP, you can achieve SSD levels of performance for your workload while paying for SSD storage for only a small fraction of your data.

Managing your data with FSx for ONTAP is easier because you can snapshot, clone, and replicate your files with the click of a button. In addition, FSx for ONTAP automatically tiers your data to lower-cost, elastic storage, lessening the need for you to provision or manage capacity.

FSx for ONTAP also provides highly available and durable storage with fully managed backups and support for cross-Region disaster recovery. To make it easier to protect and secure your data, FSx for ONTAP supports popular data security and antivirus applications.

For customers who use NetApp ONTAP on-premises, FSx for ONTAP is an ideal solution to migrate, back up, or burst your file-based applications from on-premises to AWS without the need to change your application code or how you manage your data.

As a fully managed service, FSx for ONTAP makes it easier to launch and scale reliable, high-performing, and secure shared file storage in the cloud. With FSx for ONTAP, you no longer have to worry about:

- Setting up and provisioning file servers and storage volumes
- Replicating data
- Installing and patching file server software
- Detecting and addressing hardware failures
- Managing failover and failback
- Manually performing backup

FSx for ONTAP also provides rich integration with other AWS services, such as AWS Identity and Access Management (IAM), Amazon WorkSpaces, AWS Key Management Service (AWS KMS), and AWS CloudTrail.

FSx for ONTAP – Deployment Types

FSx for ONTAP can be deployed in Single-AZ or Multi-AZ configurations.

The traditional EDA/Semiconductor datacenter contains compute and storage in the same building. If there is a failure in the datacenter, the compute and storage can be affected. This might be a result of a network outage, power outage or any other issue. The EDA data stored on the NFS filers are typically not replicated in real time to secondary datacenters. Some EDA volumes may be replicated or backed-up to another datacenter for disaster recovery several times per day, but not continuously.

The FSx for ONTAP Single-AZ deployment type is analogous to a HA pair deployed in a traditional datacenter described above. Single-AZ FSx for ONTAP file systems replicate your data and offer automatic failover within a single Availability Zone.

The FSx for ONTAP Multi-AZ deployment type on the other hand is more like a “MetroCluster” in NetApp speak. MetroClusters are two NetApp filers in two different datacenters separated by a distance (same city or adjacent city). Every write that comes in on one filer is immediately sent to the second filer. Only once both filers acknowledge the write is the acknowledgement sent back to the computer which initiated the write. MetroCluster is typically used where there is

no tolerance for outage – failover from one datacenter to another must be seamless, immediate, and completely fool proof.

Multi-AZ file systems provide added resiliency by replicating your data and supporting failover across multiple Availability Zones within the same AWS Region. Since writes are committed across AZ's, multi-AZ file systems have slightly higher latencies when compared to single AZ file systems-

FSx for ONTAP Sizing Considerations for EDA Workloads

Modern EDA workloads will utilize 20k to 100k parallel compute cores on a single design. High parallelism, high file count, high meta data workloads make optimizing for scale out performance critical.

There is no single EDA workload. As many as 50+ different tools may be used to design chips, from the early design specification stage, to design, timing/power/area optimization and chip finishing. The recommendations are focused on configuring FSx for ONTAP for maximum IOPS, minimum latency, and maximum parallel compute performance.

Infrastructure cost is always a consideration, but productivity of engineers and utilization of expensive EDA design tool licenses are much larger expense than the IT/Cloud infrastructure costs. That said, once performance requirements are understood, then performance vs cost trade-offs can be assessed.

FSx for ONTAP Disk Selection and Throughput Requirements

All-Flash (SSD based) disk performance has demonstrated its value for the demanding EDA/HPC workloads. Sub 1ms latencies have improved performance of EDA workload and has almost completely removed the need for performance tuning storage as was common practice for 10k SAS and SATA drives.

SATA drives are still common in EDA, but only for backup and cold data storage. Active workloads are almost exclusively on high performance All-Flash storage platforms on-prem. The same is true for cloud based NFS storage. In a do-it-yourself NFS storage solution on AWS, pairing disk

storage throughput performance to match the EC2 instance throughput is critical to maximizing workload performance.

FSx for ONTAP eliminates the above-mentioned operational undifferentiated heavy lifting by providing NetApp ONTAP as a fully managed AWS managed service. It provides feature-rich, fast, and flexible shared file storage in two tiers primary pool for performance and capacity pool, cost optimized for infrequently accessed data.

Amazon FSx for NetApp ONTAP provides sub-millisecond file operation latencies with solid state drive (SSD) storage, and tens of milliseconds of latency for capacity pool storage. Above that, FSx for ONTAP has two layers of read caching on each file server—NVMe (non-volatile memory express) drives and in-memory—to provide even lower latencies when you access your most frequently-read data. Each FSx for ONTAP file system provides up to multiple GB/s of throughput and hundreds of thousands of IOPS. These features make FSx for ONTAP a great solution for EDA on AWS.

EDA Storage Architecture Patterns on AWS

While not a comprehensive list, EDA architectural patterns broadly fit into the following three categories:

- Full chip design on AWS
- Burst capacity onto AWS
- Backup/Disaster Recovery (DR) on AWS

ONTAP Data Management - Full Chip on AWS

To design full chip on AWS – primary volumes, snapshots, and backups, all reside on AWS. In addition to provisioning primary volume/s on FSx for ONTAP, native snapshot and backup capabilities can be used to architect for disaster recovery (DR) and business continuity.

Furthermore, NetApp’s SnapMirror/SnapVault data replication tool is also available out-of-the-box on FSx for ONTAP. This feature, already familiar to many customers, allows replication relationships to be established across different AWS regions, if needed.

Business continuity requirements, Recovery Time Objective (RTO) and Recovery Point Objective (RPO) will determine the best approach to architecting DR for full chip on AWS.

NetApp’s FlexCache technology can be used to lazy load active data on-demand onto a FSx for ONTAP file system from another file system. The feature enables performant, cross AWS-region EDA. The methodology can also be used cross-AZ or within the same AZ to horizontally scale file systems. FlexCache can be created from primary or DR volumes.

The architectural patterns for full chip on AWS are depicted in Figure-5(a) below:

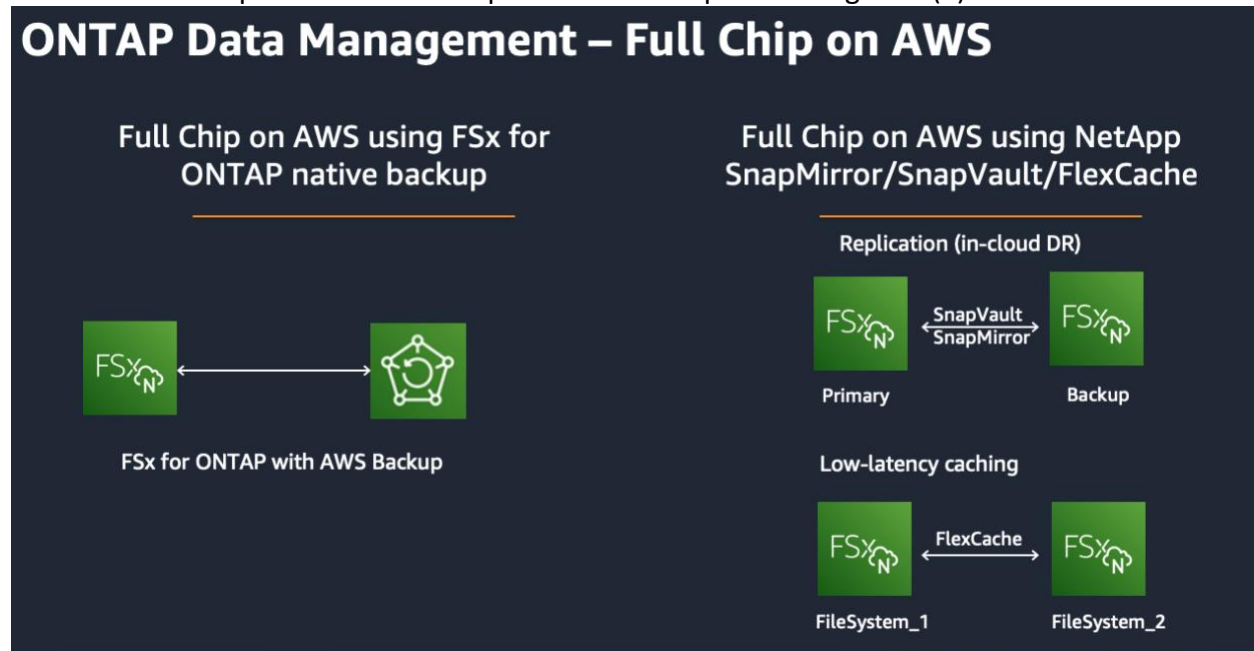


Figure-5(a): ONTAP Data Management for Full Chip on AWS

ONTAP Data Management - Burst capacity onto AWS

Hybrid EDA – Burst compute to AWS

Several customers have used AWS to serve EDA burst capacity. Complementing on-premises compute, virtually Infinite compute capacity on AWS has made it possible for EDA customers to enable faster chip tape-outs. In this architecture data located on premises can be cached remotely by an FSx for ONTAP file system using FlexCache. An example of this would be an Electronic Design Automation (EDA) tools library stored on premises in Austin, TX, and remote users in Singapore that need low latency access to those tools. By utilizing a FlexCache volume on FSx for ONTAP file system in Singapore linked to the origin volume in on-premises NetApp, customers can burst their workflow to leverage compute in Singapore and provide low-latency access to cached data for end users and compute. FlexCache supports both NFS and SMB file caching. Files can be cached between ONTAP file systems if desired.

Hybrid EDA - Storage on AWS with On-Premises Caching

Though most customers prefer to migrate compute and storage to AWS - inverse of the cloud bursting use case we just reviewed is also an architecture made possible with FSx for ONTAP. In this model, the tools/library resides on an FSx for ONTAP file system. Using a FlexCache volume on premises linked to the source volume on FSx for ONTAP file system, customers can cache data on premises for low-latency access. This architecture could be useful for sharing EDA directories (e.g., Tools/Library releases) across geos with AWS, when connectivity between on-premises sites does not exist or does not enable sharing with ease.

Architectural patterns discussed above for modeling EDA burst on AWS are depicted in Figure-5(b) below:

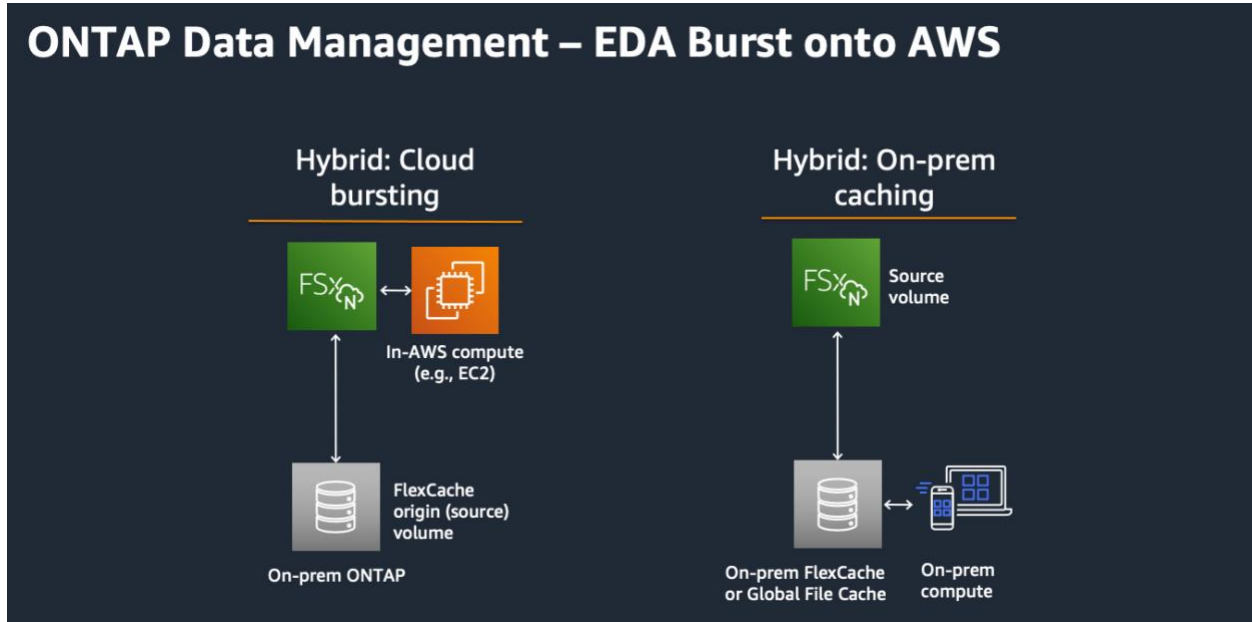


Figure-5(b): ONTAP Data Management for EDA Burst onto AWS

ONTAP Data Management – DR/Archive on AWS

NetApp ONTAP is widely deployed as an on-premises EDA storage solution. With FSx for ONTAP, customers can use already familiar ONTAP features for data replication. Hence storing EDA Backup, DR, and Archive data on cost-efficient FSx for ONTAP is a popular starting point for EDA on AWS. ONTAP's SnapVault replicates snapshots from on-premises NetApp storage array

to FSx for ONTAP file system for longer term retention on the destination. SnapMirror replicates the active file system and snapshots for DR purposes. Data can also be replicated across AWS regions/AZs with SnapVault and SnapMirror. DR architecture on AWS is depicted below in Figure-5(c).

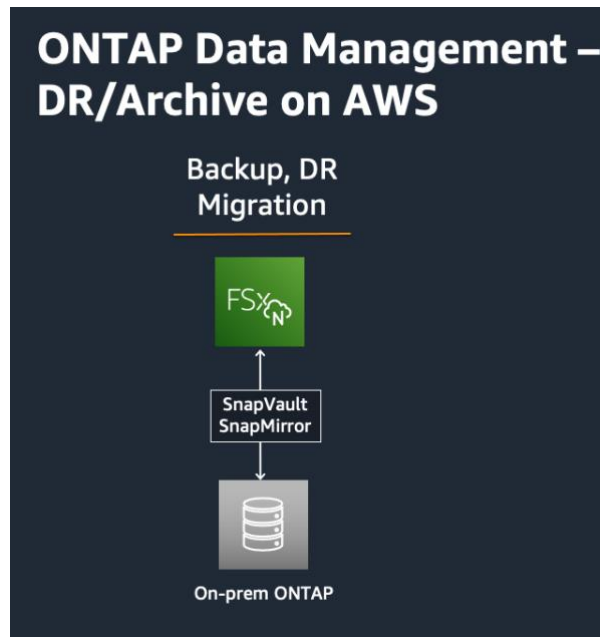


Figure-5(c): ONTAP Data Management for EDA Burst onto AWS

Multi-Site EDA Architectural Patterns

FSx for ONTAP helps realize complex EDA architectures with relative ease. For the purposes of this document, a 'site' could be a single on-premises data center facility or a single AWS region. Hence, a multi-site architecture could comprise of a combination of on-premises data center/s and one or more AWS region/s. Illustrated below in Figure-6a & b is Multi-Site EDA architecture using AWS for multi-project, multi-team organizations:

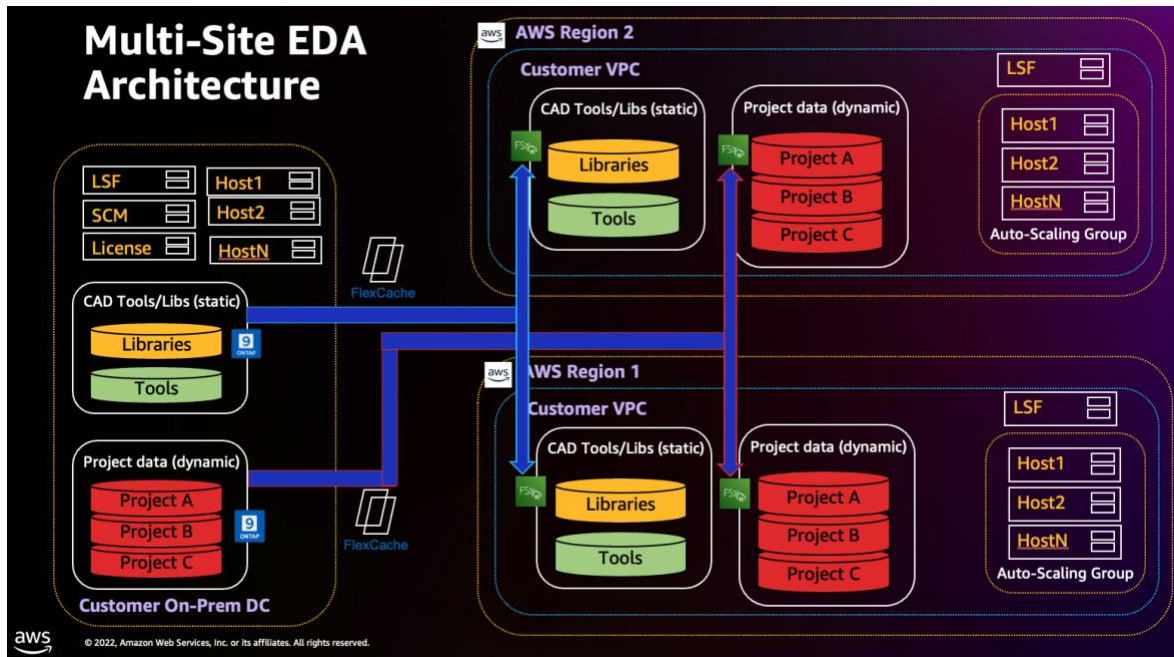


Figure-6a: Multi-Site FSx for ONTAP with on-premises architecture

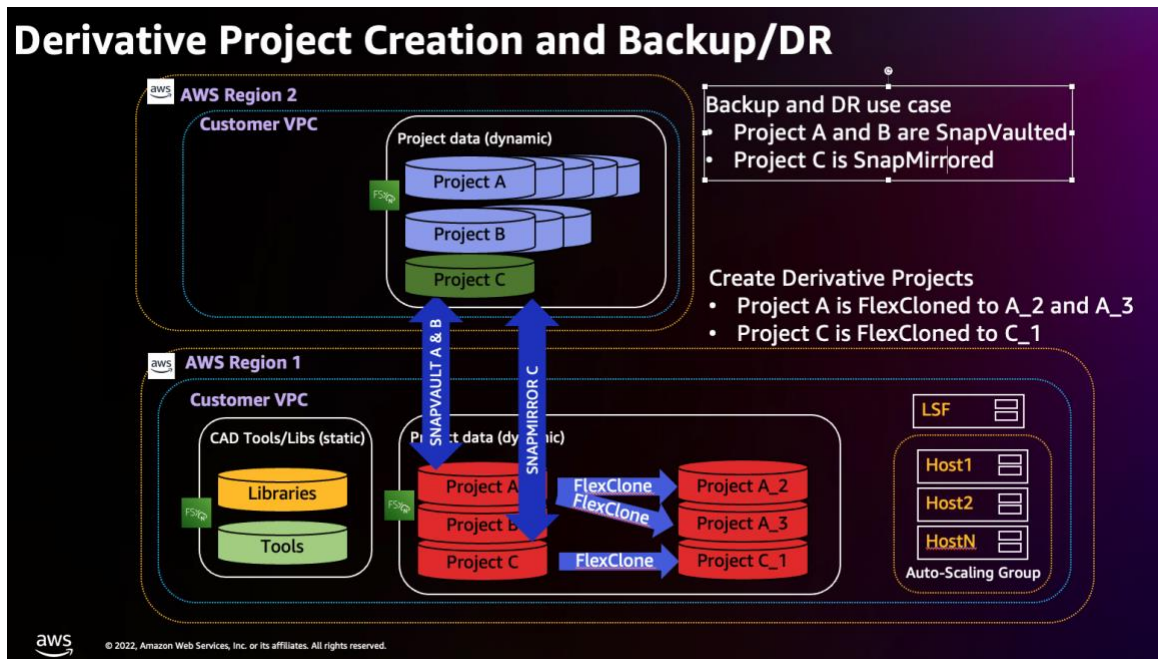


Figure-6b: Multi-Site (Multiple AWS regions) architecture using FSx for ONTAP**NOTE: A *derivative project* is one which starts from data from previous project as against all-new fresh design data. More specifically, the EDA directories (Project, Tools/ Libraries) will be replicated between either 2 AWS regions or on-premises and AWS as depicted in Figure-7 below:

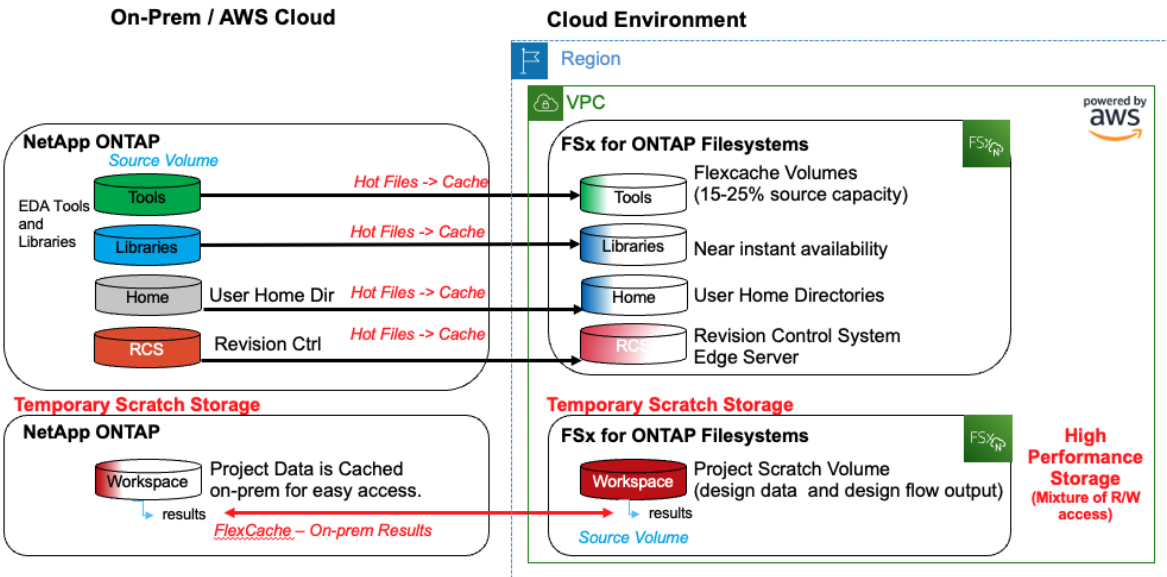


Figure-7: EDA Directory Replication using ONTAP technologies in Multi-Site Setup

Multi-Region Data Management Architecture at Enterprise Scale

Scaling Tools and Library

Tools and libraries are typically shared resources on premises and therefore should also be in the cloud. Advanced semiconductor manufacturing processes (7nm, 5nm and 3nm) increase design complexity and scale. Design technology libraries have become very large as the complexities of representing transistor performance, power and geometries have caused the files to grow significantly. As such library loading times have become a challenge.

The demand on the library volumes when tool flows start-up on 10k-100k of server's cores simultaneously can create a "boot storm" on the storage filers leading to high read latency conditions. This is further compounded by sharing Tools and Library volumes across multiple design flows, projects, and teams. Scaling FlexCache volumes to spread the read load out across the parallel compute cores can dramatically improve library and tool load times.

It is recommended to configure Tools and Library volumes (FSx for ONTAP file systems) in a multi-AZ configuration. This configuration provides resiliency against FSx for ONTAP node and AZ failure. It is best practice to replicate Tools and Libraries across multiple FSX for ONTAP file systems to achieve performance at scale and increased availability. In addition, it is

recommended to use FSx for ONTAP native backup for Tools/Libs volumes for improved reliability, durability, and uptime.

Scratch Volumes at Scale

Scratch volumes require maximum READ/WRITE performance. Best practice is to use FlexGroup volume on a Single-AZ FSx for ONTAP file system. FSx for ONTAP file system performance metrics is provided [here](#) for reference.

FSx for ONTAP file systems should be deployed on a per flow, team, or project basis. Figure-8 below shows how shared FSx for ONTAP file systems can be scaled horizontally by using FlexCache volumes for tools, libraries, and other read-heavy workloads. As more and more EC2 instances are elastically scaled up to meet increasing design workload requirements, the Tool and Library FlexCache volumes can be quickly scaled-up and UNIX mount points DNS load balanced to optimize performance. As EC2 instance are scaled backdown, FlexCache volumes can be scaled back as well.

Cloud Architecture for Multi-Project / Multi-Team Deployments

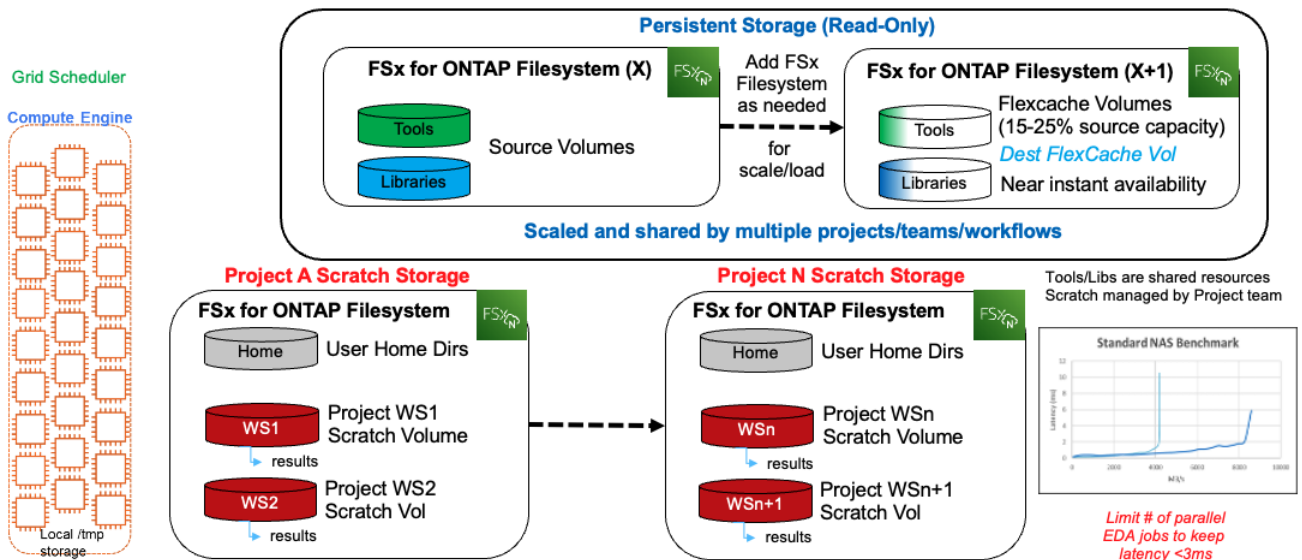


Figure-8: FSx for ONTAP architecture at scale for EDA directories

Amazon FSx for NetApp ONTAP Setup

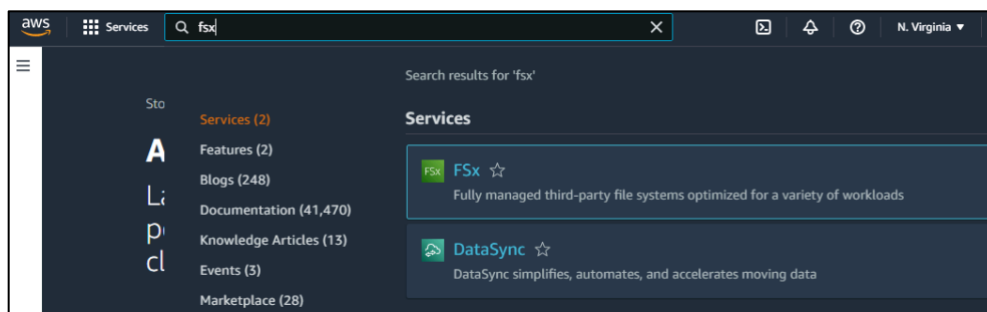
The purpose of this section is to document the current best practices for optimizing FSx for ONTAP for high-performance, large-scale EDA workloads.

FSx for ONTAP file system Creation Setup

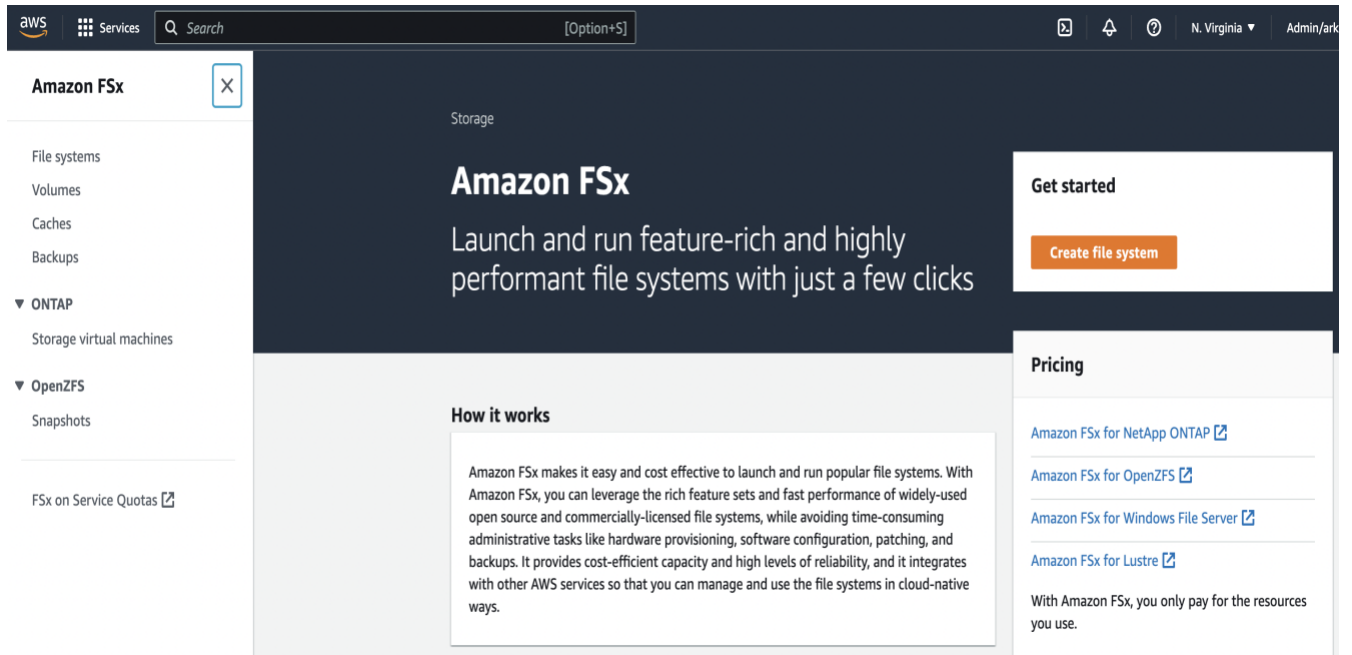
- [FSx for ONTAP file system for scale-out Tool/Library FlexCache volumes](#)
This configuration is optimized for very high Read-Only performance as well as high availability.
- [FSx for ONTAP Optimized for Maximum IO for Scratch volumes](#)
This configuration is optimized for max IO throughput for R/W performance at scale.

Open FSx Console

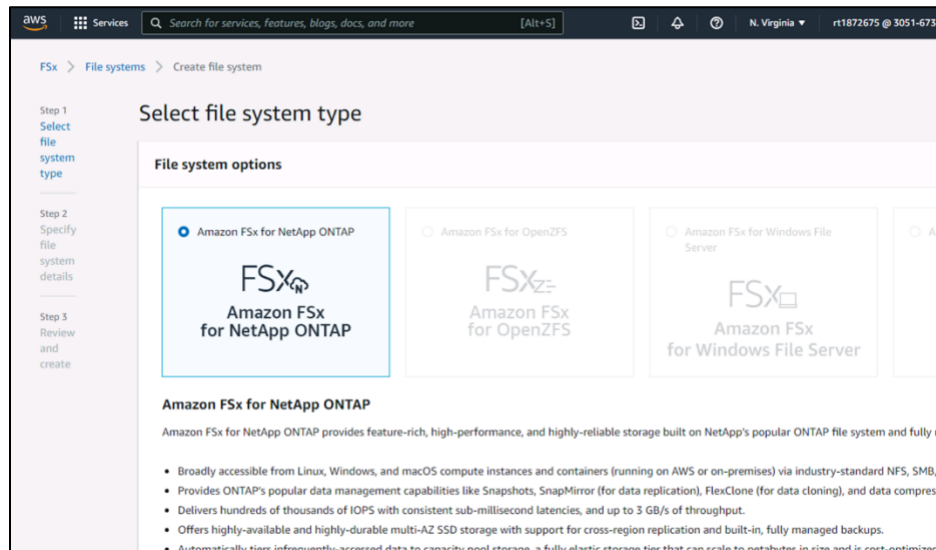
Open the FSx for ONTAP console by clicking [here](#)



Click on the FSx service under Services.



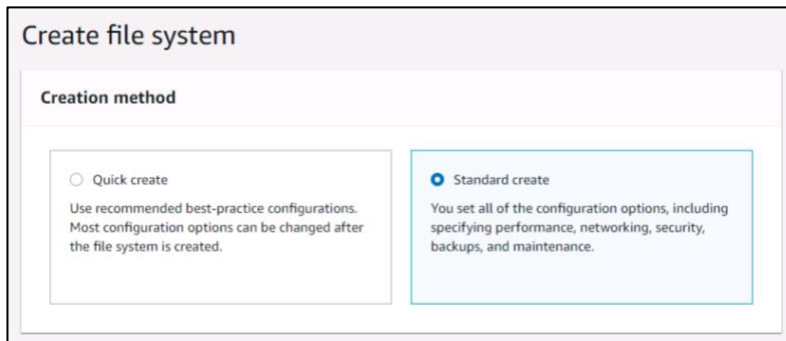
Create a new FSx for ONTAP file system.



1. Create new FSx file system
Select -> **Create file system**
then
Select -> **Amazon FSx for NetApp ONTAP**

2. Creation Method

Select -> **Standard create**



The screenshot shows a console window titled "Create file system". Under the heading "Creation method", there are two options:

- Quick create
Use recommended best-practice configurations. Most configuration options can be changed after the file system is created.
- Standard create
You set all of the configuration options, including specifying performance, networking, security, backups, and maintenance.

Choosing Standard create will provide additional customization options needed for optimal EDA performance setup.

File system details

1. Specify a File system name

This setting is optional, but best practice is to provide a meaningful name for the file system.

File system details

File system name - optional [Info](#)
eda_scratch_filer

Maximum of 256 Unicode letters, whitespace, and numbers, plus + - = . _ : /

Deployment type [Info](#)

Multi-AZ

Single-AZ

SSD storage capacity [Info](#)

81920 GiB

Minimum 1024 GiB; Maximum 192 TiB.

Provisioned SSD IOPS

Amazon FSx provides 3 IOPS per GiB of storage capacity. You can also provision additional SSD IOPS as needed.

Automatic (3 IOPS per GiB of SSD storage)

User-provisioned

160000

Maximum 160,000 IOPS

Throughput capacity [Info](#)

The sustained speed at which the file server hosting your file system can serve data. The file server can also burst to higher speeds for periods of time.

Recommended throughput capacity
1024 MB/s

Specify throughput capacity

Throughput capacity

4096 MB/s

2. Select single-AZ deployment type for scratch volume, multi-AZ for tools/libs/projects

Note: Both types of deployment (Single- & Multi-AZ) support local NVMe for read caching for detail <https://docs.aws.amazon.com/fsx/latest/ONTAPGuide/performance.html>

Specify Capacity and Storage Performance Requirements

For EDA workloads, best practice for peak performance is to start with FSx for ONTAP file system configured for maximum IOPS and Throughput performance, then to test workloads at scale. If

the workloads drive higher volume latencies than desired on a FSx for ONTAP file system provisioned with maximum IOPS and throughput, then you may need to load balance workloads across multiple file systems. Cost optimization opportunity might exist if file system performance is above expectation. Provisioned SSD IOPS and throughput capacity properties can be increased or decreased after file system creation to save storage costs.

It is typically easier to run more EDA jobs and overload the file system to find the maximum load the largest FSx for ONTAP file system can support, then scale back the number of EDA jobs running in parallel. If the system is not sized performant enough, the FSx for ONTAP volume metrics in CloudWatch will show high volume latencies which impact the total number of parallel jobs that can run before hitting the knee of the latency curve.

The recommendation is to start at full scale performance since in most cases file system performance will be the limiting constraint for EDA workloads.

3. Provision Required SSD Storage Capacity

Specify capacity in GiB. Maximum 192 TiB

Example: To create an 81 TiB volume (81 TiB x 1024 GiB), type 82944 in the SSD storage capacity dialog box.

4. Specify Provisioned SSD IOPS

Select **User-Provisioned IOPS**

Then type **160000** into the **User-provisioned** textbox. 160,000 IOPs ([Check AWS documentation for region availability](#)) the max allowed value.

Note: Min 5TB space needs to be provided to choose 160K IOPS

5. Specify Throughput capacity

Select -> **Specify throughput capacity**

Then Select -> **4096 MB/s** from the **Throughput capacity** pulldown list. This is the maximum allowed value

File system details

File system name - optional [Info](#)

Maximum of 256 Unicode letters, whitespace, and numbers, plus + - = . _ : /

Deployment type [Info](#)

- Multi-AZ
 Single-AZ

SSD storage capacity [Info](#)

 GIB

Minimum 1024 GiB; Maximum 192 TiB.

Provisioned SSD IOPS

Amazon FSx provides 3 IOPS per GiB of storage capacity. You can also provision additional SSD IOPS as needed.

- Automatic (3 IOPS per GiB of SSD storage)
 User-provisioned

← Select Maximum User Provisioned IOPS

Maximum 160,000 IOPS

Throughput capacity [Info](#)

The sustained speed at which the file server hosting your file system can serve data. The file server can also burst to higher speeds for periods of time.

- Recommended throughput capacity
1024 MB/s

Specify throughput capacity

Throughput capacity

← Select Maximum Throughput Capacity

Network and security

No EDA specific best practices requirements. Options are specific to the VPC design.

- VPC
- VPC Security Group
 - Refer [documentation](#) for VPC security group configuration
- Subnet

Network & security

Virtual Private Cloud (VPC) [Info](#)
Specify the VPC from which your file system is accessible.

Core_VPC_rt1872675 | vpc-00e4d3554a3562dfc ▼

VPC Security Groups [Info](#)
Specify VPC Security Groups to associate with your file system's network interface.

Choose VPC security group(s) ▼

sg-0c7645c02caa689d6 (default) ✕

Subnet [Info](#)
Specify the subnet in which your file system's network interface resides.

privateSN_us-east-1c_rt1872675 | subnet-0607dd12889d380ad (us-east-1c) ▼

Security and encryption

FSx for ONTAP supports nitro-based encryption in transit when accessed from nitro instances in supported regions. This encryption happens without any performance impact. Hence nitro-based instances are recommended for customers requiring encryption in transit.

FSx for ONTAP also supports Kerberos-based encryption is supported on FSX for ONTAP with NFSv4. NFSv4 is less performant compared to NFSv3 for EDA workloads. Use of Kerberos has further performance impact.

Further information on Encryption of Data in transit is available on [FSx for ONTAP documentation](#).

Security & encryption

Encryption key [Info](#)
 AWS Key Management Service (KMS) encryption key that protects your file system data at rest.

aws/fsx (default) ▼

Description	Account	KMS key ID
None	None	None

File system administrative password
 Password for this file system's "fsxadmin" user, which you can use to access the ONTAP CLI or REST API.

Don't specify a password
 Specify a password

Password

Confirm password

Default storage virtual machine (SVM)

1. Specify a Storage virtual machine name

Select -> **Storage virtual machine name** textbox

Then

Type a valid SVM name

2. SVM Administrative password

There is no EDA specific recommendation for SVM password settings.

3. Active Directory

There are no EDA specific recommendations for Active Directory.

Default storage virtual machine configuration

Storage virtual machine name

SVM administrative password
Password for this SVM's "vsadmin" user, which you can use to access the ONTAP CLI or REST API.

Don't specify a password
 Specify a password

Active Directory
Joining an Active Directory enables access from Windows and MacOS clients over the SMB protocol.

Do not join an Active Directory
 Join an Active Directory

Default volume configuration

1. Volume Configuration

Fill in the desired Volume name, Junction path, volume size as shown below.

Default volume configuration

Volume name

Maximum of 203 alphanumeric characters, plus _ .

Junction path

The location within your file system where your volume will be mounted.

Volume size

Minimum 20 MiB; Maximum 104857600 MiB

Storage efficiency
Select whether you would like to enable ONTAP storage efficiencies on your volume: deduplication, compression, and compaction.

Enabled (recommended)
 Disabled

Capacity pool tiering policy
You can optionally enable automatic tiering of your data to lower-cost capacity pool storage.

2. Storage efficiency

EDA workloads benefit from enabling FSx for ONTAP in-line storage efficiencies.

Compression, deduplication, and compaction leads to better storage utilization, increased

operational efficiency resulting in lower storage cost.

Select -> **Enabled** (Recommended)

3. Capacity pool tiering policy

EDA workloads often contain 60-80% cold data. Enabling Auto tiering policy for non-scratch volumes can help reduce storage costs.

For scratch volumes it is recommended to set **Capacity pool tiering policy** to 'auto'. This enables a cost-efficient configuration. However, capacity pool tiering policy can be set to 'None', if performance is the priority.

Select -> **Capacity pool tiering policy**

then

Select **Auto** from the pulldown box.

Default volume configuration

Volume name

Maximum of 203 alphanumeric characters, plus _ .

Junction path

The location within your file system where your volume will be mounted.

Volume size

Minimum 20 MiB; Maximum 104857600 MiB

Storage efficiency
Select whether you would like to enable ONTAP storage efficiencies on your volume: deduplication, compression, and compaction.
 Enabled (recommended)
 Disabled

Capacity pool tiering policy
You can optionally enable automatic tiering of your data to lower-cost capacity pool storage.

Backup and maintenance

Snapshots are enabled by default on your volumes using the default snapshot policy. It is best practice to review and revise snapshot policies when creating or updating volumes. Scratch volumes may or may not benefit from regular automated snapshots. Scratch volumes often see a high file change rate which can result in snapshots consuming a significant amount of available storage capacity. Snapshots can be enabled or disabled after volume creation.

FSx for ONTAP also offers a daily automatic backup feature. Daily backups leverage Snapshots. Default snapshot policy is enabled even without enabling daily backups. The FlexCache source volume should be backed up and FlexCache destination does not need backup because it caches the data from source volume. Scratch volumes do not need to be backed up. Backup for other volumes (Tools/Libraries, projects, and home dir) is recommended.

▼ **Backup and maintenance - optional**

Daily automatic backup [Info](#)
Amazon FSx can protect your data through daily backups

Enabled
 Disabled

Daily automatic backup window [Info](#)
 No preference
 Select start time for 30-minute daily automatic backup window

Automatic backup retention period [Info](#)
Choose the number of days that Amazon FSx should retain automatic backups for this file system.

days
Minimum 1 days; Maximum 90 days.

Weekly maintenance window [Info](#)
When patching needs to be performed, Amazon FSx performs maintenance on your file system only during this window.

No preference
 Select start time for 30-minute weekly maintenance window

Tags

Creating Tag key/value pairs can provide searchable labels for billing, file system ownership, etc.

There are no EDA specific recommendations for Tags.

Example tags in an EDA environment

Key: Value pair

project_name: projecta

eda_workflow: vcs_nightly_regressions
workflow_owner: projecta_verification_team

▼ **Tags - optional**

Tag key	Value	
<input type="text" value="Enter key"/>	<input type="text" value="Enter value (optional)"/>	<input type="button" value="Delete tag"/>
<input type="button" value="Add another tag"/>		

Cancel

Review and Approve

Review Create file system Summary.
If the settings are correct Select -> **Create file system** button.

Summary		
Verify the following attributes before proceeding		
Attribute	Value	Editable after creation
File system type	Amazon FSx for NetApp ONTAP	⊗
File system name	eda_scratch_filer	⊙
Deployment type	Single-AZ	⊗
Storage type	SSD	⊗
SSD storage capacity	82,944 GiB	⊙
Minimum SSD IOPS	160000 IOPS	⊙
Throughput capacity	4096 MB/s	⊙
Virtual Private Cloud (VPC)	vpc-0a33d2062627d42a0	⊗
VPC Security Groups	sg-00c944a3d5e416882	⊙
Subnet	subnet-06dfde8daf842d337	⊗

Tags

< 1 >

Key	Value
No data	

Storage virtual machines (SVMs)

Attribute	Value	Editable after creation
Name	svm_eda_1	✔
Fully qualified domain name		✔
DNS server IP addresses		✔
Service account username		✔
Delegated file systems administrators group - optional		✔

Cancel Back Create file system

Post FSx for ONTAP Creation Setup

The following steps will be run on each FSx for ONTAP file system after the file system creation process completes.

EDA Workloads significantly benefit from FlexGroup volumes for maximum scale and performance. FSx for ONTAP does not yet support creation of FlexGroup or FlexCache volume types from the AWS Management Console (this may change with later FSx for ONTAP releases)

FlexGroup and FlexCache volumes must be provisioned via the FSx for ONTAP CLI or via ONTAP REST API.

[ONTAP CLI commands](#) are shown below.

Recommended EDA-specific options (pre-volume provision)

NetApp recommends enabling 64-bit identifiers as follows:

```
ONTAP cli> vserver nfs modify -vserver [VSERVER_NAME] -v3-64bit-identifiers enabled
```

Create FlexCache volume for Tool/Library volumes

FlexCache volumes can enable almost instant replication (caching) of on-prem Tools and libraries volumes to AWS. Each on-premises source tool or library volume will require a FlexCache volume. If there are multiple tools and library volumes on-prem, then multiple FlexCache volumes will be needed in the cloud.

Setting up ONTAP cluster and SVM peering relationship between two ONTAP file systems is prerequisite for creating FlexCache. The following NetApp ONTAP documentation provides instructions for setting up peering relationships: <https://docs.netapp.com/us-en/ontap/flexcache/create-volume-task.html>

Once the Cluster and SVM peering relationships have been setup and verified, if required - FlexCache (burst to AWS or Multi-Region on AWS), volumes may be created using the following command:

```
ONTAP cli> volume flexcache create -origin-vserver [ORIGIN_SVM_NAME] -origin-volume [ORIGIN_VOLUME_NAME] \  
    -vserver [DESTINATION_SVM_NAME] -volume [DESTINATION_VOLUME_NAME] \  
    -size [VOLUME_SIZE] -junction-path [JUNCTION_PATH] \  
    -aggr-list aggr -aggr-list-multiplier 8 \  
    -policy [EXPORT_POLICY]
```

Set atime-update on Origin/Source to False. (ONTAP 9.11.1 or earlier)

Set atime-updates on the Origin (or source volume) to False. To avoid invalidations on files that are cached when there is only a read at the origin, turn off last accessed time updates on the origin volume.

If the Origin/Source volume is running ONTAP 9.11.1 or later, the creation of a FlexCache, will automatically set atime-update to False.

From the Origin/Source Volume ONTAP CLI (ONTAP 9.11.1 or earlier)

```
ONTAP cli> volume modify -vserver origin-svm -volume vol_fc_origin -atime-update false
```

In releases earlier than NetApp ONTAP 9.11.1, best practice is to disable atime-update on the

FlexCache origin volumes to reduce the amount of data being updated in the cache because of read access. Failing to disable atime-update on the source volume may lead to overly aggressive cache eviction. For example, simple reads at the source volume will evict the data from the cache.

NetApp ONTAP 9.11.1 and later provides an improved workflow that can allow for atime-updates to be propagated across the environment and not impact evictions at the cache. Even reads at the cache will be able to update atime on the origin. The cache never is an authority on atime, only the origin is.

After upgrading the Origin/Source volume to ONTAP 9.11.1 or later, if there is no use for atime, then leave atime-update off. If there are workflow decisions made based on atime, then use the knobs created in 9.11.1 appropriately for your need.

Create FlexGroup volume for R/W heavy workload

Create one or more FlexGroup volumes for scratch. If required, create for Tools/Libraries and Project as well.

FSx for ONTAP file system

```
ONTAP di> volume create -vserver [SVM_NAME] -volume [VOLUME_NAME] \  
    -size [VOLUME SIZE] -junction-path [JUNCTION_PATH] \  
    -aggr-list aggr1 -aggr-list-multiplier 8 \  
    -policy [EXPORT_POLICY] \  
    -files-set-maximum true [-supporttiering true]
```

More information on FlexGroup volume creation can be found [here](#).

FlexGroup volumes do not support backups today. If backups are needed, customers should consider using SnapVault to another FSx for ONTAP file system.

Large file considerations

The aggr-list-multiplier specifies the number of constituent volumes per node or aggregate pair. Eight-constituent volumes are recommended for EDA workloads. If an 80TB volume is provisioned, each constituent will be 10TB in size. A large file is any file that is 1% to 5% of the member constituent volume. In this example, a file large is any file larger than (80TB/8 x 5%) or

0.5TB (500GB). In general, it is better to provision larger thin-provisioned FlexGroup volumes than smaller volumes.

Increase iNode Count

EDA environments may have millions of files, requiring increase in inode value to the maximum in many cases. It is best practice to constantly monitor the inode utilization using [Amazon CloudWatch metrics](#) (Files Used Volume metrics can be used for the purpose). It is recommended to increase inode count when 80-90% utilization is observed.

Inode defaults and maximums according to FlexVol size.

FlexVol volume size	Default inode count	Maximum inode count
20MB*	566	4,855
1GB*	31,122	249,030
100GB*	3,112,959	24,903,679
1TB	21,251,126	255,013,682
7.8TB	21,251,126	2,040,109,451
100TB	21,251,126	2,040,109,451

*FlexGroup member volumes should not be any smaller than 100GB in size.

Mount Volumes

NFSv3 continues to be the standard for EDA workloads since it is more performant than NFSv4.x. The following sections provide mount commands for READ heavy Tool and Library volumes and for providing the fastest possible performance for READ/WRITE Scratch volumes.

Refer to the Appendix for Technical Reports (TRs) for NFS and FlexGroup Best Practices for further guidance for optimizing NFS mounts for EDA/HPC workloads.

Mount command for Read Optimized Tool/Library volumes

To improve file and metadata loading times as well as decrease the load on the ONTAP storage system, it is recommended to use aggressive compute server caching for Tool and Library volumes which are mostly READ-ONLY in nature. The NFS mount options "nocto, actimeo=600" have been shown to reduce metadata IO of EDA workloads by as much as 90%. It can also dramatically improve the load time of tools and libraries which is becoming an increasingly big challenge as semiconductor technology library sizes are growing dramatically in number and size at 7, 5, and 3nm design nodes (aka transistor size).

EDA workflows often make heavy use of Makefile or Makefile like dependencies tracking. The metadata IO is 60-80% of the overall IOPS. In addition to Makefile like dependency checks, EDA and scripting tools like Perl and Python often scan long files or directory search paths in the modules to load. The result is additional metadata IO transactions used in the "search" for files and directories. This can put an unnecessary load on the file system.

Aggressive NFS client-side caching works based on the assumption that Tools and Library volumes are static or unchanging. New Tools and Libraries are typically installed in directories next to older versions and are seldom over-written or modified. It is considered poor practice for EDA workloads to pull in the "latest" libraries mid analysis. EDA flows typically specify a set of tools and libraries for a given job, then those versions are used throughout the whole run.

Blog Article: [What does actimeo mean during NFS mount in Linux?](#)

The following mount command is recommended to maximize server-side caching using the "nocto and actimeo" options.

Review and apply this best practice per the EDA flow's requirements. Overly aggressive server-side caching may lead to slow cache updating issues.

From the client side:

```
%> mount -t nfs \  
-o "nocto,actimeo=600,hard,rsize=262144,wsiz=262144,vers=3,tcp,mountproto=tcp" \  
    <ontap svm>:/<volname> /mnt/<volname>
```

Note: Refer EDA – [Best Practices in ONTAP](#) for more information on command line values

Mount command for mixed R/W optimized performance on Scratch volumes

The following mount option is optimized for high performance NFS Read/Write transactions. Unlike READ-ONLY Tool and Library volumes, aggressive file caching is typically not recommended.

The following NFS v3 mount command is recommended for EDA scratch volumes workloads.

```
%> mount -t nfs \  
-o "hard,rsize=262144,wsize=262144,vers=3,tcp,mountproto=tcp" \  
<ontap svm>:/<volname> /mnt/<volname>
```

Note: Refer EDA – [Best Practices in ONTAP](#) for more information on command line values


NFS v4.x is not recommended for EDA workloads and may result in as much as 20% performance penalty.

Performance Monitoring

Performance monitoring is critical for EDA workloads. The single most important metric is volume R/W latency. If you see the latency of the volume is not at acceptable level for the EDA application, it means the volumes are under heavy IO workload and it is better to scale-up and distribute the workload across multiple FSx for ONTAP file systems.

High latency conditions can occur to high IO loads from EDA jobs running in parallel. It is recommended that EDA workloads running on FSx for ONTAP be tested at scale to determine the IO limits. Below is a graph a NAS workload which is pushed to its performance limits. The knee of the performance curve is the point at which the controller is operating outside of its performance profile and latencies start to increase rapidly. Operating in this state for long periods of time will impact EDA job runtime performance (wall clock)

Note: The below graph is a generic NAS performance profile graph for illustration purposes only and does not represent actual FSx-NetApp ONTAP performance.



Quick latency monitoring with ONTAP CLI

There are multiple ways to measure real-time ONTAP performance. One way is to simply report latency using the ONTAP CLI. The following `qos statistics volume latency show` command will keep running until Ctrl-C is pressed to stop reporting and will output a report as shown below:

```
cluster1::> qos statistics volume latency show -iterations 100 -rows 3
```

Workload	ID	Latency	Network	Cluster	Data	Disk	Qos Max	Qos Min	NVRAM	Cloud	FlexCache	SM Sync	VA
vol2	999	44.24ms	44.05ms	0ms	190.00us	0ms	0ms	0ms	0ms	0ms	0ms	0ms	0ms
-total-	-	38.89ms	38.63ms	0ms	256.00us	0ms	0ms	0ms	0ms	0ms	0ms	0ms	0ms
vol2	999	64.47ms	64.20ms	0ms	266.00us	0ms	0ms	0ms	0ms	0ms	0ms	0ms	0ms
vol1	1234	27.28ms	27.03ms	0ms	253.00us	0ms	0ms	0ms	0ms	0ms	0ms	0ms	0ms
vsivol0	111	23.72ms	23.47ms	0ms	249.00us	0ms	0ms	0ms	0ms	0ms	0ms	0ms	0ms

More detail on the command can be found [HERE](#).

Start the command either before or during EDA job submission to check the real-time volume latency. Extended periods of latency greater than acceptable latency could be signs of a heavily loaded storage file system.

Monitoring and Reporting

You can use the following services and tools to monitor Amazon FSx for ONTAP usage and activity:

- **Amazon CloudWatch** – You can monitor file systems using Amazon CloudWatch, which automatically collects and processes raw data from FSx for ONTAP into readable metrics. These statistics are retained for a period of 15 months so that you can access historical information and see how your file system is performing. You can also set alarms based on your metrics over a specified time period and perform one or more actions based on the value of the metrics relative to thresholds that you specify. For more details, please refer [FSx for ONTAP monitoring documentation](#).
- **NetApp Cloud Insights** – You can monitor configuration, capacity, and performance metrics for your FSx for ONTAP file systems using the NetApp Cloud Insights service. You can also create alerts based on metric conditions.
- **NetApp Harvest and NetApp Grafana** – You can monitor your FSx for ONTAP file system by using NetApp Harvest and NetApp Grafana. NetApp Harvest monitors ONTAP file systems by collecting performance, capacity, and hardware metrics from FSx for ONTAP file systems. Grafana provides a dashboard where the collected Harvest metrics can be displayed.
- **AWS CloudTrail** – You can use AWS CloudTrail to capture all API calls for Amazon FSx as events. These events provide a record of actions taken by a user, role, or AWS service in Amazon FSx.

Capacity Reporting

Reporting on the size and utilization of a FlexCache volume can be achieved using:

`show -fields <fields to report>` command from the ONTAP CLI. There are RESTful equivalents as well.

Here is an example with a few fields populated. You will see in the below example, the FSx for ONTAP 'proj' volume is a FlexCache of the on-premises 'proj' source volume.

ONTAP: :> volume show -volume <volname> -fields size,used,available,percent-used,files-used,files

In the following example.

on-prem 'proj' is a FlexGroup volume

AWS 'proj' is a FlexCache volume. After reading (or pre-warming) some of the files in the FlexCache, you can see the source volume capacity used and the much smaller FlexCache capacity used.

```
ONTAP CLI on-Prem::> volume show -volume proj -fields size,used,available,percent-used,files-used,files
```

vserver	volume	size	available	used	percent-used	files	files-used
svm_onPrem	proj	1TB	46.47GB	15.85GB	1%	31876696	226372

```
ONTAP CLI FSx-N::> volume show -volume proj -fields size,used,available,percent-used,files-used,files
```

vserver	volume	size	available	used	percent-used	files	files-used
svm_awsFSx-N	proj	150GB	147.0GB	2.96GB	1%	4669368	792

Contributors

Contributors to this document include:

- Arunkumar Krishnan (Solutions Architect - AWS)
- Kartik Gopal (Senior Solutions Architect - AWS)
- Michael Johnson (CTO - NetApp)

Further reading

For additional information, refer to:

[FSx for ONTAP Documentation on-line](#)

[Amazon FSx for NetApp ONTAP performance](#)

[Electronic Design Automation best practices](#)

[NetApp ONTAP FlexGroup volumes best practices and implementation guide](#)

[NFS in NetApp ONTAP Best practice and implementation guide](#)

[Amazon FSx for NetApp ONTAP Monitoring Amazon FSx for ONTAP](#)

Document revisions

Date	Description
Jan 29, 2023	Brief description of revisions.
Feb 16, 2023	First publication