PURESTORAGE®
Uncomplicate Data Storage, Forever

TECHNICAL WHITE PAPER

Accelerating Secondary Genomic Analysis in a Hybrid Cloud

Illumina, Microsoft Azure, and Pure Storage® on Equinix Metal demonstrate how to scale secondary analysis in a hybrid cloud environment with cloud-adjacent storage, while maintaining fast performance and cost-effectiveness.

EQUINIX Microsoft

Contents

Introduction

Life sciences organizations are increasing their use of genomics to develop new therapies and to identify the root causes of numerous diseases. Many research organizations are interested in using hybrid cloud to augment their onsite capabilities for genomics secondary analysis, which is data-intensive.

A hybrid cloud approach provides life sciences organizations with a combination of on-site and cloud-based technology. This approach enables the organization to burst into the cloud as needed to gain more compute power and to access additional secondary analysis capabilities.

By using virtual machines in the cloud and cloud adjacent storage in Equinix Metal with Pure Storage FlashBlade®, an organization can complete secondary analysis on numerous genomic samples in parallel, rather than being limited to onsite platforms.

Deploying a hybrid cloud strategy for secondary analytics includes several challenges. Pure Storage, Illumina, and Microsoft Azure conducted testing in Equinix's Dallas data center utilizing Equinix Metal to determine how these challenges could be overcome.

Data Transmittal Challenges with Hybrid Cloud

Moving genomic data from on premises to the cloud can be time-consuming. In a typical scenario, a life sciences organization has two to five Illumina DRAGEN servers on-premises and uses them to support secondary analysis of genomic data by the Illumina DRAGEN Bio-IT platform.

Each server can complete secondary analysis on one genomics sample at a time. If the organization needs to conduct additional analysis with DRAGEN located in a cloud environment, the IT team would need to manually upload data from on-premises storage via FTP. Our testing suggests this would require approximately 90 minutes per sample (based on a sample size of 32GB each), adding significant delays to the secondary analysis pipeline.

Testing Summary

Illumina, Azure, and Pure Storage conducted testing of the Illumina DRAGEN™ Bio-IT Platform and Pure Storage FlashBlade®, an all-flash data storage platform on Equinix Metal.

We examined how FlashBlade streamlines data transmission between on-premises DRAGEN and DRAGEN in the cloud, saving time for customers by providing a simpler architecture and process. We measured the time required to complete secondary analysis of genomic data when customers rely on hybrid cloud environments for their use of DRAGEN.

Testing Results

The results demonstrated that Pure Storage can support secondary analysis in parallel in the cloud, with FlashBlade located in local and remote data centers with direct cloud-ramp. Pure Storage FlashBlade can scale as rapidly and quickly as the customer's network allows.

In the test scenario, the secondary analysis was able to scale from two samples on-premises to 50 samples running on 50 Illumina-qualified virtual machines in a hybrid cloud environment. We demonstrated that extending to the cloud made it possible to scale the number of samples analyzed in parallel by 25 times. Every virtual machine in the cloud processed a single sample to maintain a good quality of analysis.

Faster Replication to the Cloud

Utilizing FlashBlade improved replication times significantly:

- With the array-level replication of Pure Storage on Equinix Metal with FlashBlade embedded, moving input files to the cloud was 35 times faster than the traditional way of moving genome sample files manually via FTP.
- We transferred the reference and input files for one genome sample (32GB each) with FlashBlade array-level replication to the cloud environment in 2.4 minutes, moving all the files simultaneously. Moving the files via FTP took **90 minutes per sample**.
- While we used 50 virtual machines for the test, a life sciences research organization could use this same approach to scale up to 100 virtual machines while deploying a higher network speed (20-100GB Azure ExpressRoute).

Rapid Secondary Analysis in Parallel in the Cloud

During testing, DRAGEN virtual machines reliant on Pure Storage completed secondary analysis on 50 genomic samples (32GB each) in 60 minutes through parallel processing. (This performance could vary based on genome sample size, however.) Latency on the FlashBlade system remained low, ranging from 2.85 milliseconds with two virtual machines to 3.6 ms with 50 virtual machines, demonstrating scalability.

Benefits

Switching to Pure Storage FlashBlade arrays has many advantages, including:

- **Cost-effectiveness**. Compared to other storage offerings in Azure, Pure Storage is up to 50% more cost-effective with respect to capacity and bandwidth for analyzing multi-genome samples in parallel.
- **Parallel sequencing**. FlashBlade enables users to sequence up to 50 genomic samples simultaneously while maintaining consistent and fast completion times.
- **Efficiency and seamless replication between on-premises and cloud.** FlashBlade eliminates the time previously needed to copy genomic data to local storage on each virtual machine and accelerates time to answer. FlashBlade enables customers to rapidly copy large volumes of genomic data between on-premises and cloud-based DRAGEN servers, because of Snapshot-based file replication.
- **Scalability.** FlashBlade scales sequencing workflows in the cloud with more compute power and low latency.

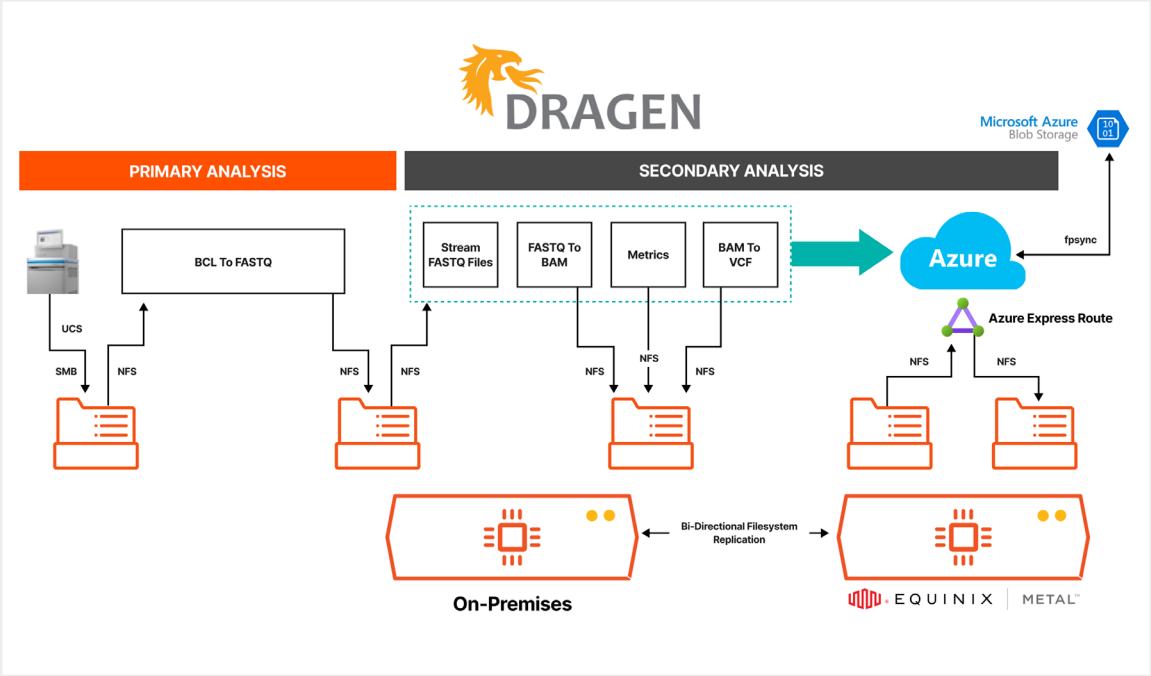

FIGURE 1 Next-generation sequencing with Illumina DRAGEN on FlashBlade on-premises and in Azure Cloud utilizing Equinix Metal.

The Business Case for Cost-effective, High-performing Storage

Resequencing and performing various clinical tests on genomic data are becoming cheaper and faster (Figure 2). This results in in rapid data growth and issues with data management and storage costs.

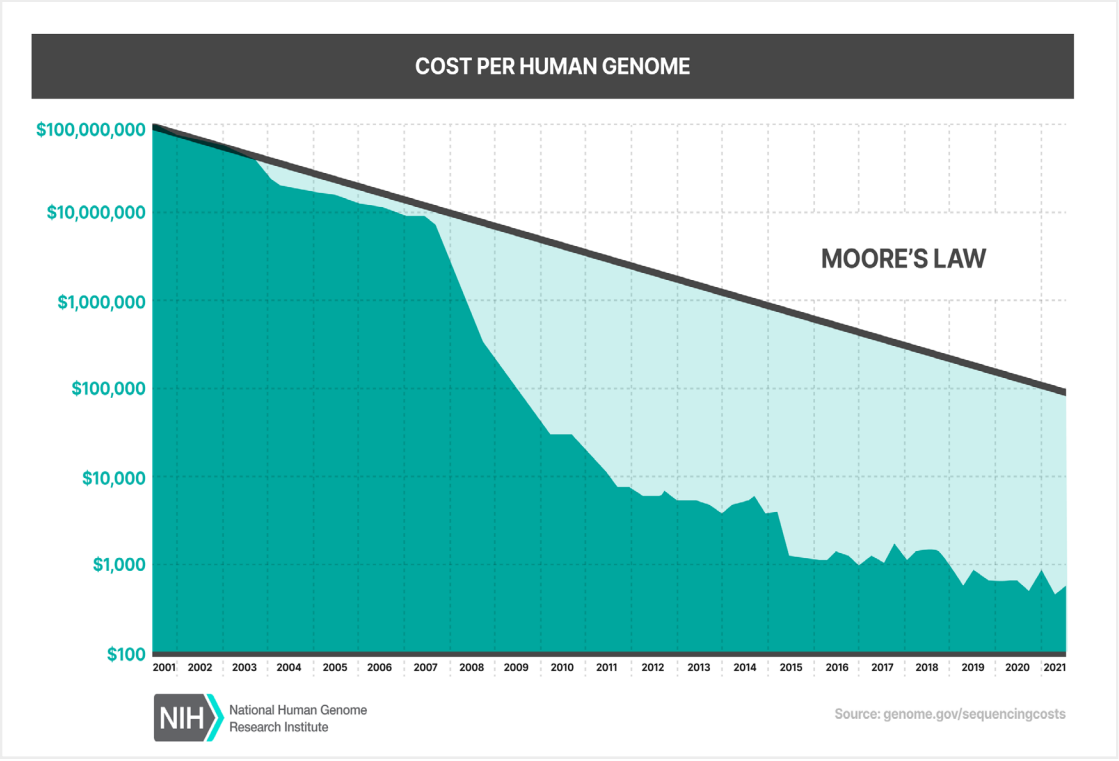

FIGURE 2 Decreasing Costs for Genome Sequencing. Source: https://www.genome.gov/about-genomics/fact-sheets/Sequencing-Human-Genome-cost

Genomics data storage infrastructure needs to handle workloads consisting of millions of small files, large files, and metadata while delivering low-latency IOPS and high throughput. Organizations want to prevent any interruptions or delays to scientific research due to data infrastructure performance.

Life sciences organizations often experience the following storage challenges for their next-generation sequencing (NGS) workloads by:

- Maintaining scalable capacity and performance for high-volume data generated during analysis of multiple genomics datasets
- Ensuring flexibility to manage and protect NGS data while analyzing data on-premises and bursting into public clouds with scalability and limited data friction.

The Data Storage Challenges of Secondary Analysis

Following the production of base-pair sequencing reads in primary analysis, organizations accumulate petabytes of reads that are assembled into whole genomes. Secondary analysis involves assembling the reads and aligning them against a reference genome to identify variants. To complete secondary analysis, organizations typically need to move raw data from interim storage to a high–performance computing (HPC) environment for assembly and analytics. Legacy HPC environments are not always capable of supporting secondary analysis workloads, leading to high-latency and low throughput.

During secondary analysis, the FASTQ, binary alignment map (BAM), and variant call format (VCF) files require considerable storage space. While the VCF files are smaller in size, scientists and researchers often need to reprocess the raw NGS data, such as binary base call (BCL) files, for future analysis.

Use of DRAGEN with FlashBlade Data Storage

FlashBlade is a symmetric, distributed, and scale-out data platform designed to load-balance client requests from a high performance computing (HPC) cluster and process data at scale. FlashBlade provides a secure and standard platform for genome sequencing workflows by eliminating data silos.

The sequencer output data can be stored or copied to FlashBlade over Server Message Block (SMB), providing a centrally managed storage platform that connects to the DRAGEN compute cluster over Network File System (NFSv3) for on-premises secondary analysis (Figure 1).

FlashBlade provides a shared data platform to collect and store the raw NGS data, and provides protection through array-level snapshots. Customers can simply add blades to scale capacity and performance for NGS analysis of thousands of genome samples in HPC environments.

If an organization consistently processes genome sample sizes that are 1-2TB of data, they can be well-supported with local data storage using SSD or faster NVMe-based storage. FlashBlade offers scalability for Next-Generation Sequencing (NGS) data sets greater than 2TB.

Solutions Tested

Here are the solutions that were tested:

- Illumina DRAGEN (dynamic read analysis for genomics) provides accurate, comprehensive, and efficient secondary analysis of NGS data.
- Azure is a cloud computing platform operated by Microsoft. Azure NP series virtual machines (VMs) use the AMD/Xilinx Field Programmable Gate Array (FPGA) acceleration card that enables Illumina DRAGEN processing.
- Pure Storage FlashBlade is an all-flash, consolidated data storage solution for both file and object workloads. FlashBlade is designed to support data-heavy, unstructured workloads efficiently, providing unequaled density, capacity, and performance.
- Equinix is the world's digital infrastructure company, enabling digital leaders to fast-track competitive advantage across clouds, networking, storage, compute, and software.
- Pure Storage on Equinix Metal with FlashBlade embedded provides a high-performance, full-stack, hosted bare-metal infrastructure platform. Enterprises can take advantage of the platform to build best-of-breed, cost-optimized, hybrid-cloud solutions. The platform is fully interconnected with every major public-cloud provider. It features more than 356,000 global interconnects to nearly every telecom and ISP provider, creating a true hybrid/multi cloud environment with unparalleled performance.
- The Equinix Metal platform embeds the full portfolio of Pure Storage products, including FlashArray™, FlashBlade®, and Portworx®. The solution enables native block, file, object, and container storage on dedicated, bare-metal hardware, physically managed by Equinix Metal.

Simplify data storage and infrastructure management into a single fully integrated, on-demand platform. Reimagine the cloud journey and accelerate business and digital transformation. Gain the flexible infrastructure to quickly respond to changing business needs and to help unlock critical business insights.

Testing Methodology: Illumina DRAGEN, Azure, and Pure Storage FlashBlade

For the purpose of this validation, Pure Storage partnered with Microsoft Azure to run the secondary analysis on Azure NP20s VM at scale on FlashBlade (as shown in Figure 3). FlashBlade was available in the Equinix Metal location in Dallas and connected to Azure VMs in San Antonio using a 10Gbps ExpressRoute connection. There was an 8ms round-trip latency between the Azure VMs and FlashBlade and a 3 millisecond round-trip latency for parallel workloads.

NOTE: *Organizations that are equipped with Azure ExpressRoute network configured to the native data center do not need to use FlashBlade in Equinix Metal locations.*

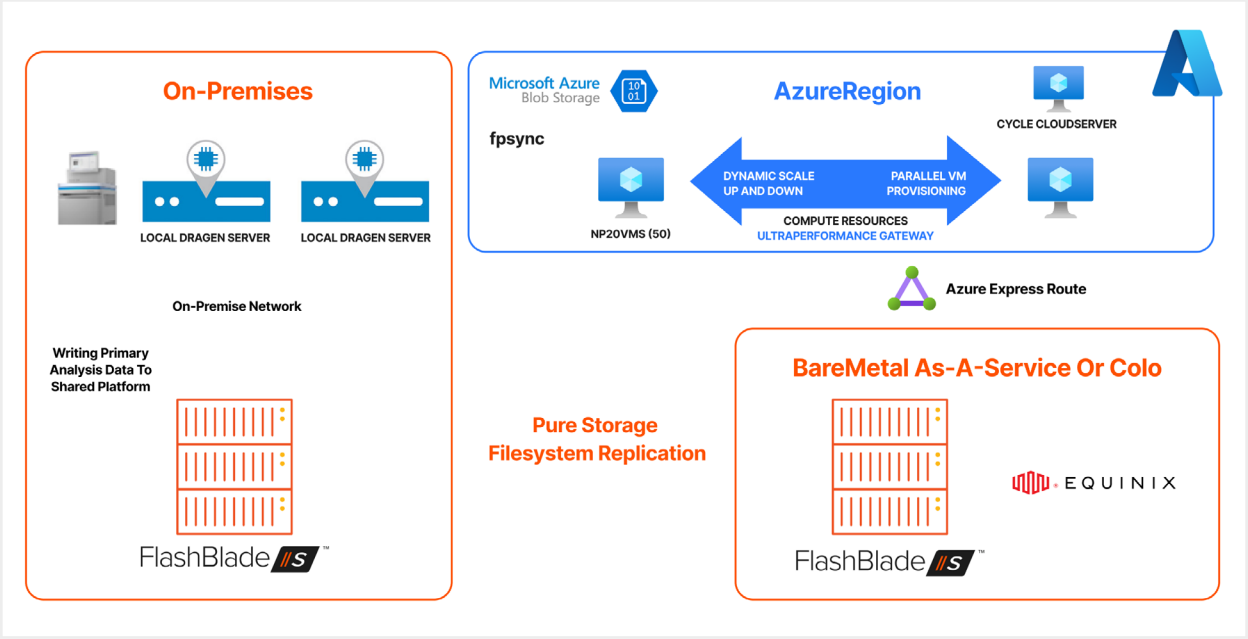

FIGURE 3 Scaling NGS secondary analysis to Azure Cloud and long-term data retention

Testing Environment

After moving the NGS data and input files to FlashBlade with seven blades in the Equinix Metal location, we used the Azure CycleCloud Server to create a Slurm cluster to scale from 2 NP20s VMs up to 50 VMs in a geometric sequence. All the Azure VMs used the Illumina DRAGEN image available in the Azure Marketplace (built on CentOS 7.9) with the following best practices for NP20s VM kernel optimization and NFS mount options for optimum performance. The read-only file system for input files and the read/write file system for output files were mounted over NFSv3 to all the Azure VMs, respectively. We used the DRAGEN v4.0.3 in the test environment.

```
192.168.16.16:/azure-dragen-test  /mnt/azure-dragen-test  nfs  hard,rw,bg,vers=3,nolock,top,timeo=600

# echo "sunrpc.tcp_max_slot_table_entries = 128" >> /etc/sysctl.conf
# sysctl -pctl -p
```

Test Results and Observations

The goal was to measure the secondary analysis phase completion time for multiple samples processed in parallel and to identify any impediments that slow the completion time at scale. The tests simulated running an Azure VM to process a single genome sample. The final test represented one genome sample with the reference data analyzed 50 times from the FlashBlade to demonstrate scalability in Azure cloud.

The chart in Figure 4 shows that the bandwidth (represented in the X-axis) from the FlashBlade scaled linearly up to 16 VMs and then dropped while scaling to 32 and 50 VMs respectively. The reason: The 10Gb/sec ExpressRoute(ER) started to saturate during the 32 VM load thus queuing up the IO request in the network layer. Increasing ER speed to 100Gb/sec can increase the parallel access to multi-datasets on FlashBlade for faster results.

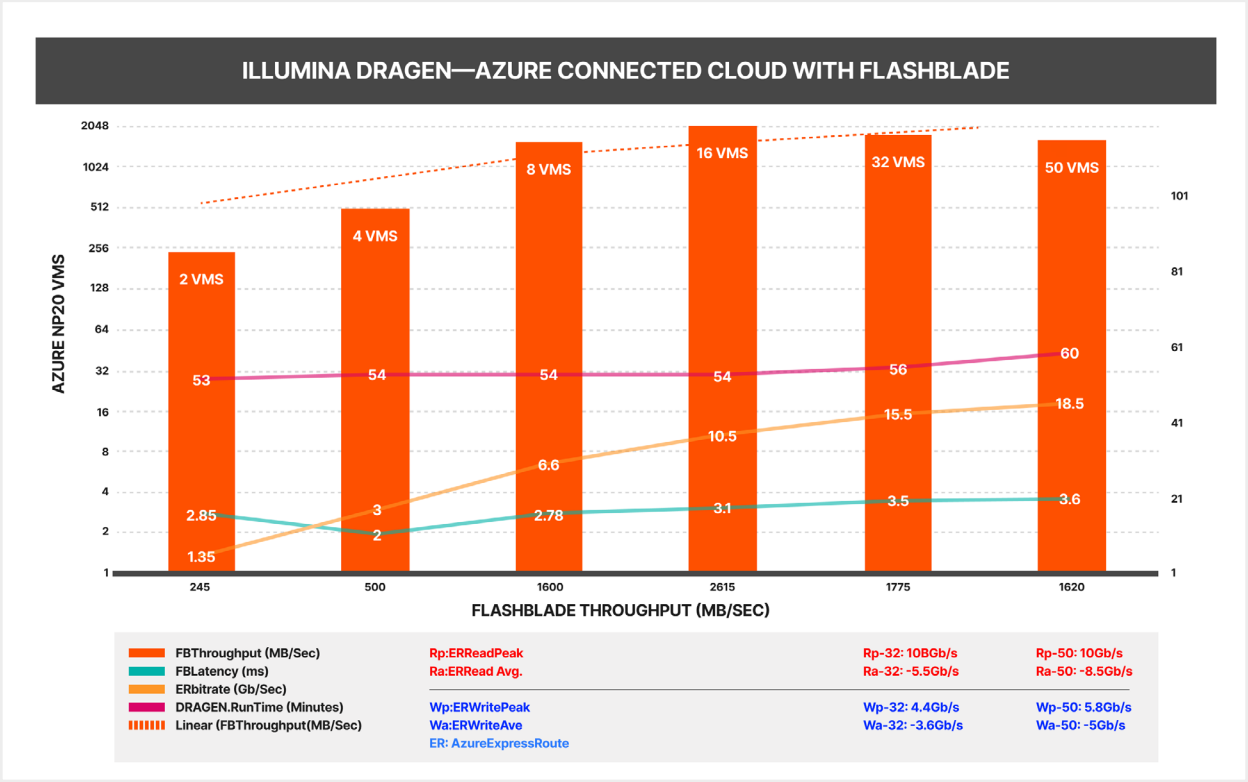

FIGURE 4 Scaling genome data in Azure Connected Cloud to FlashBlade

The brown line (in the middle) in the chart above shows the ER saturation at scale. Adding more network speed to the ER should improve the throughput of the FlashBlade system. Figure 5 shows the read/write peaks and averages reported from ER network speed as the Azure VMs are scaled from 2-50. The final test with 50 VMs was less optimal as the network was completely saturated. The secondary analysis (redline on the top) took ~54 minutes to complete until it started to increase to 60 minutes once the ER network saturation was met.

Recommendation: Address Network Connectivity

Network bandwidth is a potential limiting factor to the speed of analyzing multiple genomic samples in parallel in the cloud. During our testing, we scaled parallel processing to the point of saturating network connectivity on a 10Gbps network connection. Investing in additional bandwidth from the cloud provider could address the saturation issue.

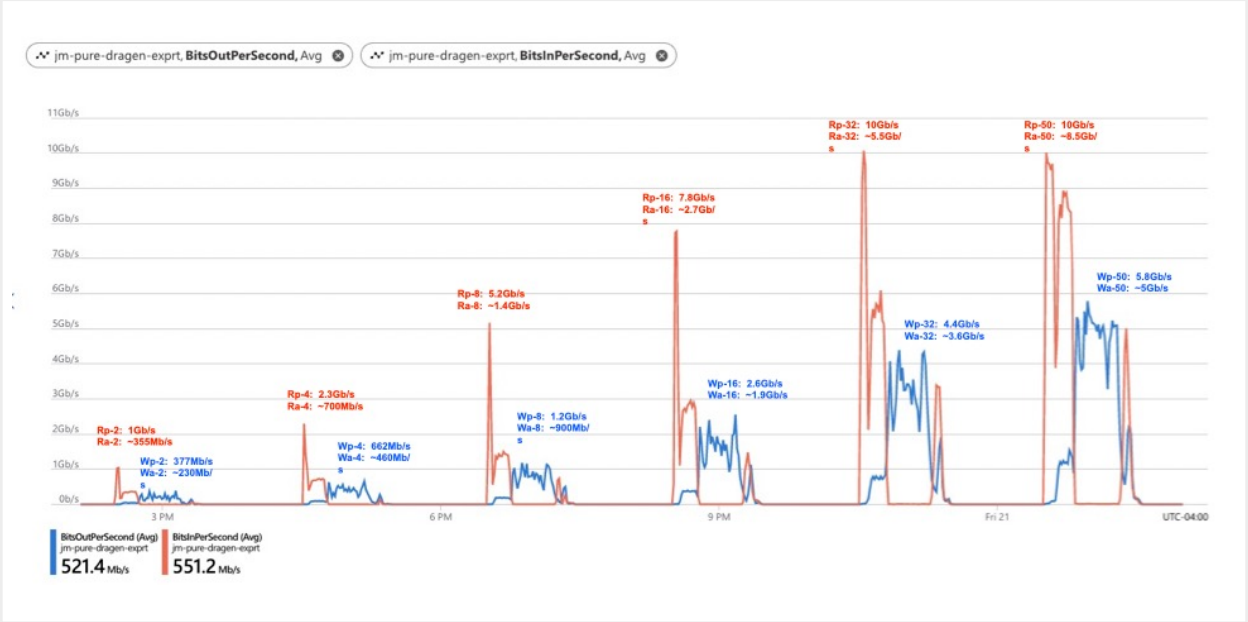

FIGURE 5 **Azure ExpressRoute network traffic while scaling from 2-50 VMs**

The read/write latency (teal line at the bottom) on the FlashBlade in Figure 6 stayed at approximately 3ms, considering the round trip time between the Azure VMs and the FlashBlade was 8ms. The following figure shows a linear scaling of the secondary analysis workload on the FlashBlade as the number of genome samples scaled from 2-50 during the test process. Even though the ER saturation started to slow down the analysis completion, the FlashBlade had plenty of headroom to accommodate more genome samples at scale.

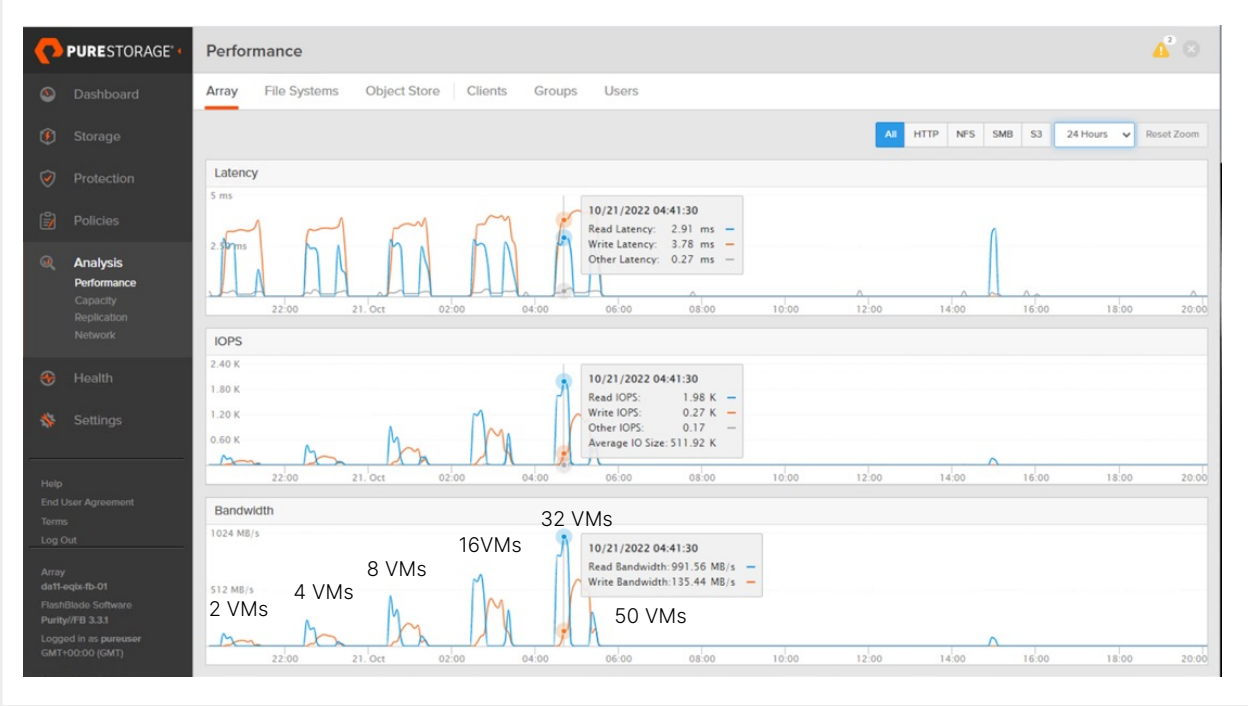

FIGURE 6 **Latency, IOPS, and bandwidth dashboard**

A single sample was copied to the local storage in Azure VM to measure the secondary analysis completion time before starting to scale the Azure VMs in the connected Cloud to FlashBlade. The completion time was recorded to be 53 minutes from the Azure NP20 VM local storage.

The takeaways from the Azure and Pure Storage on Equinix Metal with FlashBlade embedded test results were:

1. There was no or very little impact to the secondary analysis completion time compared to the Azure NP20s VM local storage. The consistent turnaround time (TAT) for scaling high volume of genome sample sequencing in Azure Cloud on FlashBlade provides business owners faster results to research data.
2. Sequencing multiple genome samples in parallel over NFS from FlashBlade will produce a consistent secondary analysis completion time until the ER gets saturated.
3. Using Pure Storage on Equinix Metal with FlashBlade embedded eliminates the overhead to copy the samples to local storage on each of the Azure VMs, thereby slowing down the overall sequencing process.

Testing Illumina DRAGEN with an On-premises FlashBlade

To test an on-premises scenario, Illumina tested a FlashBlade and two DRAGEN servers in their on-premises data center. They compared analysis time for a genome sample supported by FlashBlade storage to the analysis time of a genome sample with local storage on the DRAGEN servers. Both the DRAGEN servers were running on DRAGEN v4.0.3 with CentOS 7. Each DRAGEN server consisted of Intel(R) Xeon(R) Gold 6226 CPU @ 2.70GHz with Dual CPU/12 cores each and 256GB memory. The FlashBlade had 14 blades in a single chassis for this test. The input files for the secondary analysis were in Original Read Archive (ORA) compression format for the sample genome data.

There were two test scenarios:

1. Copy the genome sample to local storage on the DRAGEN server and measure the completion time from each of the DRAGEN servers.
2. Copy the genome samples to FlashBlade and measure the completion time on one server and scale to two DRAGEN servers over NFSv3.

Test Results: Scalability and Time-saving Operational Benefits

The FASTQ.ORA input files were copied to the local storage on the on-premises DRAGEN server. The time to copy the FASTQ.ORA input files of the test sample to the DRAGEN server local storage took 84 minutes. Adding or scaling to another DRAGEN server to accelerate the sequencing process took only incremental time to copy the sample input files to each of the DRAGEN server local storage.

In Figure 7, the chart on the left shows the base and incremental time to copy the FASTQ.ORA input files of the test sample across one and more DRAGEN servers. The hypothetical third DRAGEN server (not used in the testing) in the chart illustrates the additional time it may take to copy the input files to the local storage on the server when the DRAGEN servers are scaled incrementally in the sequencing process.

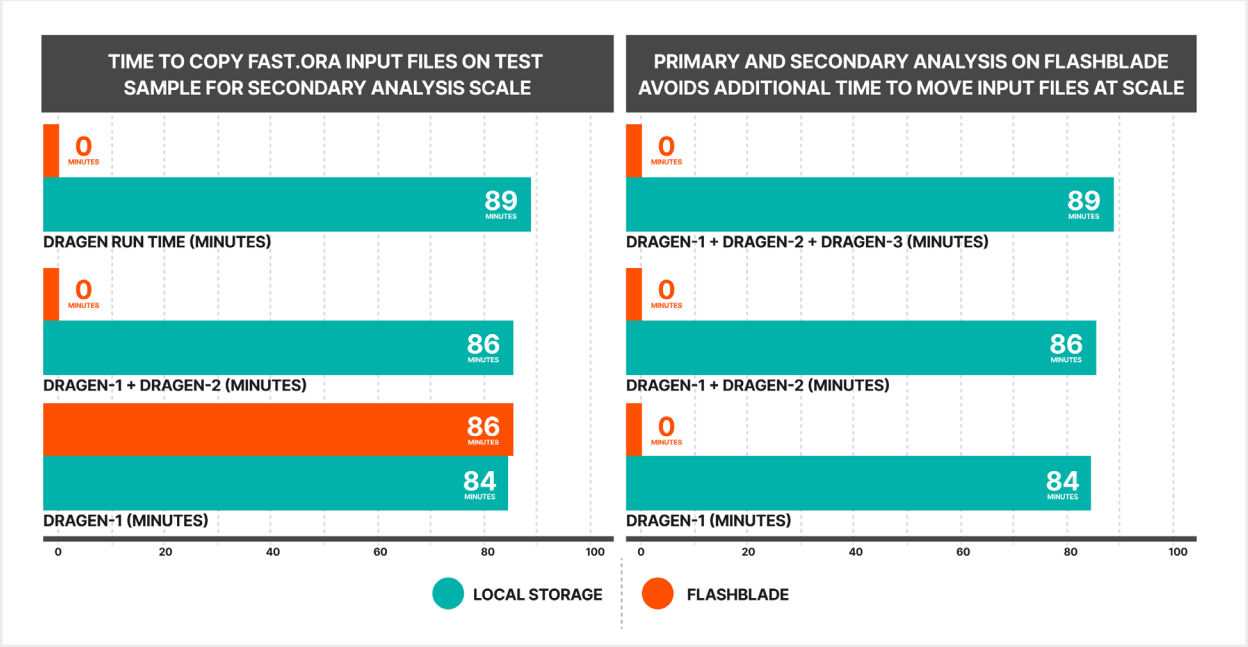

FIGURE 7 Accelerating genome sequencing with FlashBlade at scale

However, when the FASTQ.ORA inputs were copied to the FlashBlade, adding DRAGEN servers and pointing to the common dataset became fast and seamless to scale. The initial copy of the input files of the test sample to the FlashBlade took 86 minutes to complete over NFSv3. The second and subsequent DRAGEN servers accessed the input files from FlashBlade, thus reducing the time spent on copying files to each DRAGEN server at scale.

Input test sample files on a shared data platform like FlashBlade accelerated the sequencing process to achieve faster test results. It is recommended to copy the FASTQ input files to FlashBlade for scalable capacity and performance along with improved data management and protection.

Moreover, the chart on the right in Figure 7 shows that copying the input files for every genome sample to the FlashBlade can be eliminated further when the primary analysis and secondary analysis phases are configured on FlashBlade as a standard data platform.

FlashBlade provides the data continuity between the sequencing phases as the datasets of various genome samples can be reused, protected, managed, and stored for long-term retention at very low operational costs.

Performance Benefits

The secondary analysis for the sample genome data completed in 38 mins from a single DRAGEN local storage and it took 41 mins to complete from the FlashBlade. However, when the second DRAGEN server was added to the test, there was hardly any difference in the completion time for the secondary analysis between the data accessed from local DRAGEN storage compared to FlashBlade, notwithstanding the NFSv3 protocol overhead.

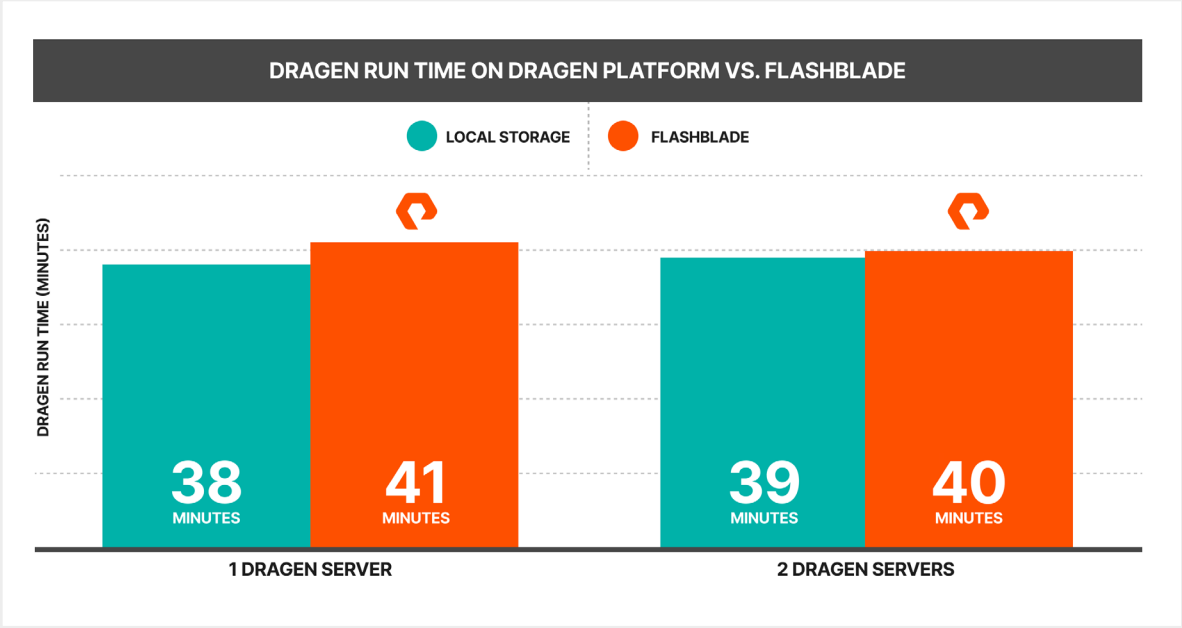

FIGURE 8 Secondary analysis completion: DRAGEN local storage vs. FlashBlade

The test result clearly shows that the FlashBlade can scale linearly when a greater number of genome samples are sequenced from many DRAGEN servers, providing consistent completion times during the secondary analysis phase.

Effective Data Storage Architecture for NGS

A genomic data storage solution needs to meet capacity requirements for high-volume NGS data and address performance requirements for sequencing thousands of genome samples in parallel. In addition, researchers have to reanalyze the NGS data for additional clinical tests for different variants. The VCF and other output files are often stored in the production data platform to reduce cloud storage cost. While VCF files are smaller and easier to work with, this approach restricts the ability to analyze new variants.

FlashBlade reduces the operational costs and improves cost-efficiency for the raw NGS data and test results during the different phases of the genome sequencing process. Raw BCL files, the FASTQ input files from the VCF, and other test results are partitioned on the FlashBlade during different analysis phases, yielding better data management and data protection.

Partitioning the raw BCL files, the FASTQ input files from the VCF and other test results in separate file systems on the FlashBlade array during different analysis phases can yield better data management and data protection.

- Automation using simple APIs allows sequencing workflows to provision and consume storage on demand from FlashBlade.
- Multi-protocol support on FlashBlade allows data collection devices and sequencing workflows to store data over NFS and SMB.
- FlashBlade data reduction allows customers to store more analysis data and test results, while achieving a smaller storage footprint for thousands of genome samples.

With FlashBlade in use, raw NGS data can be re-sequenced more quickly. There is no need to delete the raw and intermediate files. FlashBlade is capable of handling a high number of parallel jobs from thousands of genome samples analyzed in various phases.

Benefits of Array-level Snapshots

Instant array-level snapshots and rapid file-system restore allow the researchers to protect and revert back to the original dataset, in the event of any errors while reprocessing in the analysis phase.

- The raw NGS data and the FASTQ files can be moved independently to a lower tier of storage for long term retention after the analysis of the genome samples are complete.
- The NGS data is encrypted by default on FlashBlade for privacy and security reasons. FlashBlade SafeMode™ Snapshots protect genomic data and test results in the event of a ransomware attack, making it possible to rapidly restore all data.

By using bi-directional, array-level replication, organizations can burst into the cloud and scale beyond data center boundaries. FlashBlade provides the proximity to access the data from public cloud providers like Azure along with data sovereignty and security for compliance reasons. Sequencing workflows can scale in the cloud, while benefiting from more compute power and data locality on FlashBlade.

Clinical data can remain on-premises for security reasons. Sequencing data and input files can be replicated to the target FlashBlade in the Equinix Metal location, using array-level file system replication to stage the input files before starting the secondary analysis in Azure cloud.

The array-level file-system transfer offers the following advantages for the NGS data lifecycle:

The replication target in the Equinix Metal location is a read-only copy. The input ORA files are available in the target filesystem that can be read by the compute VMs in Azure over NFSv3. The transient data and the VCF outfile files are written to another filesystem on FlashBlade during the secondary analysis phase.

1. The read-only filesystem in the Equinix Metal Location can be removed to free up the storage space and reduce the overall operation costs. Pure Storage provides an OpEx model, metered to pay only the storage space consumed, including data reduction achieved from any NGS data.
    a. If re-analysis is needed, the input files can be quickly replicated from the on-premises source into the cloud environment. This repeat process has no bearing on the output files that are stored in a different filesystem on FlashBlade.
    b. Multiple target filesystems to store the FASTQ/ORA input files with separate filesystem to store the output files for thousands of genome samples can be accomplished on the FlashBlade in the Azure connected cloud environment.

2. Customers experience the following benefits when VCF output files are stored in another read-write file system on FlashBlade. Output files can be:
   a. Used as input to the tertiary analysis phase on FlashBlade
   b. Replicated back to the on-premises location after the secondary analysis process is complete
   c. Archived in Azure Blob Storage using the "fpsync" tool for long-term data retention as shown in Figure 3

Conclusion

The cost of sequencing genome samples at a large scale can be a tradeoff from the performance and speed required to sequence the NGS data. FlashBlade includes simple APIs and cloud-like storage, providing the following benefits for sequencing multiple datasets of genome samples:

1. Speed to sequence and resequence many genome samples in parallel on-premises and to burst into the cloud at scale. The turnaround time is consistent while scaling the sequencing of high volumes of genome samples for researchers.
2. Automation to provision, manage and protect the data on-premises or in Azure Connected Cloud environments to reduce the overcall cost of the data storage.
3. Competitive and predictable total cost of ownership with a subscription-based pricing model.
4. Azure Connected Cloud to FlashBlade via Equinix Metal fabric requires no egress fees for data movement from Azure, providing additional cost savings.

Additional Resources

- Illumina DRAGEN
- Pure Storage on Equinix Metal
- Accelerate Genomics Sequencing for Precision Medicine
- McArthur Lab Fights Global Threats to Human Health
- Read more about Azure High Performance Computing
- Explore more Azure HPC technical content

About the Authors

Bikash Roy Choudhury, a director for solutions at Pure Storage, is responsible for designing and architecting solutions for DevOps workflows relevant across industry verticals including high tech, financial services, gaming, social media, and web-based organizations. He has also worked on validating solutions with Rancher/Kubernetes, GitLab, Jenkins, JFrog Artifactory, IBM Cloud Private, and Perforce using RESTful APIs and integrating them with data platforms in private, hybrid, and public clouds. In his current role, Bikash drives integrations with strategic DevOps partners, including Rancher, Mesosphere, Perforce, GitLab, and JFrog.

PS2442-en-01 05/23