

Dell EMC PowerScale Powered by Yellowbrick and Faction to Modernize Data Warehouse

Modern data warehouse solution for lightning-fast queries on petabytes of data while supporting thousands of concurrent users.

Abstract

This paper describes the solution and implementation process of setting up modernized data warehouse for hybrid cloud analytics solution powered by Yellowbrick, Faction cloud and Dell EMC PowerScale to support lightning-fast queries on petabytes of data while supporting thousands of concurrent users. Here we also show the industry standard TPC-DS performance results of this solution along with data load speed from on-premise data lake into Yellowbrick data warehouse.

November 2020

Revisions

Date	Description
December 2020	Initial release

Acknowledgments

Author: Kirankumar Bhusanurmath, Analytics Solutions Architect, Dell EMC.
Joey Foley, Sr. Solution Engineer Yellowbrick

The information in this publication is provided "as is." Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

Copyright ©2020 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners. [12/17/2020] [Solution Guide] [S2407]

Table of contents

Revisions.....	2
Acknowledgements.....	2
Table of contents	3
Executive summary.....	4
1 Yellowbrick	5
1.1 About Yellowbrick Data Warehouse.....	5
1.2 Yellowbrick Hybrid Cloud Data Warehouse.....	5
1.3 Yellowbrick Architecture	5
1.4 Yellowbrick Data Lake Acceleration and Integration.....	5
2 Dell EMC PowerScale.....	7
2.1 All-Flash nodes	7
2.2 Isilon hybrid nodes	7
2.3 Isilon archive nodes.....	8
2.4 Hardware tiers and node generations	8
3 Solution overview: Modern data warehouse for hybrid cloud analytics	9
3.1 Faction Cloud Control Volumes	9
4 Modern data warehouse solution validation.....	12
4.1 Storage components	12
4.2 Compute services.....	12
4.3 Benchmark suite	12
4.4 Performance results	13
4.4.1 Test 1: Data loading with ybload	13
4.4.2 Test 2: Data Unloading with ybunload.....	14
4.4.3 Test 3: Data Loading with ybload, multiple NFS mount points	14
4.4.4 Test 4: Running the TPC-DS 99 queries.....	15
4.4.5 Testing Yellowbrick benchmarks with cloud only databases.....	17
4.5 Validation summary.....	19
5 Conclusion.....	20
A Technical support and resources	21
A.1 Related resources	21

Executive summary

The modern data warehouse for hybrid cloud analytics with Yellowbrick and Dell Technologies multicloud storage with PowerScale supports all today's requirements, including superior price/performance regardless of data scale, provides a path to the future with flexible deployment options and expands-as-you-grow architecture. This ideal solution supports several key requirements:

- **Price/performance.** One big problem with traditional data warehouses is their high price. Legacy data warehouse vendors have struggled to refresh their platforms in a way that produces good price/performance as data volumes grow and concurrent users increase in numbers.
- **Scalability.** The one thing that is constant with data warehouse deployments is that the volume of data will continue grow, as will the number of users and types of queries. Therefore, when evaluating a more modern platform, it is critical to understand how easy it is to add more data or support more users.
- **Real-time data support.** The ability to ingest and query real-time data (for example, using Kafka) is now a critical requirement. The platform should support complicated analysis on streaming data supporting millions of messages per second.
- **Practical support for cloud migration without lock-in.** A modern platform should support a flexible range of deployment options, so that organizations can select the lowest-risk way of migrating to the cloud (for example, to respect security and data gravity concerns). Some organizations will want to deploy workloads on premises or move to the cloud in a gradual way. A modern solution should run identically across all environments—from on-premises deployments, to a single public cloud deployment, to a hybrid cloud deployment. It should not require an all-or-nothing move to the cloud, and as important, it should not lock users into a specific cloud platform or impose significant financial or time investments to move data off the platform in the future.
- **Streamlined, consistent management.** Legacy data warehouses typically involve lots of specialized tuning, indexing, workload management, and overall management. A modern platform should be easy to manage, with as few operational tasks as possible to ensure good performance. And whatever management is required should be consistent across all deployment platforms, from on premises to cloud.
- **Predictable OPEX pricing.** While most enterprises now avoid CAPEX as a general policy, their need for accurate forecasting is incompatible with the hidden and complex costs typical of cloud-only alternatives. A consumption model that solves for both needs is important.

This innovative hybrid cloud analytics reference architecture with Yellowbrick and Dell Technologies multicloud storage with PowerScale radically expands data bandwidth to support lightning-fast queries on petabytes of data for thousands of concurrent users, as a result customers can:

- Rapidly import data at massive rates, in bulk or as a stream, with data immediately query-able and actionable—giving enterprises instant access to data for timely decisions
- Enable lightning-fast, subsecond ANSI SQL queries across petabyte-size datasets with latencies up to orders of magnitude faster than alternatives—increasing the richness (for example, spanning multiple months of historical data) and rate of insights
- Support parallel queries by hundreds or even thousands of users in familiar BI and data science tools including Tableau, SAS, Micro-Strategy, R, and Python, preserving investments in existing tools
- Eliminate mundane tasks that consume valuable administrative time, such as tuning, creating indexes, and reclaiming storage space—streamlining and simplifying data management
- Ensure flexibility over support for hybrid and multicloud by letting organizations run mixed workloads wherever it makes the most sense: In on-premises data centers, private clouds and multicloud.

1 Yellowbrick

1.1 About Yellowbrick Data Warehouse

Performance breakthroughs in intensive data processing have always come from systems innovation, not just software. Yellowbrick Data has invested years of R&D to deliver a truly innovative optimized instance that is battle-hardened to bring industry-leading price/performance at tens of terabytes to multiple petabyte scale—far beyond what legacy platforms or virtual machines running on generic public cloud instances can ever achieve.

Yellowbrick instances can be deployed in data centers or private clouds, as a service endpoint in AWS, Azure and GCP, and all the above, with data asynchronously replicated for disaster recovery between instances over unified hybrid architecture.

Disclaimer: © Yellowbrick Data, Inc. All rights reserved. Yellowbrick Data and the Yellowbrick logo are trademarks of Yellowbrick Data. All other trademarks are the property of their respective owners. Information is subject to change without notice.

1.2 Yellowbrick Hybrid Cloud Data Warehouse

Yellowbrick built the only modern hybrid cloud data warehouse from the ground up to maximize performance and scale for modern SQL workloads. The database was designed to solve today's business challenges like ingesting and querying real-time data with historical data and support business critical applications with interactive workloads. Yellowbrick allows businesses to scale from a few terabytes to petabytes of data while providing in-memory performance for all data.

Yellowbrick Data Warehouse is an ANSI-SQL compliant database and modernizes processes used by business analyst allowing questions to be asked of the data that could not be done before. Yellowbrick complements data lakes and Hadoop environments for businesses by allowing fast data movement between the two environments. This seamless integration allows Yellowbrick to provide the SQL speed layer for end users to meet demanding subsecond SLAs and lightning-fast dashboard and ad hoc queries.

1.3 Yellowbrick Architecture

Yellowbrick Data Warehouse is a scale-out MPP architecture based on optimized instances. Clusters can scale from four instances to dozens of instances working together. Yellowbrick Data Warehouse is the first data analytics platform designed to take advantage of flash memory, NVMe architecture, and modern hardware accelerators for in-memory speed, with optimizations in the storage, kernel/driver, data path, and database layers, and is the only database technology with “flash-native queries.” The Yellowbrick database itself is a hybrid row- and column-oriented implementation, allowing it to provide the scan throughput of petabyte scale column stores with real time ingest of data. With this approach, limits on query latency, processing capacity, and concurrency associated with data-starved multicore CPUs in traditional architectures are removed, thereby allowing nearly complete utilization of CPU resources. All components of a cluster are redundant for high availability and all data is protected using erasure coding.

1.4 Yellowbrick Data Lake Acceleration and Integration

Yellowbrick Data Warehouse is designed with optimized instances delivering a purpose-built database, hardware, and software for maximum performance. In order to achieve this performance, we believe that

data must be tightly coupled to the processing. Separating compute and storage across current networking topology causes delays in loading to local storage (and adds delays) before data can be processed.

What about data that resides on external storage or a data lake which needs to be processed internally by Yellowbrick?

The Yellowbrick Data Warehouse was designed to ingest data at the speed of today’s modern data volumes. Whether feeds are streamed in from Kafka, bulk loaded from external file systems, like Dell EMC PowerScale, or semi or unstructured data in object storage, like Dell EMC ECS, S3 or Azure Blob, that needs to be enriched, Yellowbrick was built to handle all modern data motion scenarios and works out of the box with common data motion tools like Informatica, Talend, Denodo, Striim, and Oracle GoldenGate.

When we look at methods for data ingest and data export, the Yellowbrick utilities ybload and ybunload are the recommended approaches. Both utilities can parallelize across the backend nodes and can move data at wire-speeds of the network. Customers can achieve rates as high as 10 TB/hr. These utilities parallelize across the external server’s CPU cores they run on, so ensure you have a fast disk subsystem and plenty of cores available. Our standard server for this task has 40 CPU cores and a RAID 1 configuration of multiple NVMe drives with a minimum of 10 Gb network or higher.

Unstructured and semistructured data in a data lake can be incorporated with the Yellowbrick Spark connector. Spark is like a Swiss Army Knife when it comes to processing unique data structures in the data lake. Spark can be used to traverse JSON or XML, and unpack data stored in Parquet, ORC, or Avro. Once Spark has read the data and flattened the data into columns and rows, Spark can stream the data to YBRelay. YBRelay is the Yellowbrick Spark loader and scheduler which streams data to ybload instances. ybrelay can handle multiple simultaneous data streams and can queue and prioritize data streams to the Yellowbrick Data Warehouse.

Yellowbrick Data Warehouse utilities ybload and ybunload can read and write data to many different platforms. Storage systems, like Dell EMC PowerScale, can interact with the utilities over NFS mountpoints. These utilities can also interact with S3 and S3 compatible platforms object stores.

Data Lake Augmentation with Yellowbrick

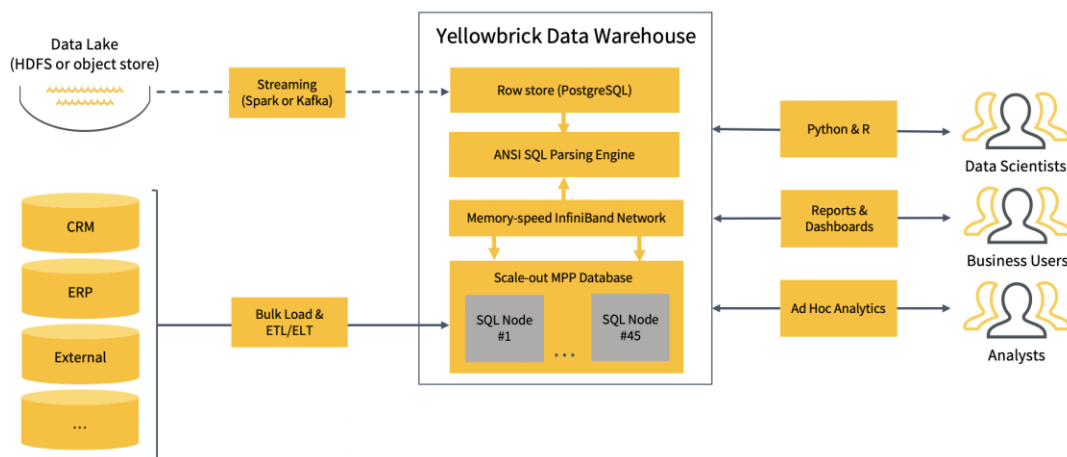


Figure 1 Yellowbrick architecture

2 Dell EMC PowerScale

PowerScale is the next evolution of OneFS – the operating system powering the industry’s leading scale-out NAS platform. The PowerScale family includes Dell EMC PowerScale platforms and the Dell EMC Isilon platforms configured with the PowerScale OneFS operating system. OneFS provides the intelligence behind the highly scalable, high-performance modular storage solution that can grow with your business. A OneFS powered cluster is composed of a flexible choice of storage platforms including all-flash, hybrid, and archive nodes. These solutions provide the efficiency, flexibility, scalability, security, and protection for you to store massive amounts of unstructured data within a cluster. The new PowerScale all-flash platforms co-exist seamlessly in the same cluster with your existing Isilon nodes to drive your traditional and modern applications

2.1 All-Flash nodes

New PowerScale all-flash storage platforms - powered by the OneFS operating system - provide a powerful yet simple scale-out storage architecture to speed up access to massive amounts of unstructured data while dramatically reducing cost and complexity. They deliver extreme performance and efficiency for your most demanding unstructured data applications and workloads. Powered by the new OneFS 9.0 operating system, the all-flash platforms are available in four product lines:

- **PowerScale F200:** Provides the performance of flash storage in a cost-effective form factor to address the needs of a wide variety of workloads. Each node allows you to scale raw storage capacity from 3.84 TB to 15.36 TB per node and up to 3.8 PB of raw capacity per cluster. The F200 includes inline compression and deduplication. The minimum number of PowerScale nodes per cluster is three while the maximum cluster size is 252 nodes. T
- **PowerScale F600:** With new NVMe drives, the F600 provides larger capacity with massive performance in a cost-effective compact form factor to power the most demanding workloads. Each node allows you to scale raw storage capacity from 15.36 TB to 61.4 TB per node and up to 15.48 PB of raw storage per cluster. The F600 includes inline software data compression and deduplication. The minimum number of nodes per cluster is three while the maximum cluster size is 252 nodes.
- **Isilon F800:** Provides massive performance and capacity. It delivers up to 250,000 IOPS and up to 15 GB/s aggregate throughput in a single chassis configuration and up to 15.75M IOPS and up to 945 GB/s of aggregate throughput in a 252 node cluster. Each chassis houses 60 SSDs with a capacity choice of 1.6 TB, 3.2 TB, 3.84 TB, 7.68 TB, or 15.36 TB per drive. This allows you to scale raw storage capacity¹ from 96 TB to 924 TB in a single 4U chassis and up to 58 PB raw storage in a single cluster.
- **Isilon F810:** Provides massive performance and capacity along with inline data compression and deduplication capabilities to deliver extreme efficiency. The F810 delivers up to 250,000 IOPS and up to 15 GB/sec aggregate throughput in a single chassis configuration and up to 15.75M IOPS and up to 945 GB/s of aggregate throughput in a 252 node cluster. Each F810 chassis houses 60 SSDs with a capacity choice of 3.84 TB, 7.68 TB, or 15.36 TB per drive. This allows you to scale raw storage capacity from 230 TB to 924 TB in a 4U chassis and up to 58 PB of raw storage in a single cluster.

2.2 Isilon hybrid nodes

Isilon hybrid storage platforms powered by the OneFS operating system uses a versatile yet simple scale-out architecture to speed access to massive amounts of data. The hybrid platforms are highly flexible and presses the balance between large capacity and high-performance storage to provide support for a broad range of enterprise file workloads. The hybrid storage platforms are available in four product lines:

- **Isilon H400:** Provides a balance of performance, capacity, and value to support a wide range of file workloads. The H400 delivers up to three GB/s bandwidth per chassis and provides capacity options ranging from 120 TB to 960 TB per chassis.
- **Isilon H500:** This versatile hybrid platform delivers up to five GB/s bandwidth per chassis with a capacity ranging from 120 TB to 960 TB per chassis. The H500 is an ideal choice for organizations looking to consolidate and support a broad range of file workloads on a single platform.
- **Isilon H5600:** Combines massive scalability – 1.28 PB (raw) per chassis – and up to eight GB/s bandwidth – in an efficient, highly dense, deep 4U chassis. The H5600 also includes inline compression and deduplication capabilities. The H5600 is designed to support a wide range of demanding, large-scale file applications and workloads.
- **Isilon H600:** Designed to provide high performance at value, delivers up to 120,000 IOPS and up to 12 GB/s bandwidth per chassis. The H600 is ideal for HPC workloads that do not require the extreme performance of all-flash.

2.3 Isilon archive nodes

Isilon offers two highly efficient and massively scalable archive storage solutions. Both nodes use a modular architecture while dramatically reducing cost and complexity and both platforms use a dense hardware design that provides four nodes within a single 4U chassis

- **Isilon A200:** Is an ideal active archive storage solution that combines near-primary accessibility, value, and ease of use. The A200 provides between 120 TB to 720 TB per chassis and scales to 45 PB in a single cluster.
- **Isilon A2000:** A2000 is an ideal solution for high density, deep archive storage that safeguards data efficiently for long-term retention. The A2000 stores up to 960 TB per chassis and scales to over 60 PB in a single cluster.

2.4 Hardware tiers and node generations

Dell EMC PowerScale OneFS clusters can be architected with a wide variety of node styles and capacities, in order to meet the needs of a varied dataset and wide spectrum of workloads. These node styles encompass several hardware generations and fall loosely into four main categories or tiers. The following table illustrates these tiers, and the associated hardware generations and models:

Tier	I/O Profile	Drive Media	Node Type
Extreme Performance	High Perf, Low Latency	All-flash	F800 F810 F600 F200
Performance	Transactional I/O	SAS & SSD	H600 H5600
Hybrid / Utility	Concurrency & Streaming Throughput	SATA/SAS & SSD	H500 H400
Archive	Nearline & Deep Archive	SATA	A200 A2000

Figure 2 Hardware tiers and node generations

3 Solution overview: Modern data warehouse for hybrid cloud analytics

Dell Technologies Cloud Storage enables connecting file storage, consumed as a service, directly to the Yellowbrick hybrid cloud Data warehouse. This is achieved through native replication from on-premises Dell EMC Isilon storage to a managed service provider location. Dell Technologies has partnered with Faction Inc. to deliver a fully managed, cloud-based service for Dell EMC storage to address various cloud use cases.

Faction, Inc. is a Dell Technologies Gold Cloud Service Provider (CSP) and Tech Connect Select partner founded in 2006 and headquartered in Denver, Colorado. Faction is a multicloud platform-as-a-service provider and VMware partner that offers multicloud-attached storage from various co-locations (Equinix, Coresite, and Digital Reality). Faction has expanded globally to London and Frankfurt. In this hybrid cloud data warehouse solution, we use Factions Cloud Control Volumes (CCVs) storage offerings as storage layer or data lake for Yellowbrick data warehouse.

Below figure shows the solution diagram of modern data warehouse for hybrid cloud analytics.

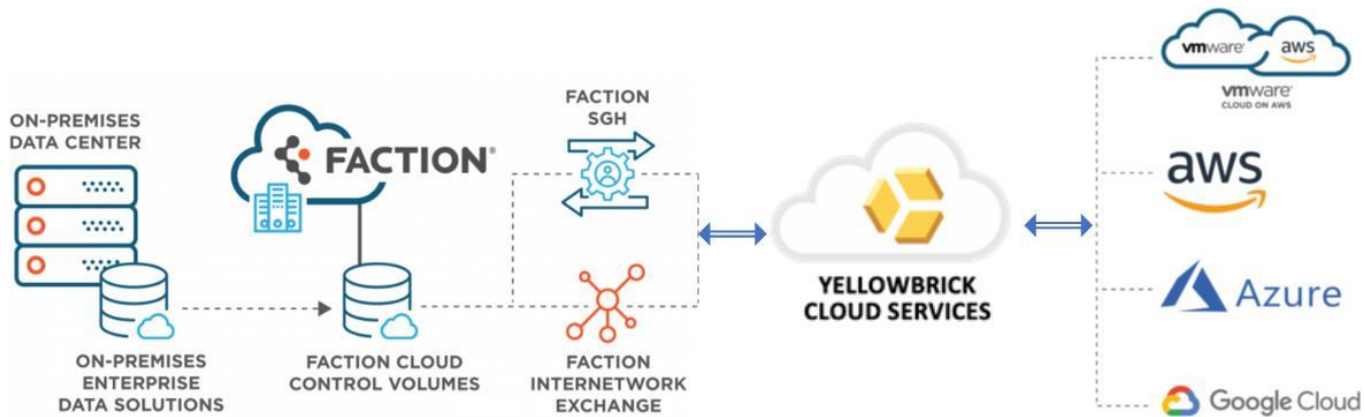


Figure 3 Solution diagram

3.1 Faction Cloud Control Volumes

Cloud Control Volumes (CCVs) provide durable, persistent, cloud-attached, and cloud-adjacent storage directly connected to the Yellowbrick hybrid cloud.

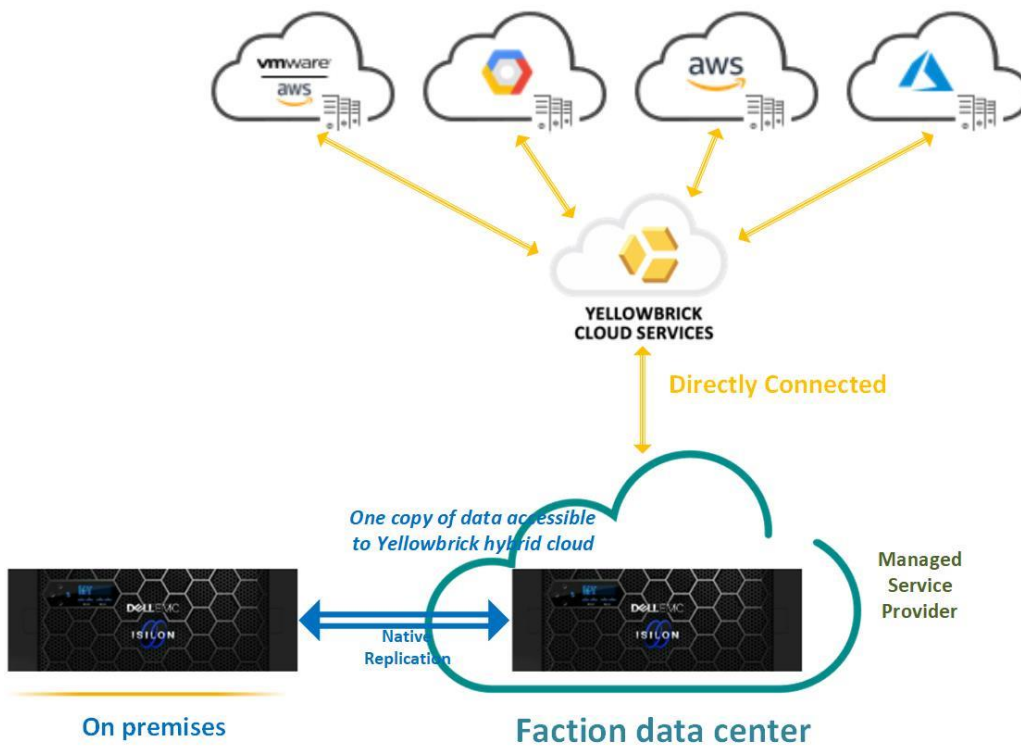


Figure 4 Array-based replication of volumes to Faction directly attached as CCVs across one or more clouds through NFS

Use cases for CCVs could be transient in nature, such as performing data analytics on a large or complex data footprint. A variety of tiers of CCV storages are available in Faction data center. Storage tier specifics are ultimately determined by the Dell EMC arrays and use cases as shown in the below figure.

		Archive	Standard	Premier	Elite
ISILON FILE SCALE OUT	Base Network Connectivity				
		10 Gb/s	10 Gb/s	40 Gb/s	80 Gb/s
	Storage Scaling				
		Base includes 540 TB	Base includes 163 TB	Base includes 162 TB	Base includes 130 TB
		Scale in 300 TB increments	Scale in 90 TB increments	Scale in 90 TB increments	Scale in 77 TB increments
	Workloads				
		Best for workloads and data that requires infrequent access. These include use cases like long-term records retention, write-once-read-never, video retention, and web content management	Delivers a powerful solution for workloads that are like the Archive Tier, but with a need for more active access of data.	Best for applications that require high amount of read access, including video streaming, rendering, test/dev, big data use cases such as genomics, and replacing on-premise file servers	Backed by all flash storage for high performance workloads like critical stream analytics, real-time inference with machine learning, and time sensitive data warehouse.
	Isilon A2000	Isilon A200	Isilon H500	Isilon F800	

Figure 5 File scale-out CCV details

Modern data warehouse for hybrid cloud analytics, organizations need to migrate volume data from an on-premises data center to a Faction data center. Array-based replication is configured between on-premises Isilon storage and a similar Isilon storage array owned and managed by Faction in the Faction data center.

It is the customer's responsibility to manage the network between their on-premises data center and the Faction data center. A dedicated circuit should be opted for a dedicated connection for replication traffic between their facility and Faction. Customers may also use a VPN as redundancy to a dedicated link. Faction can source and manage the dedicated link, or the client can work with their carrier directly.

CCVs are presented in close proximity to Yellowbrick hybrid cloud provider while leveraging redundant connectivity with multiple 10 Gb Ethernet connections and redundant switches to provide highly available connections. Link Aggregation Groups (LAGs) are used to scale to higher levels of bandwidth into the Yellowbrick cloud.

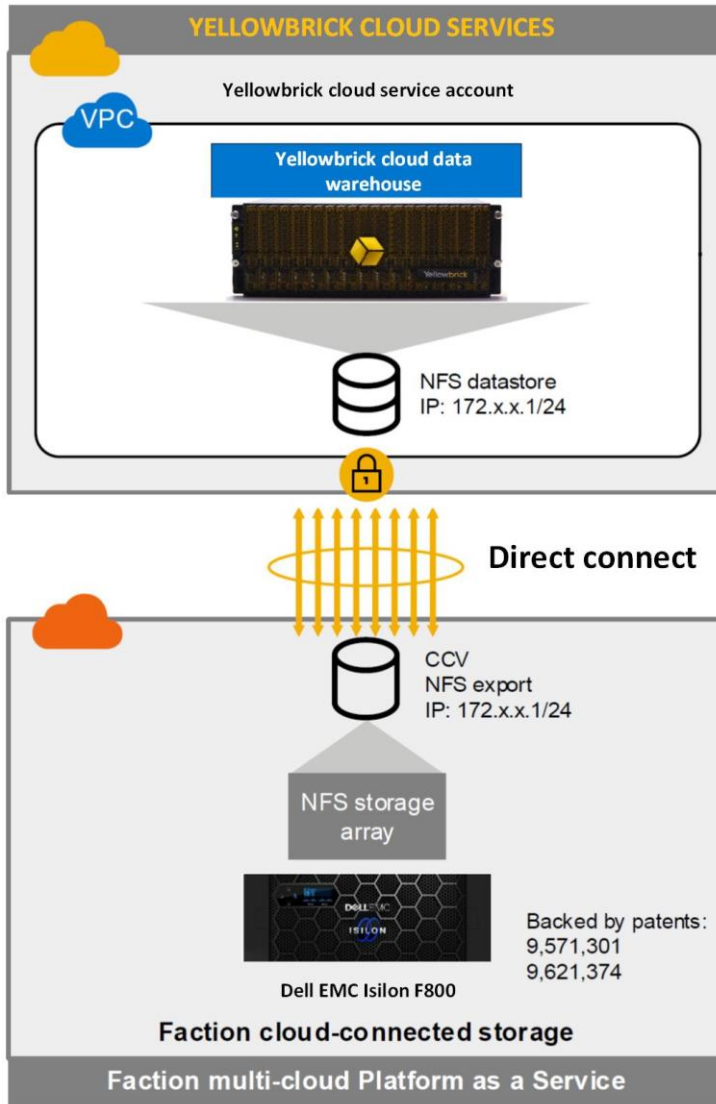


Figure 6 Link Aggregation Groups yield higher levels of bandwidth into Yellowbrick clouds.

The Faction network infrastructure provides full layer-2 and layer-3 isolation between the storage service and the Yellowbrick hybrid cloud infrastructures.

Latency between the Faction data center and the cloud providers varies by public cloud and region. Faction provides the current latency numbers at the page [Faction Latency Information](#). The latency numbers can be used for planning and comparison purposes.

4 Modern data warehouse solution validation

4.1 Storage components

The solution is validated using Dell EMC Isilon F800 4 node cluster, for more information about the Isilon F800 see section 2.1.

4.2 Compute services

Yellowbrick’s Cloud Data Warehouse was used for the testing. The cluster was configured with 15-node instances and two manager nodes connecting the cluster to the network with 40 Gb fiber connections. A Linux-based server was configured with 36 CPU cores and 16 Gb memory to support the ybtools for load and unload testing and running the TPC-DS queries. Multiple clusters were configured for the TPC-DS query testing portion, including a 4-node, 15-node, and 30-node all with the same two manager nodes and networking.

The following diagram illustrates Yellowbrick architecture and connection to the external data source.

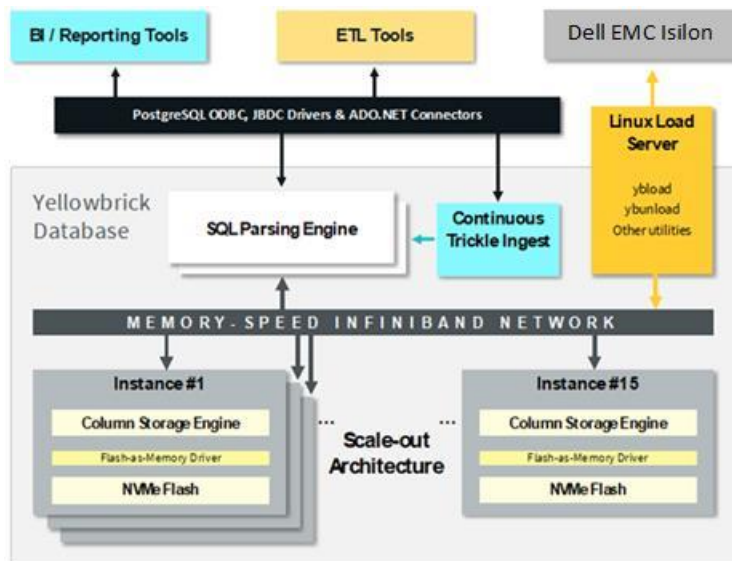


Figure 7 Yellowbrick data warehouse architecture

4.3 Benchmark suite

The industry-standard TPC-DS test suite was used to benchmark Yellowbrick performance. The TPC-DS benchmark data was modeled on the decision support functions of a retail product supplier. TPC-DS consists of seven fact tables and 17 dimensions.

TPC-DS consists of 99 queries which are divided into four broad classes:

- Reporting queries
- Ad hoc queries
- Iterative OLAP queries
- Data Mining queries

The dataset size of 10 TB in text file format is used to test the scale-out factor. A 10TB text file format dataset is generated by using TPC-DS dsdgen utility. Data was placed on PowerScale F800 storage and connected as NFS mount to a Linux-based server that hosted the Yellowbrick ybtools utilities ybload, ybunload, and ybsql. The queries in the TPC-DS test suite were executed from the Linux server directly against Yellowbrick Data Warehouse using ybsql, a command-line utility used to submit SQL. Yellowbrick is ANSI-SQL compliant and can successfully run all 99 queries in the TPC-DS test suite.

Testing between Dell storage and Yellowbrick consisted of using ybload to load the 10TB dataset using a single 10Gb interface. Next, the 10TB dataset was unloaded using ybunload to high-speed test data movement. Further, multiple NFS shares were configured to the Linux load server to push higher throughput from the PowerScale F800. Finally, the 99 query benchmark was run against three Yellowbrick clusters: a 4-node, a 15-node, and a 30-node. To show the full scalability of the 30-node system, the TPC-DS 100 TB dataset was used.

4.4 Performance results

4.4.1 Test 1: Data loading with ybload

Using the ybload utility to load the catalog_sales fact which is 3.2 TB in text files took 1:06:53.

```
2020-10-18 17:07:24.687 [ INFO] <main> ABOUT CLIENT:
  app.cli_args          = "-h" "<HOSTNAME>" "-d" "dell_demo" "--read-sources-
concurrently" "ALWAYS" "--logfile" "<FILENAME>" "-t" "sf10000.catalog_sales"
  app.name_and_version = "ybload version 5.0.0-20201010013014"
  java.version         = "1.8.0_262"
  jvm.memory           = "981.5MB (max=14.2GB)"
  jvm.name_and_version = "OpenJDK 64-Bit Server VM (build 25.262-b10, mixed
mode)"
  jvm.options          = "-Xmx16g, -Xms1g, -Dapp.name=ybload, -Dapp.pid=13683,
-Dapp.repo=/opt/ybtools/lib, -Dapp.home=/opt/ybtools, -Dbasedir=/opt/ybtools"
  jvm.vendor           = "Oracle Corporation"
  os.name_and_version  = "Linux 3.10.0-1127.19.1.el7.x86_64 (amd64)"

2020-10-18 17:07:24.689 [ INFO] <main> Logfile written to
sf10000.catalog_sales.log
2020-10-18 17:07:24.797 [ INFO] <LFSourceContext> Gathering metadata on input
files
2020-10-18 17:07:24.851 [ WARN] <LFConcurrencyContext> Unable to read all 73
sources concurrently. Limiting the number of concurrent sources to 50
2020-10-18 17:07:24.852 [ WARN] <LFConcurrencyContext> Limiting the number of
concurrent sources to 36 in order to match --num-readers
2020-10-18 17:07:24.948 [ INFO] <Choosing apx source encoding> Assuming source
encoding matches database server encoding: LATIN9
2020-10-18 18:14:18.025 [ INFO] <main> SUCCESSFUL BULK LOAD: Loaded 14399964710
good rows in 1:06:53 (READ: 767.0MB/s WRITE: 689.2MB/s)
```

4.4.2 Test 2: Data Unloading with ybunload

Using the ybunload utility to unload the catalog_sales fact which is 3.2 TB in text files took 54:38 minutes.

```
2020-10-18 16:11:18.810 [ INFO] <main> ABOUT CLIENT:
  app.cli_args      = "-h" "10.252.252.40" "-o"
"/mnt/storage/faction2/home/sf10000/catalog_sales" "--truncate-existing" "--
prefix" "catalog_sales" "--logfile"
"/mnt/storage/faction2/home/sf10000/catalog_sales.log" "--logfile-log-level"
"INFO" "-t" "sf10000.catalog_sales"
  app.name_and_version = "ybunload version 5.0.0-20201010013014"
  java.home           = "/usr/lib/jvm/java-1.8.0-openjdk-1.8.0.262.b10-
0.el7_8.x86_64/jre"
  java.version        = "1.8.0_262"
  jvm.memory          = "512.0MB (max=6.0GB) "
  jvm.name_and_version = "OpenJDK 64-Bit Server VM (build 25.262-b10, mixed
mode) "
  jvm.options         = "-Xms512m, -Xmx6g, -XX:+UseG1GC, -Dapp.name=ybunload,
-Dapp.pid=9955, -Dapp.repo=/opt/ybtools/lib, -Dapp.home=/opt/ybtools, -
Dbasedir=/opt/ybtools"
  jvm.vendor          = "Oracle Corporation"
  os.name_and_version = "Linux 3.10.0-1127.19.1.el7.x86_64 (amd64) "
2020-10-18 16:11:18.812 [ INFO] <main> Logfile written to
/mnt/storage/faction2/home/sf10000/catalog_sales.log
2020-10-18 17:05:56.676 [ INFO] <main> Transfer complete
2020-10-18 17:05:56.676 [ INFO] <main> Transferred: 3.23 TB Avg Network BW:
812.76 MB/s Avg Disk write rate: 987.60 MB/s
```

4.4.3 Test 3: Data Loading with ybload, multiple NFS mount points

Using the ybload utility to load across four NFS mount points, the test consisted of loading catalog_returns (223.1 GB), catalog_sales (2.94 TB), web_returns (103.2 GB) and web_sales (1.46 TB) in parallel with each traversing a single mount point. All loads were run across the Linux-based load server with 36 CPU cores and 16 GB of memory to a 15-nodes cluster with two manager nodes with 40 Gb networking over fiber.

The table below show combined read and write rates for the tables with parallel execution, note smaller tables finished faster and thus more sources would be given to larger write tasks over longer run times.

Table 1 ybload time for all tables

Table Name	Read Rate	Write Rate
Catalog_returns	452.9 MB/s	428.9 MB/s
Catalog_sales	767.0 MB/s	689.2 MB/s
Web_returns	380.1 MB/s	345.6 MB/s
Web_sales	736.5 MB/s	665.2 MB/s
Combined Throughput	2336.5 MB/s	2128.9 MB/s

Catalog_returns

2020-10-19 17:15:49.020 [INFO] <main> READ:223.1GB(452.9MB/s). ROWS G/B: 1440033112/0(2.72M/s). WRITE:211.3GB(428.9 MB/s). TIME E/R: 0:08:24/ 0:00:00

2020-10-19 17:15:49.021 [INFO] <main> SUCCESSFUL BULK LOAD: Loaded 1440033112 good rows in 0:08:24 (READ: 452.9MB/s WRITE: 428.9MB/s)

Catalog_sales

2020-10-19 17:14:18.024 [INFO] <main> READ: 2.94TB(767.0MB/s). ROWS G/B: 1439996471 0/0(3.42M/s). WRITE: 2.64TB(689.2MB/s). TIME E/R: 1:06:53/ 0:00:00

2020-10-19 17:14:18.025 [INFO] <main> SUCCESSFUL BULK LOAD: Loaded 14399964710 good rows in 1:06:53 (READ: 767.0MB/s WRITE: 689.2MB/s)

Web_returns

2020-10-19 17:12:02.717 [INFO] <main> READ:103.2GB(380.1MB/s). ROWS G/B: 720020485/0(2.47M/s). WRITE:93.87GB(345.6MB/s). TIME E/R: 0:04:38/ 0:00:00

2020-10-19 17:12:02.718 [INFO] <main> SUCCESSFUL BULK LOAD: Loaded 720020485 good rows in 0:04:38 (READ: 380.1MB/s WRITE: 345.6MB/s)

Web_sales

2020-10-19 17:12:10.027 [INFO] <main> READ: 1.46TB(736.5MB/s). ROWS G/B: 7199963324/0(3.29M/s). WRITE: 1.32TB(665.2MB/s). TIME E/R: 0:34:45/ 0:00:00

2020-10-19 17:12:10.029 [INFO] <main> SUCCESSFUL BULK LOAD: Loaded 7199963324 good rows in 0:34:45 (READ: 736.5MB/s WRITE: 665.2MB/s)

4.4.4 Test 4: Running the TPC-DS 99 queries

Yellowbrick Data Warehouse can be configured with varying numbers of optimized instances to scale to the data volume and performance a business needs to meet demanding SLAs. In trying to show the capabilities of different cluster sizes, the 10 TB TPC-DS dataset was loaded on a 4-node and 15-node cluster. The 4-node cluster ran all 99 tests in 5152.55 seconds. The 15-node cluster ran all 99 tests in 1860.66 seconds and by tripling the cluster size we see the performance is nearly 3x.

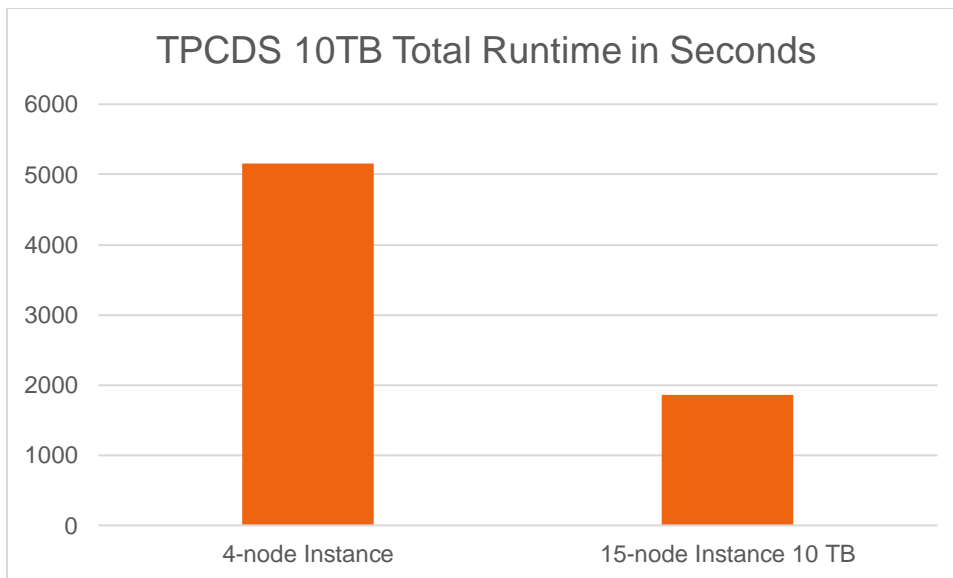


Figure 8 TPC-DS 10 TB 99 query results demonstrate both the performance and linear scalability of additional instances.

The MPP architecture and scale-out design of Yellowbrick Data Warehouse allows for increases in instances for increases in performance and increases in the amount of data the system can process. By adding more instances, you add more compute resources and scale the speed at which you access the data and the interconnections of the nodes in the cluster. This effect combined with the more efficient data path in Yellowbrick has a multiplying effect on data access speed and scan rates.

To showcase the effect of faster data access, a third test was built to showcase the speed of data path access and scaling the node count of a cluster. The next test compares a 10TB TPC-DS dataset on a 15-node cluster to a 100TB TPC-DS dataset on a 30-node cluster. To scale the data by a factor of 10x and the nodes by a factor of 2x, we maximize the throughput of data access in Yellowbrick Data Warehouse. With 10x the data and 2x nodes, one would expect with linear scaling of nodes that runtime should be 5x longer on the 30-node cluster. Yellowbrick Data Warehouse completed all 99 tests in 2336.15 seconds, a speedup of 4x when compared to the expected linear scaled runtime. Yellowbrick Data Warehouse has a unique architecture allows companies to process data volumes at speeds not possible before.

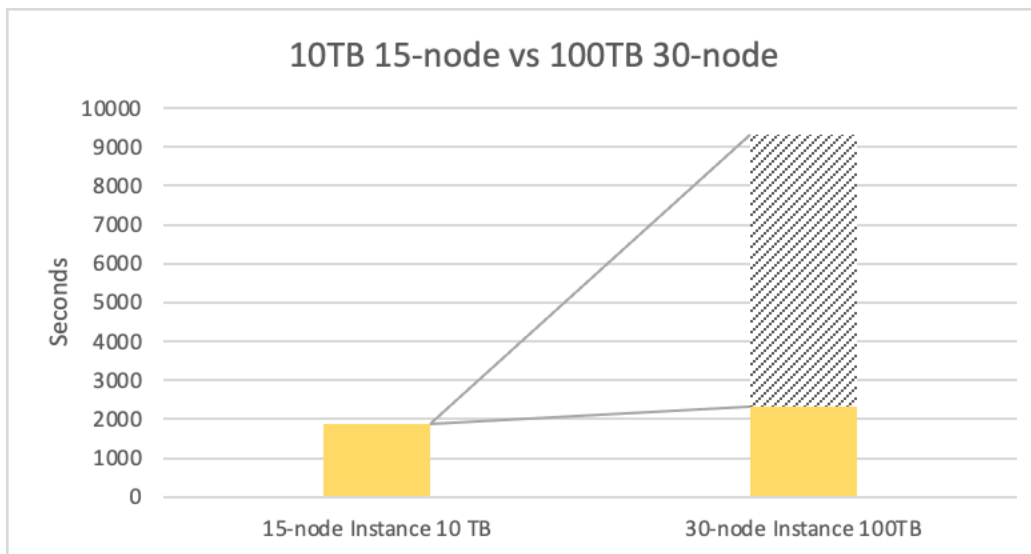


Figure 9 TPC-DS 10 Tb vs 100 TB 99 queries execution time, yellow is actual runtime, gray is expected runtime based on linear scaling of nodes.

4.4.5 Testing Yellowbrick benchmarks with cloud only databases

For the benchmark tests, Yellowbrick decided to compare Yellowbrick Data Warehouse compare with the top cloud vendor's databases using the TPC-DS 30 TB dataset and 99 queries. Yellowbrick used a 16-node instance cluster against similar-size/cost configuration for a leading Cloud Vendor EDW, Config 1, (CPU, Memory, Price). This approach provided a direct performance/cost comparison with various Cloud EDW configurations. The benchmark test showed that Cloud EDW performed these specific complex-query workloads 6X slower than Yellowbrick.

Yellowbrick had planned to scale the Yellowbrick cluster to see how other configurations compared, however; the 16-node instance cluster performance was unmatched, no matter the configuration. In the second comparison, Config 2, Cloud EDW's configuration was doubled but ran 3X slower. In the third, Config 3, Cloud EDW was scaled up to their largest configuration (4X the size of Yellowbrick's Memory and CPU) but still ran 2X slower.

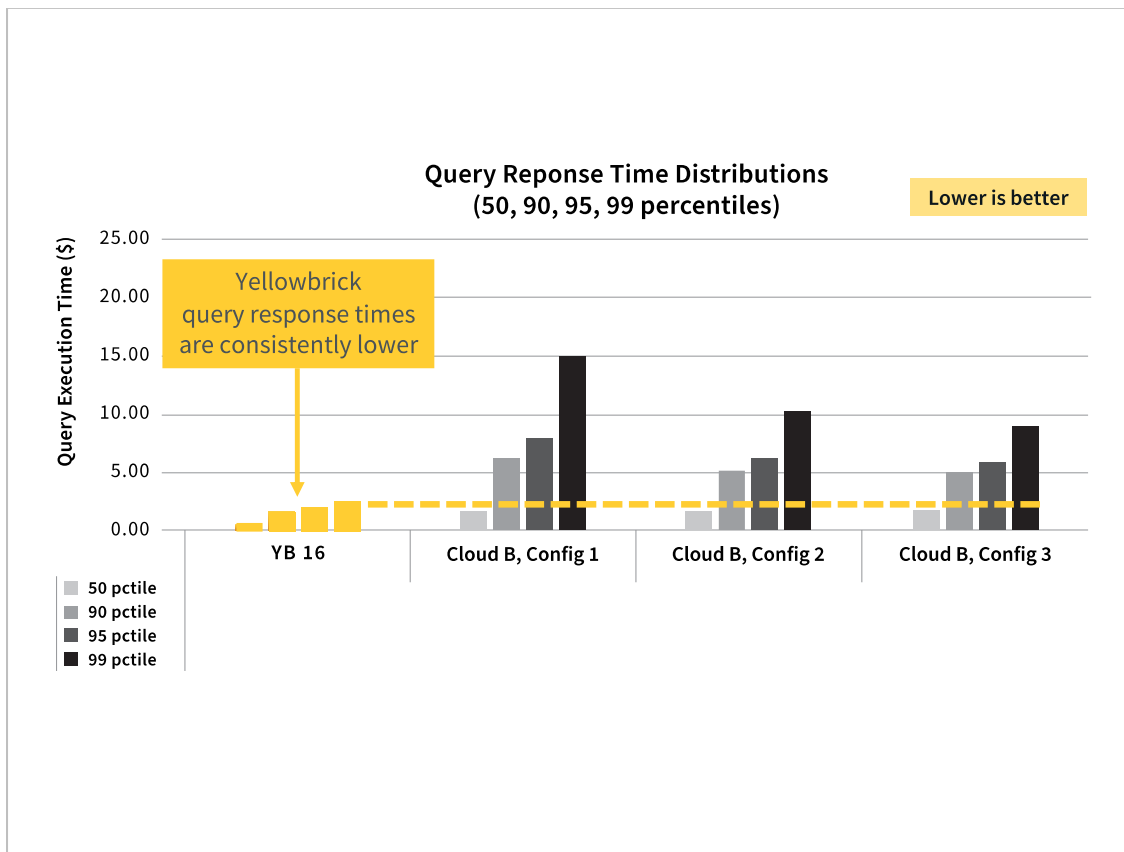


Figure 10 Yellowbrick 16-node instance compared to leading Cloud Vendor, each cloud cluster doubles the number of compute nodes.

Yellowbrick finally compared concurrency against the leading cloud vendor. The leading cloud vendor touts the ability to add clusters on demand for peak concurrency loads. This approach was not used for this test as we wanted to see the efficiency of the single cluster. The concurrency test selected tests of 5, 10, 20, 50 and 100 users running a subset of queries from the TPC-DS 99 queries. Each user ran the same set of queries in a random order and caching mechanisms were turned off for the cloud vendor so they would need to run the queries and not return results only.

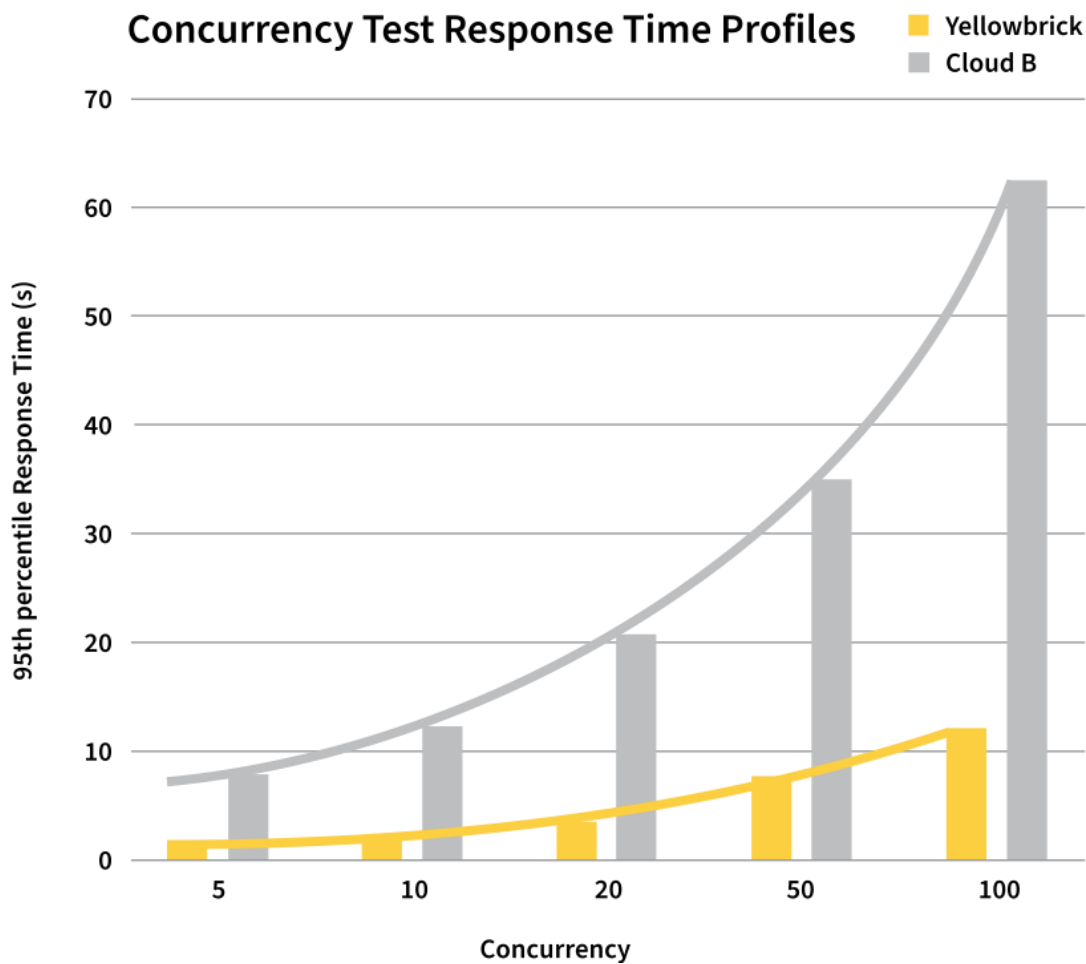


Figure 11 Yellowbrick 16-node instance concurrency test against a leading cloud vendor similarly configured cluster

Yellowbrick Data Warehouse could support 10x the number of concurrent users on a single cluster for the same user response time. Further, Yellowbrick demonstrated consistently low query response times as concurrency increased where the cloud vendor response times increased drastically as concurrency increased. When accounting for the cloud vendor ability to scale on demand, the associated costs would increase. To support a similar workload, concurrency and response time as the Yellowbrick 16-node instance, the cloud vendor would need 10 clusters, increasing cost 10x to complete the same workload in the same time duration.

4.5 Validation summary

Yellowbrick Data Warehouse and Dell EMC PowerScale F800 solutions provide an advanced platform for modern data lakes with high-speed SQL execution. This solution could demonstrate:

1. High-speed data movement into and out of the Yellowbrick Data Warehouse
2. Amazing Query Performance for standard TPC-DS queries
3. Massive SQL Scale and performance using 100 TB TPC-DS dataset

5 Conclusion

The modern data architecture for enterprises must effectively address the data deployment challenges and costs associated with the storage and consumption of data for insights. The solution presented herein combines the strengths of Dell EMC PowerScale powered by Faction multicloud Platform-as-a-Service and Yellowbrick Data Warehouse to offer enterprises both multicloud and hybrid-cloud flexibility, leading to deployment freedom, superior performance and industry-leading costs. With this reference architecture, enterprises can share and leverage data across public clouds in both an agile and secure fashion, more efficiently use cloud compute, eliminate cloud lock-in and reduce cloud egress costs.

In addition to deployment flexibility and superior price-performance, the modern data architecture for enterprises must also meet a number of tactical demands. First, it must effectively process data by using solutions like Spark and Presto for data discovery and AI/ML. Second, it must effectively support real-time streaming with the ability to scale to high message rates and large datasets. Third, it must satisfy the concurrency requirements resulting from today's business intelligence solutions.

While enterprise customers demand for more data with faster access to insights will continue to grow, traditional architectural approaches relying on commodity solutions built on virtualized instances will fall short of these stringent demands and will come at a premium price. As such, enterprise customers are best served by leveraging optimized instances and purpose-built modern data warehouse solutions like this to achieve the flexibility and best price performance for today's multicloud challenges.

A Technical support and resources

[Dell.com/support](https://dell.com/support) is focused on meeting customer needs with proven services and support.

[Storage technical documents and videos](#) provide expertise that helps to ensure customer success on Dell EMC storage platforms.

A.1 Related resources

Provide a list of documents and other assets that are referenced in the paper; include other resources that may be helpful.

- [Reimagining Traditional Data Warehouses with Yellowbrick and Dell Technologies](#)
- [Dell Technologies Cloud Storage Hybrid Disaster Recovery as a Service](#)
- [Dell technologies cloud storage for multi-cloud – powered by Faction](#)
- [Dell EMC PowerScale](#)
- [Yellowbrick](#)
- [Yellowbrick hybrid cloud data warehouse](#)