WHITE PAPER

Teradata VantageCore IntelliFlex® Design Principles

A Question of Balance

teradata.

Table of Contents

A Question of Balance: Teradata VantageCore IntelliFlex

Since Teradata's beginning, our engineering lab's goal has been to deliver balanced systems. Balanced, scalable analytics systems can deliver huge system throughput, fast performance, solid uptime, and low cost per query. It's in Teradata's cultural DNA. Every week, hundreds of Teradata engineers whiteboard plan for these attributes. Other vendors like to talk about speeds and feeds such as CPU cores, IOPS, and terabytes. But those are table stakes in today's world. Any vendor can boast those same speeds and feeds. A bag-of-parts does not guarantee business value. It's line-of-business people and data scientists who convert database queries into business value. Fast performance is the foundation of productivity benefits. It can help a data scientist find new customer prospects worth $5 million in revenue. Or it can enable running thousands of queries per hour, each one shaving $10-$1000 in costs off the bottom line. Teradata VantageCore IntelliFlex installations do that and more every day. They can do it because the system is balanced and easy to use. It's engineered to deliver business facts and insights fast, so organizations can solve important challenges.

But what exactly do we mean by "balanced systems"? A Teradata hardware performance expert invented the LH ratios in the late 1980s after the performance analysis team discovered that waiting for disk IO had zero business value. Delivering insights to the business user was—and continues to be—the primary goal.

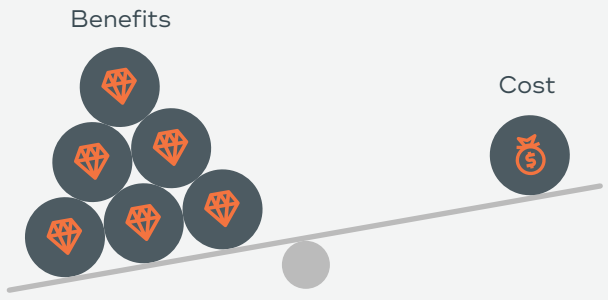

Figure 1: Providing Systems That Deliver More Value, Less Cost

teradata.

LH ratios balance compute power with disk input-output speed, saturating the Intel CPUs. Today we call this system utilization. At Teradata, we balance costs and value by squeezing every ounce of performance from disk IO, the most expensive component of the machine, and making sure there is enough CPU to meet the query needs of the workloads. Teradata experts acutely understand the art and science of achieving the perfect balance of economic and performance for our customers.

When the LH ratio capabilities first debuted, Teradata's vertical industry experts conservatively provided prospects with return on investment (ROI) estimates of 50%-100%. They were off—by a lot. In the late 1990s, a major analyst house discovered early data warehouses delivered an incredible ROI of 400%-1000%. After 20 years of solid ROI, the word "analytics" became synonymous with data warehouses, business intelligence, and decision support. Later, data science was added to analytics.

Economically Superior Systems

True data warehouses are economically superior to data marts and data lakes. They consolidate and run many use cases from the same shared data. Large pools of shared data are easier to manage, secure, and govern than tens or hundreds of data marts. As one CIO noted, "I can manage one of anything better than two of anything." Data warehouses remove duplicate data, pipelines, and systems from the cost environment. This in turn eliminates many unnecessary data transfers and time delays. Fewer systems and copies of data also reduces IT complexity. Put another way, a single data warehouse is lower cost than dozens of "my-data-marts."

Teradata has long been known for implementing the biggest databases in the world. Over 100 Teradata customers have petabyte-sized systems. Among these installations are many banks, manufacturers, retailers, telecommunications, insurance companies, and government agencies. Most are household brand names in their countries. Below is a sampling of Teradata's large on-premises deployments across various industries.

Industry	Servers	Storage
Banking	305	17 PB
Insurance	177	7 PB
Manufacturing	312	27 PB
Manufacturing	390	24 PB
Telecommunications	126	10 PB

Table 1: Example Petabyte Class Customers as of 2022

These data warehouses are 24/7 powerhouse systems supporting thousands of concurrent users with continuous loading data streams. Teradata database systems are always integrated with data lakes and cloud workloads. Many customers have three to ten Teradata systems spread across corporate divisions and geographies. Most are real data warehouses with 4,500 to 7,000 database tables. That's much more sophistication than a 20 or 30 table snowflake schema.

However, big systems aren't the only way to drive value and insights. VantageCore IntelliFlex 2.5 scales from the smallest data mart to the largest data warehouse, which means small systems can yield the same benefits as petabyte class systems. In fact, due to the powerful capabilities and ROI, a third of all Teradata installations are for small and medium businesses. It is also common for large corporations to have smaller Teradata database systems for departmental workloads.

Standard High-Volume Manufacturers

Teradata hardware systems are assembled from leading manufacturer's high-quality components. There is nothing proprietary in VantageCore IntelliFlex systems. It takes thousands of hours of painstaking research, testing, and experimentation to combine these widely available components into a balanced, high-performance system. The secret to building a powerful, high value, low-cost system is knowing how to select, combine, and balance the right "ingredients" for the best results, which is where the Teradata engineering team excels.

Teradata led the database industry in using standard high-volume (SHV) components for data warehouses in the late 1990s. SHV components include hardware like Intel and Dell servers, Nvidia InfiniBand, NetApp disk arrays, and many others. Our engineers selected SHV components for Teradata systems because of the unique way these components balance cost with performance. SHV vendors sell in bulk, which ensures volume discounts that translate into savings for customers. Dealing in volume also means SHV vendors can iterate quickly by supporting thousands of customer workloads, learning what works best and what doesn't. As a result, SHV vendors invest heavily in product quality to prevent costly warranty returns and recalls. All these factors make SHV components nearly impossible to beat on price for performance and quality, which is how they earned their place in Teradata systems.

However, it takes tremendous skill to select the best SHV components and balance them in the final configuration. In any given year, Intel and Dell both have up to 50 server variations based on memory sizes, CPU counts, internal storage, networking, and other options. NetApp and other storage vendors have hundreds of variations for disk arrays, storage devices, and cabling. Only a select few of these components map well to data warehouse workloads.

Few SHV vendors or system integrators grasp the fundamentals of high throughput data warehouses. That's where Teradata Engineering excels. Our engineers spend considerable effort sifting through available components, then assembling and stress testing them to ensure every new system is better than the prior generation. Every step includes balancing between price and performance to match customer workloads.

This is just one part of designing for high levels (80-95%) of analytic server utilization. High utilizations mean customers can get maximum value from their investment. These choices intentionally prevent the need for excess servers to meet service level goals.

Deep Hardware and Software Allies

Teradata Engineering is an insider when it comes to working with SHV hardware suppliers. Teradata has strong working relationships with Intel®, Dell®, InfiniBand®, NetApp®, Nvidia®, and other SHVs. We build strong relationships with these vendors by breaking their products. Sound strange? Most enterprise products are designed for fast access transactional workloads—think online transactional processing (OLTP) and software as a service (SaaS). These workloads are the equivalent of motorcycles zipping around major freeways and highways. Analytic workloads tend to be more like heavy hauler trucks carrying huge loads of data and algorithms. Teradata systems run both (i.e., mixed workloads) which more than doubles hardware stress.

Teradata Engineering takes stress testing new product samples to the limits. Our team applies skills, data, and software that SHV suppliers just don't have. That's because our customers push Teradata database systems to extreme limits. Data warehouse workloads are throughput oriented, not simple transaction processing of eight kilobytes. Scanning, sorting, and joining two 500TB tables is not an OLTP or SaaS profile. Teradata Engineering shares analytic workload profiles and stress test results with SHV suppliers which helps them improve their products.

VantageCore IntelliFlex designers also build strong relationships with Cloud Service Providers (CSPs). One CSP shut down Teradata VantageCloud, believing they detected a denial-of-service attack because VantageCloud was driving all the disks at maximum utilization. But it's common for Teradata SQL to run full table scans by driving extreme levels of throughput. Teradata's parallelism can fetch data from all the disks drives on all servers simultaneously. The goal is to get the answer to the business user faster. The CSP software flagged the behavior as a serious breach, not knowing that's what Teradata SQL customers need. That CSP now has Teradata engineers in their "canary in the coal mine" process so they can learn about high-capacity analytic stress points. These kinds of working relationships distinguish Teradata Engineering from competitors. With patience and respect, we earn our time at the whiteboard with cloud and SHV leaders.

teradata.

Figure 3. Teradata Intelliflex™ 2.5

Teradata VantageCore IntelliFlex

VantageCore IntelliFlex 2.5 begins as two primary configurations: balanced and high-CPU. Balanced systems maintain the LH ratios between huge amounts of input-output per second (IOPS) and compute capacity. This is best for most data warehouses. The high-CPU systems have less need for cold data storage per server. This is popular for data marts, departmental databases, and data science workloads.

The VantageCore IntelliFlex cabinet is a 42U SHV rack. VantageCore IntelliFlex 2.5 uses Intel 1U servers. The system uses only solid-state storage. VantageCore IntelliFlex comes in various configurations. The most popular are six servers or twelve servers plus one hot stand-by node: 6+1 or 12+1, also known as cliques. Smaller configurations as low as 2+1 are available for production or test and development purposes. Every cabinet contains two ethernet switches. The ethernets provide system management and data transfers to business users. Most of the cabinet space is taken by disk arrays and storage devices. A few cabinets contain InfiniBand switches which send and receive BYNET messages and data blocks. BYNET protocol provides bi-directional broadcast, multicast, point-to-point communication, and merge functions.

Teradata released the 4600 system in 2010[1]. It was the world's first all-solid-state drive (SSD) analytic database system. Because of SSD costs, hybrid hard disk drive (HDD) and SSD systems became the norm in the prior decade. Today, to match customer expectations, SSDs have increased capacity and lower prices. VantageCore IntelliFlex 2.5 uses NetApp disk arrays populated with either 6.4TB or 12.8TB SSDs. All are standard high volume (SHV) components.There are 100TB SSDs available today that are five times the cost per terabyte versus Teradata VantageCore IntelliFlex SSDs. When high-quality 100TB SSDs fall far enough in price, they will be included in Teradata systems.

A few specialty servers are common in Teradata configurations. The Parsing Engine (PE) node supports huge user communities (i.e., more than 5,000 concurrent users). PE nodes also help with frequent, extra-complex SQL optimization. FICON channel servers load transactions from IBM mainframes into the data warehouse. A memory dump server supports security lockdown policies for Teradata Customer Support secure access. Teradata Multi-purpose Servers (TMS) run database system administration tasks such as Viewpoint and DSA (backup and restore).

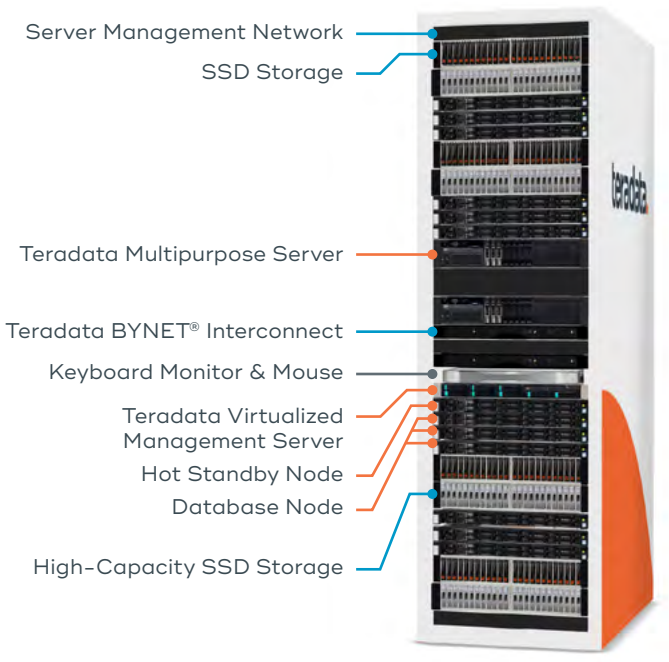

Figure 2. Anatomy of VantageCore IntelliFlex

1 Teradata adds flash, denser packaging https://www.theregister.com/2011/04/11/teradata_edw_6650_6680/

teradata.

BYNET V6 and InfiniBand

Teradata's VantageCore IntelliFlex 2.5 platform utilizes a converged fabric architecture. It is built with NVidia InfiniBand™ hardware. It carries both BYNET™ protocol messages and disk array storage messages. BYNET protocol passes Teradata Vantage™ commands and data between server nodes. InfiniBand also connects disk storage arrays to compute nodes, sending commands and data blocks using low overhead RDMA. Thus, InfiniBand acts like two networks while separating compute servers from storage. This was another design decision driven by standard high-volume (SHV) technology price and performance.

Nvidia InfiniBand and BYNET protocol are both non-blocking message passing subsystems. Non-blocking means messages always go through in both directions. Messages are never stalled by other users or messages.

In contrast, many ethernet devices still contain vestiges of message blocking. That means one sender can lock out all others for periods of time. InfiniBand accelerates results to the business user. InfiniBand EDR switches provide data traffic between cliques. This assures high performance for data loading and extreme table joins across the system.

Telemetry Data

Teradata Vantage has always collected hardware and software statistics. They are used for customer support, repairs, billing, and forecasting workload growth. Most of this data comes from the resource usage (ResUsage) logfile statistics.

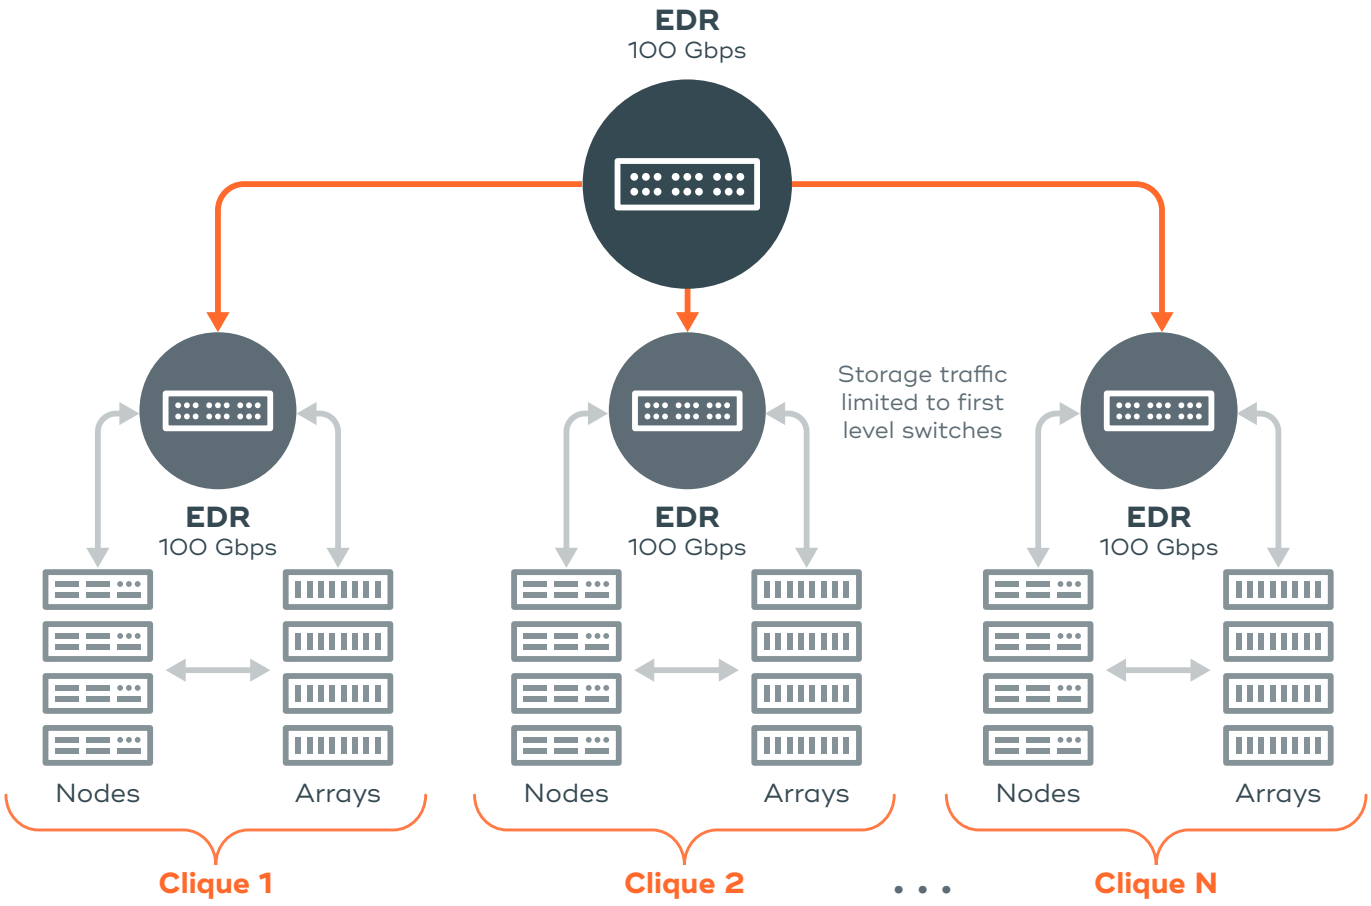

Figure 4. BYNET V6 InfiniBand Fabric, Servers, and Storage Arrays

teradata.

ResUsage data enables:

- System performance monitoring
- Component and system performance measurements
- Planning installations, upgrades, and migrations
- Detecting performance degradations and improvements
- Finding system bottlenecks, parallel inefficiencies, and congestion

Many database administrators review ResUsage to diagnose business user analytics problems. Teradata Customer Support analyzes ResUsage data to resolve hardware and software problems. Teradata engineers use ResUsage to profile workload trends for next generation systems. The data also reveals workload differences across industries, geographies, and calendars.

Benchmarking Real Customer Workloads

Teradata Engineering uses TPC benchmarks and a few dozen internal benchmarks to test new components. Among many hardware performance tests is the active data warehouse concurrent workloads test. It includes severely complex queries, thousands of sub-second queries, and continuous terabyte data load batch jobs. Most databases cannot support any two of these workloads.

Any new hardware device or software change can unexpectedly upset the system balance. One way to detect this is to run real-world workloads. Retailers, insurers, services, and other industries have donated complete workloads for benchmark testing. Donations include SQL queries, batch jobs, workload schedules, and data. Those who donate benefit by knowing their workloads are performance tested for them before they get a new Teradata Vantage system. Some benchmarks are part of Teradata Engineering's daily DevOps cycles. Every new device and software change gets run through all these grinders hundreds of times. Hardware quality assurance performance tests run every day for a year before a system is ready to ship.

How Teradata Designs New Systems

As suppliers release new components, Teradata's engineers apply strenuous, component-level bench tests to quantify performance and reliability variables. Whether the component is a server, a CPU, memory, or storage, it goes through extensive scrutiny and testing. The items that survive are candidates for the next generation of on-premises systems. Next comes prototype systems assembly and stress testing. Throughout these efforts, a lot of information is exchanged with SHV engineers. We learn from each other constantly.

Then comes the hard part: balancing all the new hardware, software, costs, and performance ratios. Imagine a master chef at a Michelin 3-star restaurant. You might shop in the same grocery store as they do, watch their videos over and over, and follow their recipe exactly. But despite your best efforts, your results lack something: talent and skills from decades of practice with other master chefs. Only the master chef delivers exceptional cuisine.

Teradata Engineering has that "master chef" talent and the skill from decades of building massively parallel data warehouse machines. We've learned from thousands of customer workloads what is truly needed. These unique skills and experience squeeze enormous performance and value from SHV components. Our engineers get a 3% performance boost here, 5% there, and so on until they've created the best performing system possible. There are insights learned every year that accumulate in the designs, the culture, and the engineers. Competitors can duplicate the parts but not the process or the expertise.

Which brings us to the last SHV vendor: Flex Manufacturing. Teradata designs our hardware systems, but Flex manufactures them. This multinational electronics contract manufacturer is the third largest in the world. They build the systems, test, and ship them. Every system is different, customized to the customer's goals, yet the quality and turn times are exceptional, and the build costs are low. Flex managed the supply chain with expert planning through the pandemic.

How Teradata Delivers New Systems

Does a salesperson configure a customer's next proposed system? We hope not, although it's common among some vendors. Teradata Engineering has a secret weapon to ensure each customers' system proposal exceeds their expectations: Global Sales Solutions (GSS). This proven process provides the best configuration choices for that customer's workloads. It begins with the ResUsage data, the profile for what is happening on customers' current system. ResUsage data reveals peak utilization, surges, seasonality, and other influences. To this, GSS adds the organization's plans. Customers may want to shorten the hours spent in ETL data loading and cleansing. Or they may expect to add 500 users between sales and marketing. GSS can segment out that ResUsage data and predict the changes in those workloads. That translates into a more accurate hardware configuration. Notice how GSS uses MPP analytics to configure MPP systems.

Migrating from legacy platforms to completely new architectures brings new challenges. Imagine switching from hard disks to SSDs. Everything just goes faster, right? It will if more CPU cores are added. Otherwise, SSD input-output might sit in a queue, limiting the value of the device. Plus, when data compression is added, more CPU capacity is needed. The hardest configurations are moving clients from two generations back into the newest machines. Many old expectations are void and new ones must begin. To address this challenge, our Solution Platform Adviser tool models the current system and maps it to the future system goals.

What if you're a new customer with no ResUsage data to start with? Ideally, your organization should do a proof-of-concept benchmark on an VantageCore IntelliFlex system. It's the best way to calculate accurate performance expectations. To streamline the process, GSS acts as the liaison between your organization, Teradata engineers, and pricing experts. GSS experts help answer the complex "why" questions. The final balance of hardware, software, and customer requirements is backed by data analysis as much as possible.

Teradata Vantage

Portability across hardware and software platforms has been the definition of open software for three decades. Teradata Vantage runs production on Teradata hardware systems, Amazon Web Services (AWS), Microsoft Azure, Google Cloud, and VMware private clouds. Our clients have the open platforms they need to solve their most important business challenges. With these choices comes the need for sophisticated data movement tools between on-premises data warehouses and cloud platforms. Teradata QueryGrid accesses or exchanges data in parallel between Teradata systems and cloud platforms. Data can be structured or unstructured. It can reside in databases or object storage. Teradata Data Mover transfers bulk data between on-premises and cloud systems. These data sharing tools enable cloud bursting: connecting to a public cloud for additional computing resources without service interruption. Teradata Vantage is designed to support both hybrid cloud and multi-cloud deployments.

Teradata Vantage consists of a parallel relational database plus an analytic algorithms subsystem called ClearScape Analytics™. Teradata Vantage Analytics Database runs on-premises, public clouds, and private clouds. Vantage supports terabytes-to-petabytes scale. No special effort by business analysts or data scientists is necessary. Business users don't even need to know SQL. They simply click through with BI tools such as Tableau, Microsoft PowerBI, or Qlik. Data scientists can run more than a hundred data science algorithms in parallel or they can bring their own models. Dozens of popular third-party tools also connect to Teradata Vantage for data cleansing, governance, loading, and other tasks. Numerous programming languages also run inside the parallel engine such as Java, R, and Python.

Integrated Hardware and Software

Hardware and software are interdependent, symbiotic technologies. Many hardware and software changes affect the balance between CPUs, memory, and storage. The existence of two subsystems—Teradata Optimizer and Teradata Active System Management—is pivotal to ensure optimum teamwork between hardware and software.

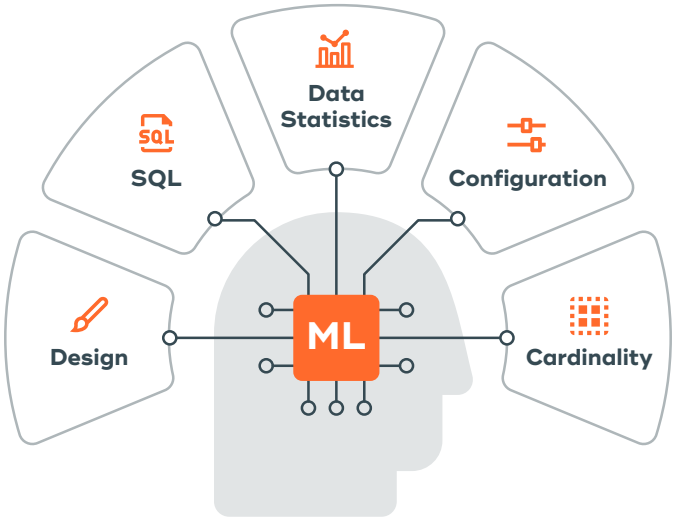

Figure 5. Teradata Differentiator: The Best Cost Based SQL Optimizer

Teradata Optimizer

Optimizers are a branch of machine learning. Teradata Optimizer has 40 years of investment behind it—a formidable barrier to entry for competitors. Even being 10 years behind is insurmountable. Consider that a 10-table relational join has 10 factorial (10!) join plans, which equates to about 3.6 million plans to choose from. Teradata Optimizer easily handles 20 to 50 table joins every day. The Teradata Optimizer uses the Teradata database configuration plus hardware performance coefficients to calculate the best SQL plan possible. Competitor databases, data lakes, and object storage have no or negligible query optimization, which is why they are limited to simple, star-schema joins that are manually optimized by programmers. Every generation of Teradata hardware comes with new hardware coefficients fed into the Optimizer's algorithms.

Because of the Optimizer's design, less hardware is needed for a workload. Teradata Optimizer works smarter, not harder, in subseconds on every task.

Teradata Active System Management

In 2001, Teradata was the first to load streaming data into a data warehouse. Customers loved it. Soon, they demanded continuous loading all day long. Meanwhile, hundreds of concurrent OLTP-tactical user queries per second were being added to the data warehouse. Like a four-lane highway with big trucks, motorcycles, and family cars, a five-kilometer traffic jam emerged inside the system. Teradata Active System Management (TASM) emerged in 2002 to manage this "traffic jam." To organize traffic lanes by priority, TASM looks deep into the Optimizer's plan. It uses hardware coefficients like estimated processing time, estimated row counts, memory use, and workload type (i.e., tactical, loading, and complex query) to manage tasks in priority lanes. TASM is the traffic director, the run-time optimizer ensuring the best throughput. The on-ramps are organized, and the heavy hauler work must stay in its lane while the family cars zip through. Less hardware is needed for a workload than competing databases without high-quality workload management.

As a result of the Teradata Optimizer and Teradata Active System Management subsystems, Teradata system utilization is unmatched for analytic workloads. One customer database administrator (DBA) told our engineers "Your Viewpoint tool blinks red when the system is running over 95% utilization, but we wish it didn't. If the machine is racing at 95-100%, we don't need an alarm—we're getting our money's worth!" Teradata engineers removed the alarm.

It's common for VantageCore IntelliFlex to run 60-85% utilization all day with 90-95% surges in daily and seasonal workloads. Many competitor systems rarely hit 25%-35% utilization, which means money is spent but most of it sits idle.

teradata.

High Availability and Disaster Recovery

A large portion of all hardware and software design is aimed at maintaining uptime and recovery.
All the wonderful performance is forgotten when the system is down. For the business analysts, availability is simple: either the entire system is up or its down. For hardware engineers, it's a career.

High Availability Uptime

Hardware vendors strive to provide extensive high availability to minimize downtime. There are hot-swappable components so Teradata Customer Support can replace devices without shutting down the system. Device redundancy enables the hardware to self-heal by switching out failing components automatically. SHV manufacturers also build in dozens of redundant components that heal automatically. Thus, Teradata VantageCore IntelliFlex commonly operates continuously at 99.99% uptime per year or better. Well-managed installations achieve 99.999% uptime. That's roughly one hour of unplanned downtime per year.

With Teradata databases, a group of physical servers with access to the same disk arrays is a clique. Teradata databases use hot standby nodes (HSNs) in cliques to ensure performance continuity. HSNs take over database tasks whenever a database node fails. HSNs eliminate performance degradation and system restarts during node failures. The use of InfiniBand connectivity to disk arrays allows any server in the clique access to any disk array. This enables up to 12 servers and one hot standby server in a clique. Performance is assured and the newer hardware configuration costs are much lower. One CSP likes Teradata's hot standby architecture and has added cliques to their instances. Another CSP is planning to do the same.

Scheduled downtime during upgrades has been dramatically improved. Previously, adding servers and storage could cost 5 to 24 hours of scheduled downtime. Multiple Hash Maps (MHM) software changed this into a background process that runs when the DBA chooses. MHM reconfiguration runs concurrent with production workloads. Customers upgrading VantageCore IntelliFlex report outstanding satisfaction, faster results, and 10–15 minutes of downtime.

Hot Swap Features	
• BYNET switches • LAN switches • Disk array controllers • Power supplies • Storage devices	• Power supplies • Array controllers
Device Redundancy	
• Hot standby node • Dual AC input • Dual BYNET • Dual network cards • CMIC failover • Redundant disk access • Raid storage	• Power supplies • Multiple fans • Fallback on all tables

Table 2. Self-Healing Hardware

Disaster Recovery

Investments in disaster recovery plans are an insurance policy against severe downtime and CIO/DBA unemployment. Yet, like most insurance policies, some organizations let them lapse or spend their budget elsewhere. And like most insurance policies, cost optimization is the better choice.

To reduce disaster recovery costs, Teradata offers Business Continuity Manager software (BCM). BCM uses a hot standby Teradata system with only the database tables required for continuity. It need not be a duplicate configuration with 100% mirrored data. Updates to the primary production system are sent to the hot standby system, which uses one of three synchronization methods.

teradata.

- SQL multicast broadcasts SQL to the standby system to ensure consistency and availability.
- Change data multicast runs SQL on the primary system, then ships changed data to the hot standby system.
- Content replication copies tables to hot standby systems using QueryGrid. This provides bulk data loading and self-healing tables.

BCM runs on a second Teradata hardware system or in Amazon Web Services (AWS) instances.

On-premises or in the cloud?

While database management system (DBMS) growth in cloud platforms is surging,on-premises database software continues to see steady gains, including $2.2 billion of revenue growth in 2021, per Gartner.[2] According to technical journalist Joe McKendrick, "we still live in an on-premises world…Half of the [IT] executives surveyed said their organization's total IT environment is 'some on cloud but mostly on-premises.'"[3] Look no further than industry cloud giants AWS and Google Cloud for proof that on-premises has plenty of growth ahead: both CSPs are selling on-premises hardware to gain market share.[4]

What factors should organizations consider when choosing cloud or on-premises? Workload type, security requirements, uptime needs, performance speeds, and financial governance needs are crucial starting points.

Workloads

Some workloads fit best in the cloud while others are better suited for on-premises. Cloud migrations for large, mission-critical workloads are complex and expensive. Data gravity and application sophistication also play a key role in large workload placement.

Security Requirements

Security is another crucial factor to consider. It can be a cloud inhibitor for organizations in highly regulated industries, such as financial services, healthcare, and government. However, for other industries, CSP security is often better than their on-premises data center.

Uptime Needs

Hardware uptime is a benefit of on-premises systems. AWS, Azure, and Google Cloud only offer 99.9% uptime in their contracts. VantageCore IntelliFlex routinely delivers 99.99% uptime and often achieves even better results.

Performance

Cloud performance is often slower than Teradata's on-premises solutions, as explained in Teradata's white paper, Driving Forces for Taking Analytics into the Cloud.[5] Slow cloud performance can be caused by commodity hardware, hypervisors, and intentional CSP throttling.

Financial Governance

Most runaway cloud costs are compute instance costs, not storage costs. FinOps software recently emerged to provide cloud cost governance. Gartner explains reasons[6] Why Cloud Budgets Don't Stay in Check. Three are relevant here: ungoverned costs, too much production headroom, and wrong sizing production.

2 DBMS Market Transformation 2021: On-Premises DBMS Revenue Grows, Gartner, Merv Adrian, June 16, 2022. https://blogs.gartner.com/merv-adrian/2022/06/16/dbms-market-transformation-2021-on-premises-dbms-revenue-grows/

3 We still live in an on-premises world for now, Joe McKendrick, April 23, 2022. https://www.zdnet.com/article/cloud-computing-has-all-the-momentum-but-we-still-live-in-an-on-premises-world-for-now/

4 AWS fights on-prem hardware vendors, June 2, 2022 https://www.theregister.com/2022/06/02/aws_launches_two_new_challenges/; Google Cloud brings AI into the datacenter, June 16, 2022 https://siliconangle.com/2022/06/16/google-expands-distributed-cloud-platform-anthos-premises/

5 Driving Forces for Taking Analytics into the Cloud, Brobst, Elliot, Freund, Le, Suryamurthy; see Figure 1; Teradata, April 22, 2022. https://www.teradata.com/Resources/White-Papers/Driving-Forces-for-Taking-Analytics-into-the-Cloud

6 Gartner, Why Cloud Budgets Don't Stay in Check – And How to Make Sure Yours Do, January 28, 2022, https://www.gartner.com/en/articles/why-cloud-budgets-don-t-stay-in-check-and-how-to-make-sure-yours-do

teradata.

In contrast, VantageCore IntelliFlex expansions are well governed by the customer's executives, procurement staff, and DBAs. Production headroom is optimized by TASM workload manager while delivering peak performance with shared servers. Teradata's GSS configurations are meticulously sized to match current production workloads and near future growth

Teradata VantageCore IntelliFlex Pricing

For many years, pundits and competitors complained about storage "cost-per-terabyte." But costs and price perceptions changed as "big data" emerged, easy cloud provisioning arrived, and burgeoning departmental technology budgets became widespread.

Teradata VantageCore IntelliFlex systems support multiple forms of cloud-like elasticity while providing organizations with flexible pricing that enables control over costs. Teradata's TCore pricing has been lauded by analysts as a true pay-for-performance pricing model. Organizations can tailor their configuration to exactly what they need now, and easily grow by activating only what they need via the capacity on demand (CoD) function. This enables rapid expansions of computing power or storage capacity in budget friendly increments. There are several methods of managing CoD.

The Elastic Performance on Demand (EPOD) program gives customers the opportunity to buy their hardware resources up front at a reduced price. To meet the dynamic needs of the business, pay-per-use pricing is available for immediate access to increased performance.

Another COD option is Elastic TCore which enables VantageCore IntelliFlex installations to dynamically turn CPU cores on or off in every node and adjust IO throughput levels independently, as needed by the active workload demand, to deliver on a software-defined TCore.

Summary

Teradata engineers have decades of experience deploying analytic workloads via standard high-volume components. It's a rare set of skills, as many organizations have discovered. The balanced technologies inside VantageCore IntelliFlex deliver a cost-effective, scalable, and highly available platform—one that's capable of boosting productivity for tens or thousands of users while empowering them to deliver maximum business value.

About Teradata

Teradata is the connected multi-cloud data platform for enterprise analytics company. Our enterprise analytics solve business challenges from start to scale. Only Teradata gives you the flexibility to handle the massive and mixed data workloads of the future, today. Learn more at **Teradata.com**.

17095 Via Del Campo, San Diego, CA 92127 Teradata.com

teradata.